A Core-Spectral Theory of Differential Hull Clustering in Nibble SPNs

Weipeng Yu**

Improved differential distinguishers for CRAFT

Abstract. Differential hull clustering, i.e., the phenomenon that the differential probability of a cipher can exceed its best single trail by many orders of magnitude, has long been observed, but a structural explanation of *where* the clustering comes from has been missing. This paper gives such an explanation for four-bit nibble SPNs: the core-restricted clustering gap is governed by the *recurrent cores* of the round function, i.e. the irreducible recurrent components of the differential transition. For a core $\mathcal{C}$ with k active S-boxes and weighted adjacency matrix $T_{\mathcal{C}}$ of spectral radius ρ, the clustering gap obeys the asymptotic law $\mathrm{gap}_{\mathcal{C}}(R) = R\,(\mu(\mathcal{C}) + \log_2 \rho) + O(1)$, where $\mu(\mathcal{C})$ is the minimum-mean-cycle weight. A central consequence is that the branch number, the standard diffusion metric, is *not* sufficient to bound clustering: it controls only the single-trail weight, not the trail multiplicity measured by ρ. We make the energy/entropy split rigorous through the variational principle, which identifies $\log_2 \rho$ as an entropy term, and we show that for a periodic core the $O(1)$ transient is the discrete Fourier transform of its cyclic classes; two branch-2 layers differ by 12.3 bits of gap while a branch-4 layer clusters less than both. We construct *new differential distinguishers* on CRAFT in the single-tweak model for 11 to 14 rounds; at each of these round counts, the probability of our distinguisher is, to our knowledge, the highest reported to date. The 11-round distinguisher reduces the data complexity from $2^{49.79}$ to $2^{49.559}$. We also report, to our knowledge, the first single-tweak differentials for 1 to 8 rounds. The theory is implemented in an open, C++-accelerated tool (`coretool`) that automates core search, spectral analysis and finite-round validation on CRAFT [1], GIFT-64 [7], PRESENT-80 [8], SKINNY-64 [9] and controlled ablations, yielding actionable design guidelines for hull-resistant linear layers.

Keywords. differential cryptanalysis, hull clustering, subshift of finite type, transfer operator, topological pressure, branch number, nibble SPN.

1 Introduction

Differential cryptanalysis of a block cipher proceeds from a chosen input difference Δ_{in}; the adversary exploits an output difference Δ_{out} whose *differential probability* is high enough to distinguish the cipher from a random permutation. The differential probability is not the probability of a single trail: it is the sum over *all* trails realizing the same pair $(\Delta_{\mathrm{in}}, \Delta_{\mathrm{out}})$. The gap between this sum and the best single trail is the *hull-clustering effect*, and in modern nibble SPNs it is routinely the dominant term of the security margin. Ignoring it leads to systematically over-optimistic security

** Correspondence: 3232010811123@glut.edu.cn.

claims. This paper attributes the clustering to a concrete structural feature of the linear layer (its *recurrent differential cores*) and proves that the branch number, the standard diffusion metric, does *not* control it.

The phenomenon is concrete and significant already for CRAFT, a tweakable nibble SPN proposed at FSE 2019 for efficient protection against differential fault analysis. Hadipour et al. [1] showed that the best eleven-round trail of CRAFT has probability 2^{-80} (40 active S-boxes), yet the differential probability of the optimal input–output pair is $2^{-49.79}$: the hull contributes roughly 2^{30} worth of additional trails. This is a 30-bit correction to the naive trail-based estimate, and it is the kind of margin that decides whether an attack or a security claim is meaningful. In Section 5, our tool locates a new 11-round differential with probability $2^{-49.559}$, higher than the $2^{-49.79}$ reported by Hadipour et al., and extends it to 12 to 14 rounds.

The linear-algebraic treatment of trail clustering has a long history. Biham and Shamir's differential cryptanalysis [13] established the trail-based framework; Knudsen's truncated differentials [11] already group trails by their active positions, anticipating the pattern-class view; Lai, Massey and Murphy [12] formalised the Markov-cipher model in which the differential transition matrix is the fundamental object; and Daemen's correlation matrices [10] express the average linear-hull effect through the spectral radius of a correlation matrix, and the same idea transfers to differentials through the DDT transition matrix; it is the backbone of several differential-hull estimates. These matrix methods, however, act globally on the state space and do not expose which internal structure of the linear layer is responsible for the clustering.

A precise account of this clustering, however, has been elusive. Hadipour et al. enumerate trails within the *optimal-pattern class*, i.e. they fix the minimal-active-S-box pattern and count the trails realizing it; this captures the clustering *within* the optimal pattern but says nothing about the contribution of patterns with more active S-boxes. Liu et al. [2] count trails by iterating correlation matrices, which yields trail *counts* but is manual and hard to scale to the full pattern space. Cui et al. [3] introduce *congruent differential clusters*—trails whose active S-boxes share equal output differences—and obtain exact probabilities for that restricted, highly structured subclass, but not for the full trail space. Beyond nibble SPNs, the clustering effect has been studied for ARX ciphers and with general-purpose search tools: Leurent et al. [4] compute the full probability distribution of a window-based trail class for Simon and Simeck, Biryukov et al. [5] automate the truncation and clustering of ARX trails, and Sakamoto et al. [6] parallelise SAT to find the best differential under clustering. All of these compute *how much* clustering occurs; none of them attributes it to a structural feature of the linear layer.

Table 1 summarises the comparison.

Method	scope	setup	result	structural cause
Hadipour et al.	optimal pattern	MILP + manual	exact (class)	none
Liu et al.	opt. + subopt.	manual matrices	counts, approx.	none
Cui et al.	congruent clusters	analytic	exact (subclass)	partial
Leurent et al.	ARX window class	analytic	exact (class)	none
Biryukov et al.	ARX truncation	automated	exact (truncated)	none
Sakamoto et al.	generic	SAT (parallel)	best diff.	none
This work	low-weight cores	automated	exact (restricted) + spectral	recurrent-core spectrum

Table 1: Comparison with prior differential-clustering work. The last column indicates whether the method attributes clustering to a structural feature of the linear layer.

Hadipour et al. enumerate trails inside a fixed pattern and therefore cannot say *why* that pattern clusters; Liu et al.'s correlation-matrix iteration yields a spectral rate but no interpretation of which

part of the linear layer produces it; Cui et al.'s congruent clusters identify a subclass of trails but not a feature of the linear layer itself. This work attributes the clustering to a concrete object—the recurrent differential core of the linear layer, and in particular its fixed points—so that the spectral radius is not merely a number but the signature of a design property, turning the theory into a design criterion rather than a numerical tool.

Existing matrix-based works can observe hull-clustering numerically, yet they cannot answer which internal linear-layer structure produces such clustering. Moreover, prior methods rely on manual setup or only handle special sub-classes. We propose the *recurrent-core decomposition*, which attributes clustering to recurrent cores of the linear layer, and prove that the asymptotic core-restricted gap obeys $\mathrm{gap}_{\mathcal{C}}(R) = R\,(\mu(\mathcal{C}) + \log_2 \rho) + O(1)$, via the Ruelle–Perron–Frobenius theory of the associated transfer operator. We derive a branch-number-irrelevance corollary and two structural lemmas, provide `coretool`, a C++-accelerated open tool for core search and spectral analysis, validate the theory on CRAFT, GIFT-64, PRESENT-80, SKINNY-64 and controlled ablations, and distill three design guidelines.

Scope and claims. The claims of this paper are structural and asymptotic. We prove that the asymptotic core-restricted hull-clustering gap of a nibble SPN is governed by the linear layer's recurrent core, through the spectral slope $\mu(\mathcal{C}) + \log_2 \rho$, and that the branch number does not bound it; we provide an automated tool that finds low-weight recurrent cores and validates the prediction, and we derive design guidelines. The framework is stated for four-bit nibble SPNs, and the core search is not a complete full-space search; these boundaries are made precise in Section 6.

Organisation. Section 2 fixes notation. Section 3 develops the theory and states the theorems and lemmas. Section 4 describes the algorithm and its implementation. Section 5 reports the experiments. Section 6 discusses limitations and relates the framework to prior work. Section 7 states the design guidelines, and Section 8 concludes with open problems.

2 Preliminaries and Notation

We consider a nibble SPN with $n = 16$ four-bit cells (all ciphers in this paper are 64-bit with 16 nibbles). The S-box $S\colon \mathbb{F}_2^4 \to \mathbb{F}_2^4$ is bijective. Its *differential multiplicities* are the integers

$$m[a][b] \;=\; \big|\{\, x : S(x) \oplus S(x \oplus a) = b \,\}\big|, \tag{1}$$

with $0 \le m[a][b] \le 16$ and $m[0][0] = 16$, and $m[a][b] = 0$ whenever $a \neq 0$ and $b = 0$ (bijectivity of S). No probability normalisation enters the theory; the normalised entry $D[a][b] = m[a][b]/16$ is used only when a probability statement is quoted. We write $w_{\min}$ for the minimal nonzero S-box weight,

$$w_{\min} \;=\; -\log_2\Big(\max_{a\neq 0,\, b} m[a][b]/16\Big), \tag{2}$$

that is, the weight of the most probable single S-box transition. The number of active S-boxes must be distinguished from the trail weight: for a non-uniform S-box the two do not coincide. Minimising the active count and then assigning each active S-box its most probable transition is not in general the lightest trail. The classical example is GIFT-64, whose nine-round optimal trail has weight 42.0 (twelve 4/16 and six 2/16 transitions), strictly below the value obtained by minimising the active count alone. The weight-aware analysis of this paper is therefore the appropriate object, and $w_{\min}$ is the weight of the single most probable transition rather than a uniform weight.

For a 64-bit difference Δ written as sixteen nibbles $\Delta = (\Delta_0, \ldots, \Delta_{15})$, the *support* of Δ is the set of indices i with $\Delta_i \neq 0$, and its *active-S-box count* is $|\mathrm{supp}(\Delta)|$. A support is exactly

the *truncated difference* of Knudsen [11]: it records which S-boxes are active while discarding their actual difference values. Throughout the paper we call a support a *mask* when it is used as a search constraint.

The linear layer is either of *cell type* or of *bit-permutation type*. For cell-type layers (CRAFT, SKINNY, and the ablations of Section 5) it is $L = P \circ M$, where M is an invertible $\mathbb{F}_2$-linear map acting column-wise on the nibbles and P is a cell permutation. For bit-permutation layers (GIFT-64 and PRESENT-80) the linear layer is a permutation of the 64 bits, which is not of the form $P \circ M$. The core-spectral theory of Section 3 applies to both types, but the core search of Section 4 handles them differently (mask pre-filtering for cell type, active-class closure for bit-permutation type). Its *branch number* is

$$\mathcal{B} \;=\; \min_{\Delta \neq 0} \big(\mathrm{wt}(\Delta) + \mathrm{wt}(L(\Delta)) \big), \tag{3}$$

where wt denotes the number of nonzero nibbles. The round differential map F_D applies the S-box layer and the linear layer L in an order fixed by each cipher: for CRAFT the linear layer is applied first and the S-box layer second, so $F_D = S \circ L$ with $L = P \circ M$ (MixColumns, then the cell permutation, then the S-box layer); for GIFT-64, PRESENT-80 and SKINNY the S-box layer is applied first and the linear layer second, so $F_D = L \circ S$. We write $F_D(\Delta)$ for the one-round image of a difference under the differential transition.

A *differential trail* is a sequence of cell differences $(\Delta^{(0)}, \Delta^{(1)}, \ldots, \Delta^{(R)})$ with $\Delta^{(r+1)}$ reachable from $\Delta^{(r)}$ under one round. The one-round differential transition is a *weighted directed graph* G_D on the 2^{64} difference states: there is an edge $\Delta \to \Delta'$ of weight $m(\Delta, \Delta') = \prod_{\text{active S-boxes}} m[\Delta_i][\Delta'_i]$ whenever the transition is possible, and the weighted adjacency matrix of G_D is denoted A. A is nonnegative but is *not* assumed row stochastic. The *R-round path-sum* from α to β is

$$(A^R)[\alpha, \beta] \;=\; \sum_{\substack{\text{trails } \alpha \to \beta \\ \text{of length } R}} \prod_{r=0}^{R-1} m\big(\Delta^{(r)}, \Delta^{(r+1)}\big), \tag{4}$$

a purely combinatorial sum of nonnegative integers. Its normalised form (the product of the $D[\cdot][\cdot]$ entries along each trail, summed over all trails) is the standard DDT-product differential of cryptanalysis; equality of this normalised sum with the exact cipher differential $\frac{1}{2^{64}} |\{\, x : F^R(x) \oplus F^R(x \oplus \alpha) = \beta \,\}|$ holds under the standard stochastic-equivalence assumption [12], and is *not* needed for the clustering mechanism proved below. The reason an assumption is needed at all is that the sequence of differences alone is *not* a Markov chain: the one-round transition $\Delta \to \Delta'$ is exact (its probability is the DDT product, uniformly averaged over the intermediate value x), but for several rounds the next difference $\Delta^{(r+1)} = F(x_r) \oplus F(x_r \oplus \Delta^{(r)})$ depends on the actual value x_r as well as on $\Delta^{(r)}$. The process tracking the pair $(x_r, \Delta^{(r)})$ is Markov but has 2^{128} states (the full cipher), and the stochastic-equivalence assumption is exactly the statement that, for the purpose of the multi-round differential, the intermediate values can be treated as independent and uniform, so that the matrix product equals the exact probability. A trail's *weight* is the negative binary logarithm of the product, along the trail, of the normalised entries $D[\Delta_i^{(r)}][\Delta_i^{(r+1)}]$. The *differential probability* of a pair $(\Delta_{\text{in}}, \Delta_{\text{out}})$ is the sum of the normalised trail probabilities realising it. For a given pair, the *clustering gap* is

$$\mathrm{gap} \;=\; -\log_2 p_{\text{trail}}^{\max} \;-\; \big(-\log_2 p_{\text{diff}}\big), \tag{5}$$

where $p_{\text{trail}}^{\max}$ is the probability of the best single trail and p_{diff} is the differential probability. For a recurrent core $\mathcal{C}$ we also define the *core-level gap*

$$\mathrm{gap}_{\mathcal{C}}(R) \;=\; -\log_2 p_{\text{trail}}^{\max}(\mathcal{C}, R) \;-\; \Big(-\log_2 \max_{i,j} (T_{\mathcal{C}}^R)_{ij}\Big), \tag{6}$$

where $p_{\mathrm{trail}}^{\max}(\mathcal{C}, R)$ is the optimal single trail inside $\mathcal{C}$ over R rounds and $\max_{i,j}(T_{\mathcal{C}}^{R})_{ij}$ is the maximal core-internal differential. This is an *aggregate* quantity over all input–output pairs of the core, and is distinct from the fixed-pair gap of Eq. (5). Theorem 1 concerns the core-level gap; the fixed-pair version requires a per-pair endpoint analysis, which is left as an open problem (Section 8).

3 Theoretical Results

3.1 Recurrent cores and the weighted adjacency matrix

Definition 1 (Recurrent core)**.** *A* closed differential set *is a set $\mathcal{C}$ of cell-difference states with $F_D(\mathcal{C}) \subseteq \mathcal{C}$. A* recurrent core *is an irreducible recurrent strongly connected component of the restriction of F_D to $\mathcal{C}$; equivalently, the transient states are removed and only the recurrent closed class is retained. Its* weighted adjacency matrix *is the nonnegative matrix defined entrywise by*

$$T_{\mathcal{C}}[i][j] \;=\; \sum_{\text{choices}} \;\prod_{\text{active S-boxes}} D[\text{in}][\text{out}], \tag{7}$$

the sum ranging over all DDT choice combinations that map the difference state i to the difference state j (one S-box output difference per active nibble), with $D[\cdot][\cdot] = m[\cdot][\cdot]/16$ the conventional DDT normalisation. $T_{\mathcal{C}}$ is nonnegative; it is not *assumed to be row stochastic, and Perron–Frobenius theory is applied to it as a nonnegative matrix, not as the transition matrix of a Markov chain.*

Definition 2 (Value graph, support, and the mask graph)**.** *The* value graph *is the weighted differential-transition digraph G_D of Section 2, whose vertices are the 2^{64} difference states and whose edge $\Delta \to \Delta'$ carries weight $m(\Delta, \Delta')$. For a cell difference Δ write* $\mathrm{supp}(\Delta) = \{\, i : \Delta_i \neq 0 \,\}$ *for its support, the set of active nibble positions. The* mask graph *is the directed graph whose vertices are all supports and which has an edge $s \to s'$ exactly when some difference Δ with* $\mathrm{supp}(\Delta) = s$ *satisfies* $\mathrm{supp}(F_D(\Delta)) = s'$*. The support map $\Delta \mapsto \mathrm{supp}(\Delta)$ is then a surjective graph homomorphism from the value graph to the mask graph: every edge $\Delta \to \Delta'$ of G_D is sent to the mask-graph edge* $\mathrm{supp}(\Delta) \to \mathrm{supp}(\Delta')$*.*

Lemma 1 (A value cycle projects to a mask cycle)**.** *Let $\Delta^{(0)} \to \Delta^{(1)} \to \cdots \to \Delta^{(\ell-1)} \to \Delta^{(0)}$ be a directed cycle in the value graph, and put $s_i = \mathrm{supp}(\Delta^{(i)})$. Then $s_0 \to s_1 \to \cdots \to s_{\ell-1} \to s_0$ is a directed cycle in the mask graph. Consequently a strongly connected value core can be supported only on masks lying in a single strongly connected component of the mask graph: a mask-level cycle is a necessary condition for a value-level recurrent core.*

Proof. For each i (indices modulo ℓ), the edge $\Delta^{(i)} \to \Delta^{(i+1)}$ gives $s_{i+1} = \mathrm{supp}\, F_D(\Delta^{(i)})$. This is exactly the condition defining the mask-graph edge $s_i \to s_{i+1}$. The projected sequence is therefore a closed walk, hence it contains a cycle, in the mask graph. For a strongly connected value core, every two of its states are joined by directed paths in both directions; projecting these paths to the mask graph shows that the supports of all states of the core are mutually reachable, so they lie in one strongly connected component of the mask graph. ■

Definition 3 (Two-level SCC decomposition)**.** *The core extraction of Section 4 proceeds through two nested strong-connectivity levels. At the* mask level*, the strongly connected components of the mask graph are the candidate supports. At the* value level*, the strongly connected components of the value graph induced on the concrete states whose support lies inside one mask-level component are computed by value closure. The recurrent core $\mathcal{C}$ is the maximal value-level strongly connected component retained after the transient states are stripped. By Lemma 1 every value-level component is*

contained in the preimage of a single mask-level component, so the two levels form a coherent hierarchy: mask-level strong connectivity is a coarse necessary filter, and value-level strong connectivity is the exact criterion that certifies the core.

Remark (mask-level versus value-level irreducibility). By Lemma 1 a recurrent *support* (a cycle in the mask graph) is necessary for a recurrent core; it is not sufficient. Perron–Frobenius theory applies to $T_{\mathcal{C}}$ only once the value-level graph on the retained concrete difference states has been verified to be strongly connected (irreducible). The value-closure BFS of Section 4 performs exactly this verification: it discards the value states whose one-round image leaves the candidate masks and reports the strongly connected component, so the matrix $T_{\mathcal{C}}$ fed to Perron–Frobenius is irreducible by construction. A mask cycle that closes to no irreducible value component is not a core (GIFT-64 exhibits a one-active mask cycle of this degenerate kind).

Remark (the full state graph is not strongly connected). The full differential-transition graph G_D over the 2^{64} difference states is far from strongly connected. One round of an SPN induces strong directionality: an edge $\Delta \to \Delta'$ by no means guarantees a return path, and many states are *transient*—once left they can never be revisited (for instance a difference that dies at an inactive S-box, or one that escapes into a high-weight regime). A recurrent core $\mathcal{C}$ is therefore not the whole graph but a strongly connected closed subgraph carved out of it: inside $\mathcal{C}$ any two states are mutually reachable, and the one-round image of $\mathcal{C}$ stays inside $\mathcal{C}$. Concretely, for the CRAFT mask $\{2, 5, 9, 13\}$ the value-closure BFS of Section 4 discards 45095 singleton transients and retains a single 2830-state recurrent core (see the walk-through below). $T_{\mathcal{C}}$ must be irreducible because the Perron–Frobenius theorem yields a unique dominant spectral radius ρ only when the associated digraph is strongly connected, so the dominant-term formula of Theorem 1 is valid for the recurrent core, not for an arbitrary reducible submatrix.

Running example. For CRAFT the mixing matrix admits a four-active fixed point, which we trace on the input difference $\{2{:}a, 5{:}a, 9{:}f, 13{:}f\}$ (all entries are nibble differences, inactive nibbles omitted). MixColumns acts column-wise: in the column $\{1, 5, 9, 13\}$ the four output nibbles are 0, $a \oplus f = 5$, f, f (nibble 1 cancels), so the intermediate state is $\{2{:}a, 5{:}5, 9{:}f, 13{:}f\}$. The cell permutation then maps $2 \mapsto 13$, $5 \mapsto 9$, $9 \mapsto 5$, $13 \mapsto 2$, giving $\{2{:}f, 5{:}f, 9{:}5, 13{:}a\}$. Finally the S-box layer, with $S(a) = 1$, $S(f) = 6$, $S(5) = b$, produces the output difference $\{2{:}6, 5{:}6, 9{:}b, 13{:}1\}$. The output has the same active-nibble mask $\{2, 5, 9, 13\}$ as the input. At the mask level the support-transition graph therefore maps this mask to itself, so the mask-level core consists of this single mask—a period-1 support cycle of size 1. The difference values, by contrast, evolve inside this mask: the value-level recurrent core built on this mask is analysed in the walk-through below (Figure 1). We return to this example throughout.

A detailed walk-through. The support-transition graph is built over masks of weight at most four; its recurrent classes isolate three four-active cores, two self-loops and one two-cycle. For the self-loop core on the mask $\{2, 5, 9, 13\}$, the value-closure BFS enumerates 47925 value states, of which 45095 are singleton transients and one 2830-state strongly connected component is the recurrent core. The weighted adjacency matrix of this component has spectral radius $\rho = 0.061628$; Karp's minimum-mean-cycle algorithm (Section 4) applied to the trail-weight matrix $W_{\mathcal{C}}$ of Lemma 2 returns minimum-mean-cycle weight $\mu = 8$, the minimum average weight of a directed cycle in the core, attained by a cycle whose four active S-boxes each realise the minimal S-box weight 2. The predicted per-round slope is therefore $\mu + \log_2 \rho = 3.98$ bits, and the measured gap at $R = 11$ is 30.4 bits, close to 11×3.98.

Definition 4 (Core-restricted weighted adjacency matrix)**.** *Let $\mathcal{C}$ be a strongly connected recurrent*

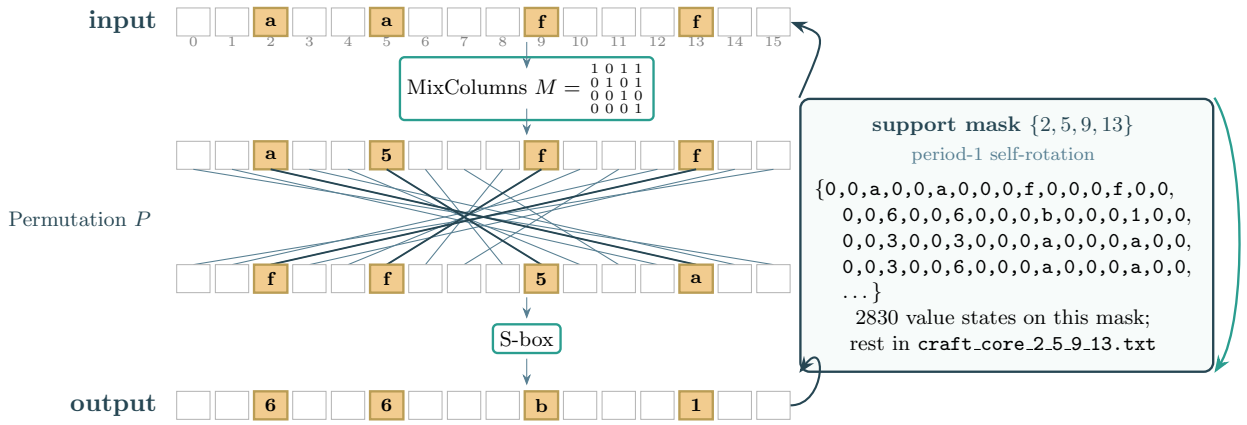


Figure 1: The four-active recurrent core of CRAFT. The support mask $\{2, 5, 9, 13\}$ (right box, enclosing its 2830 value states) self-rotates with period one, assigns the active positions of the input, and receives back the output support. The cell difference $\{2{:}a, 5{:}a, 9{:}f, 13{:}f\}$ passes through the MixColumns matrix M, the cell permutation P, and the S-box layer (in that order) to $\{2{:}6, 5{:}6, 9{:}b, 13{:}1\}$: the difference *values* change, but the active-nibble *mask* is invariant. Within this fixed support, each active S-box branches over its DDT row, and this branching accumulates trail multiplicity (spectral radius ρ).

component of the transition graph induced by the prescribed mask constraints, and let $T_{\mathcal{C}}$ be the restriction of the weighted adjacency matrix to the transitions whose image stays inside $\mathcal{C}$. In general $T_{\mathcal{C}}$ has row sums at most 1*: transitions leaving the constrained state set are discarded, so the spectral radius satisfies $\rho < 1$, and ρ measures the rate at which the path-sum* retained *in the core decays. This row-sum bound is a consequence of the restriction, not a stochasticity assumption: Perron–Frobenius is applied to $T_{\mathcal{C}}$ as an arbitrary nonnegative matrix. All core-level asymptotic statements below refer exclusively to paths whose transitions remain inside $\mathcal{C}$ for every round. If $\mathcal{C}$ were a closed class of the full process with all transitions retained, $T_{\mathcal{C}}$ would have row sums equal to* 1 *and spectral radius* 1*. We use "recurrent core" throughout for this restricted object.*

Remark (normalisation and dimensions). The combinatorial object underlying $T_{\mathcal{C}}$ is the integer-weighted adjacency matrix $A_{\mathcal{C}}$ with entries $\prod m[\cdot][\cdot]$. Every state of a recurrent core has the same active count k, so $A_{\mathcal{C}} = 16^k T_{\mathcal{C}}$ and $\rho(A_{\mathcal{C}}) = 16^k \rho$. The gap is invariant under this rescaling *only if both of its terms are measured in the same units*: in integer units the optimal trail has weight $k(-\log_2 m^{\max}) = k(w_{\min} - 4)$ per round and the integer path-sum grows as $\log_2 \rho(A_{\mathcal{C}}) = 4k + \log_2 \rho$ per round, whose difference is again $kw_{\min} + \log_2 \rho$ per round. Throughout the paper we work in the normalised units of Eq. (7), in which μ and $\log_2 \rho$ are both in bits and ρ is the quantity reported in the tables; mixing the two normalisations—for example using μ in bits together with $\log_2 \rho(A_{\mathcal{C}})$—is dimensionally incorrect.

Remark (truncation and escape). The full weighted adjacency matrix A of the 2^{64} difference states is never materialised; $T_{\mathcal{C}}$ is obtained from it by deleting every transition leaving the candidate masks. This deletion is what makes $T_{\mathcal{C}}$ sub-stochastic and is the origin of the "escape" mass: for the cores studied here the per-round escape probability $p_{\text{esc}} = \max_i(1 - \sum_j T_{\mathcal{C}}[i][j])$ is large (≈ 0.75–0.98 for CRAFT's four-active core). Consequently the core-restricted path-sum $T_{\mathcal{C}}^R$ is only a *lower bound* on the full path-sum (entrywise $A^R[i,j] \geq T_{\mathcal{C}}^R[i,j]$ for core states i, j), and the core theorem does *not* claim to approximate the full cipher. It characterises the restricted path-sum exactly;

bounding the discarded (escape-and-return) mass is a separate problem, discussed in Section 6 and left open in Section 8.

The recurrent-core abstraction is the key object of this paper. The states outside the core split into *transient* states—those on paths that enter the core, or that leave it, which can be visited only finitely many times (once left they are never revisited, or only a negligible set of paths returns)—and the *recurrent core*, which is cyclic and can be traversed arbitrarily many times. A transient or escaping path contributes at most a constant $O(1)$ to the R-round mass, whereas a path confined to the core multiplies a factor of order ρ in every round and therefore contributes a term of the form $c\,\rho^R$, exponential in R. Taking negative binary logarithms, the transient contribution remains a constant, while the core contribution becomes $R\,(-\log_2\rho)+O(1)$, a term linear in R; the per-round slope is therefore determined entirely by the core (the exponential base ρ together with the trail slope μ), while the transient and escape mass only shift the $O(1)$ intercept.

Lemma 2 (Optimal trail weight)**.** *Let $\mathcal{C}$ be a recurrent core with weighted adjacency matrix $T_{\mathcal{C}}$ (Definition 1). For each edge $i \to j$ of the core, let $W_{\mathcal{C}}[i][j]$ be the minimum weight $-\log_2 p_\pi$ over the DDT choice combinations π contributing to $T_{\mathcal{C}}[i][j]$ in Eq. (7). Let $\mu(\mathcal{C})$ be the minimum-mean-cycle weight of $W_{\mathcal{C}}$, i.e. the minimum, over all directed cycles of the core, of the average edge weight. Then, optimising over all start and end states of the core, the optimal R-round trail satisfies*

$$-\log_2 p_{\mathrm{trail}}^{\max}(\mathcal{C}, R) \;=\; R\,\mu(\mathcal{C}) + O(1).$$

Proof. Let $w(\tau)$ denote the weight of an R-round trail τ in the core, and let $w^* = \min_\tau w(\tau) = -\log_2 p_{\mathrm{trail}}^{\max}(\mathcal{C}, R)$. Put $n = |\mathcal{C}|$.

Lower bound on every trail. Any R-round trail is a walk in the finite state graph; decompose it into a simple path and a multiset of cycles. Each cycle of length ℓ has weight at least $\ell\,\mu(\mathcal{C})$ by minimality of the mean, and the simple path has length at most $n-1$; hence

$$w(\tau) \;\geq\; (R-(n-1))\,\mu(\mathcal{C}) \;=\; R\,\mu(\mathcal{C}) - K_1$$

for *every* R-round trail τ, with $K_1 = (n-1)\,\mu(\mathcal{C})$ independent of R. Minimising over τ gives the lower bound $w^* \geq R\,\mu(\mathcal{C}) - K_1$.

Upper bound from one trail. Fix a minimum-mean cycle $\gamma = (v_0, \ldots, v_{t-1}, v_0)$ of length t and weight $t\,\mu(\mathcal{C})$, choose the start state v_0, and write $R = q\,t + r$ with $0 \leq r < t$. The trail $\tilde{\tau}$ that traverses γ for q full laps and then its first r edges stays on $\gamma \subseteq \mathcal{C}$, has length exactly R, and has weight at most $q\,t\,\mu(\mathcal{C}) + r\,w_{\max} \leq R\,\mu(\mathcal{C}) + K_2$, where $w_{\max}$ is the largest edge weight and $K_2 = (t-1)\,w_{\max}$ is independent of R. Since $w^* \leq w(\tilde{\tau})$, this gives the upper bound $w^* \leq R\,\mu(\mathcal{C}) + K_2$.

Combining the two bounds yields $R\,\mu(\mathcal{C}) - K_1 \leq w^* \leq R\,\mu(\mathcal{C}) + K_2$, i.e. $w^* = R\,\mu(\mathcal{C}) + O(1)$, with both constants depending only on the finite graph. ∎

3.2 The trail space as a weighted subshift of finite type

The set of all differential trails is a one-step *subshift of finite type* (SFT): its symbols are the difference states, and a transition $\Delta \to \Delta'$ is allowed precisely when Δ' is a one-round image of Δ. The R-round trails are exactly the length-R paths of this directed graph.

To make the statistical-mechanical correspondence precise, we recall its minimal vocabulary. A statistical-mechanical system is a set of *configurations* σ with energies $E(\sigma)$; at inverse temperature β the *partition function* is $Z = \sum_\sigma e^{-\beta E(\sigma)}$, the *free energy* is $F = -\beta^{-1}\log Z$, and the *ground-state energy* is $E_0 = \min_\sigma E(\sigma)$. At zero temperature ($\beta \to \infty$) the partition function is dominated by the minimum-energy configurations, so $F \to E_0$; at positive temperature, configurations of higher

energy contribute through their *multiplicity* (the number of configurations at a given energy), which the *entropy* measures.

Our differential system instantiates this vocabulary as follows. The configurations are the R-round trails; the energy of a trail is its weight w (the negative binary logarithm of its probability); and the Boltzmann weight $e^{-\beta E}$ is replaced by the DDT-product weight 2^{-w}, i.e. inverse temperature $\beta = 1$ in binary-logarithm units; and the optimal trail is the ground state, of energy $R\mu$.

Endowing each allowed edge with the potential $\phi(\Delta, \Delta') = -w(\Delta, \Delta')$ (so that e^{ϕ} is the DDT product weight 2^{-w}) defines a *Ruelle transfer operator* whose matrix is the core-restricted weighted adjacency matrix $T_{\mathcal{C}}$, and the hull *partition function* is

$$Z_R \;=\; \sum_{\text{trails of length } R} \;\prod 2^{-w} \;=\; \mathbf{1}^{\top} T_{\mathcal{C}}^{R} \mathbf{1}, \tag{8}$$

the sum of the weights of all length-R trails. By the Ruelle–Perron–Frobenius theorem for an irreducible nonnegative matrix, $Z_R = \Theta(\rho^R)$ with $\rho = \rho(T_{\mathcal{C}})$. The free energy per round, $-\frac{1}{R}\log_2 Z_R$, therefore tends to $-\log_2 \rho$; its negation, the per-round logarithmic growth rate of the partition function, is the *topological pressure*

$$P \;=\; \lim_{R\to\infty} \frac{1}{R} \log_2 Z_R \;=\; \log_2 \rho. \tag{9}$$

The SFT formalism pays off through the *variational principle*, which gives $\log_2 \rho$ a strict thermodynamic meaning that plain matrix spectral theory does not supply. Let $\mathcal{M}_\sigma$ be the set of shift-invariant Borel probability measures on the trail space, let $h_\mu(\sigma)$ be the Kolmogorov–Sinai entropy of μ, and let $\int w\, d\mu$ be the expected per-round trail weight. The variational principle states

$$P \;=\; \log_2 \rho \;=\; \sup_{\mu \in \mathcal{M}_\sigma} \Big(h_\mu(\sigma) - \int w\, d\mu \Big), \tag{10}$$

i.e. the pressure is the maximal "entropy minus energy" over invariant measures.[1] Two consequences matter for this paper. First, the term $\log_2 \rho$ in Theorem 1 is genuinely an *entropy* term: it measures the trail multiplicity that the hull sustains, and it vanishes precisely when the optimal trail is unique. Second, the branch number constrains only the *energy* side of Eq. (10): Lemma 3 bounds $\mu(\mathcal{C})$ from below while the entropy $h_\mu(\sigma)$ is unconstrained by $\mathcal{B}$; the branch number disciplines the ground-state energy, not the entropy that drives hull clustering, which is Corollary 1.

The single optimal trail is the *ground state* of this system: its energy per round is the minimum-mean-cycle weight μ, hence its R-round energy is $R\mu$. The hull instead carries the full *free energy* $-\log_2 Z_R = -R\log_2 \rho + O(1)$, lower than $R\mu$ because the exponentially many trails contribute entropy. The clustering gap is therefore the difference between the ground-state energy and the free energy:

$$\begin{aligned} \mathrm{gap}_{\mathcal{C}}(R) \;&=\; R\mu - \big(-R\log_2 \rho\big) + O(1) \\ &=\; R\big(\mu + \log_2 \rho\big) + O(1), \end{aligned} \tag{11}$$

which is Theorem 1 in the language of thermodynamic formalism. The formulation makes the mechanism exact and model-free: clustering is the *entropic* contribution of trail multiplicity to the free energy, which is invisible to the zero-temperature (single-trail) description. The branch number

[1] The variational principle is usually stated for mixing shifts; for a periodic irreducible SFT it remains valid by passing to the period-adjusted matrix $T_{\mathcal{C}}^{t}$ or by taking the supremum over shift-invariant measures, and the identity $P = \log_2 \rho$ is unaffected.

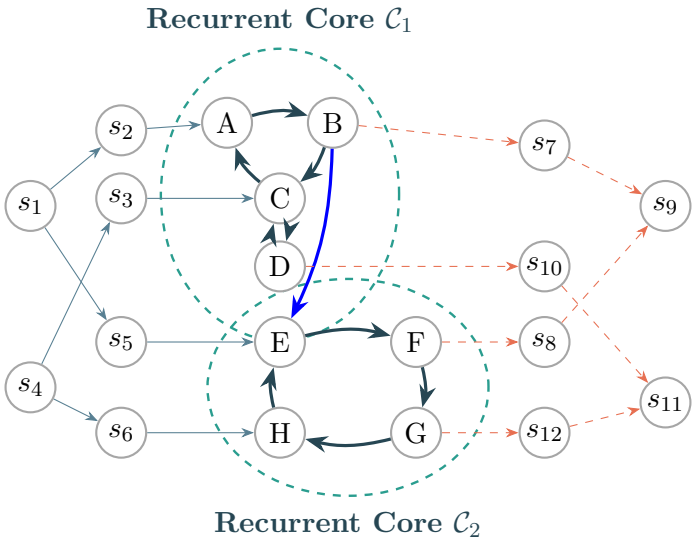


Figure 2: The differential trail space as a weighted subshift of finite type. Two recurrent cores (shaded light teal): $\mathcal{C}_1 = \{A, B, C, D\}$ contains a 3-cycle $\{A, B, C\}$ and a 2-cycle $\{C, D\}$ sharing the vertex C, and $\mathcal{C}_2 = \{E, F, G, H\}$ is a single 4-cycle. Navy solid arrows are the intra-core cycles. Slate solid arrows are transient input paths entering a core (s_1 and s_4 each branch into both cores). Coral dashed arrows are transitions leaving a core to the common output states s_9, s_{11}. The blue arrow $B \to E$ is a transition that leaves $\mathcal{C}_1$ and enters $\mathcal{C}_2$.

bounds the ground-state energy ($\mu \geq (\mathcal{B}/2)w_{\min}$, by Lemma 3) but not the pressure $P = \log_2 \rho$; the pressure is a property of the transfer operator, and two layers of identical branch number can carry transfer operators of very different spectral radius (Section 5); Figure 2 depicts this structure.

3.3 The core-spectral theorem

Theorem 1 (Core-spectral gap)**.** *Let $\mathcal{C}$ be a recurrent core with weighted adjacency matrix $T_{\mathcal{C}}$ (non-negative, not assumed row stochastic) of spectral radius $\rho = \rho(T_{\mathcal{C}})$, and let $\mu(\mathcal{C})$ be the minimum-mean-cycle weight of the trail-weight matrix $W_{\mathcal{C}}$ defined in Lemma 2. Then, as $R \to \infty$,*

$$\begin{aligned} -\log_2 \mathrm{DP}_{\mathcal{C}}(R) &= -R \log_2 \rho + O(1), \\ \mathrm{gap}_{\mathcal{C}}(R) &= R\left(\mu(\mathcal{C}) + \log_2 \rho\right) + O(1). \end{aligned} \tag{12}$$

If the recurrent cycle admits an all-optimal transition, then $\mu(\mathcal{C}) = k\, w_{\min}$ for the k active S-boxes; otherwise $\mu(\mathcal{C}) > k\, w_{\min}$.

The two lines of Eq. (12) are the hull energy and the gap. The first line follows from $Z_R = \Theta(\rho^R)$, where $Z_R = \mathbf{1}^\top T_{\mathcal{C}}^R \mathbf{1}$ is the partition function of Eq. (8) and the estimate holds by Ruelle–Perron–Frobenius for the irreducible nonnegative $T_{\mathcal{C}}$. The second line is Eq. (11), obtained by subtracting the optimal trail weight of Lemma 2, $R\,\mu(\mathcal{C}) + O(1)$, from the hull energy.

The proof uses only Perron–Frobenius theory for irreducible nonnegative matrices; in particular it never invokes row stochasticity, a Markov property, or the stochastic-equivalence assumption. The quantity $\mathrm{DP}_{\mathcal{C}}(R)$ in Eq. (12) is the maximal entry of the normalised path-sum $T_{\mathcal{C}}^R$ over core states, a well-defined combinatorial quantity.

Periodicity of the transient. For an irreducible nonnegative $T_{\mathcal{C}}$ of period $d > 1$, plain Perron–Frobenius theory only yields $T_{\mathcal{C}}^R = \rho^R P + o(\rho^R)$, which is not the correct dominant term: the

dominant factor is d-periodic rather than constant, and plain matrix spectral theory does not identify it. The Fourier expansion below does. Concretely, the error term $O(1)$ in Eq. (12) is *periodic* in R with period d rather than converging to a constant: the power $T_{\mathcal{C}}^R$ cycles among the d cyclic classes. The d eigenvalues of maximal modulus of $T_{\mathcal{C}}$ are $\rho\,\omega^k$, with $\omega = e^{2\pi i/d}$ a primitive d-th root of unity, $k = 0, \ldots, d-1$; hence

$$T_{\mathcal{C}}^R \;=\; \rho^R \sum_{k=0}^{d-1} \omega^{kR}\, P_k \;+\; E_R, \qquad \|E_R\| \;\le\; C\, R^{m-1}\, \lambda_*^R,$$

where P_k is the spectral projection onto the eigenspace of $\rho\,\omega^k$, E_R is the residual matrix, $\lambda_* < \rho$ is the next eigenvalue magnitude, m is the largest Jordan-block size of the eigenvalues of modulus λ_* ($m = 1$ when $T_{\mathcal{C}}$ is diagonalisable), $\|\cdot\|$ is any submultiplicative matrix norm, and C is a constant independent of R. The factor R^{m-1} is the only departure from the diagonalisable bound; the exact residual N^q of Proposition 1 and the numerical validation of Table 6 avoid this Jordan-block issue, which is negligible for the actual matrices. The $O(1)$ transient of the gap is therefore the binary logarithm of a sum of d Fourier modes $c_k\,\omega^{kR}$—the discrete Fourier transform of the d cyclic classes—whose period is exactly d. The leading coefficient $R(\mu + \log_2 \rho)$ is unaffected, but the $O(1)$ transient must be read as a bounded, d-periodic sequence. The core-level gap of Eq. (6) therefore takes the maximum over all core pairs, because the maximum exists in every residue class and the slope is well-defined. A *fixed* pair (i, j) would instead require R restricted to a single residue class of d (otherwise $T_{\mathcal{C}}^R[i, j] = 0$); the period-adjusted matrix $T_{\mathcal{C}}^d$ of Proposition 1 removes this oscillation.

Theorem 1 is a statement about the *restricted* subsystem: $\mathrm{DP}_{\mathcal{C}}(R)$ is the differential probability of the trails that stay inside the core, not the full cipher's differential. Extending the core-level gap to the full cipher requires bounding the total probability of the trails that leave the core and return, which is left open (Section 8). Throughout the paper we therefore distinguish three quantities: the *core-restricted gap* (only trails that stay inside the core), the *fixed-input/output gap* (the gap of one specified input and output pair, used as an experimental validation), and the *full-cipher security bound* (all trails, which additionally requires the bounds just mentioned). Theorem 1 characterises only the first quantity.

Proof. Let v be a strictly positive Perron right eigenvector, $T_{\mathcal{C}} v = \rho v$, and set $D = \mathrm{diag}(v)$. Then $B = D^{-1} T_{\mathcal{C}} D$ is nonnegative with $B\mathbf{1} = D^{-1} T_{\mathcal{C}} v = \rho \mathbf{1}$, so every row sum of B equals ρ; by induction every row sum of B^R equals ρ^R, and $\|B^R\|_\infty = \rho^R$. Since $T_{\mathcal{C}}^R = D B^R D^{-1}$,

$$\rho^R \;\le\; \|T_{\mathcal{C}}^R\|_\infty \;\le\; \|D\|_\infty \, \|B^R\|_\infty \, \|D^{-1}\|_\infty \;=\; \frac{v_{\max}}{v_{\min}}\, \rho^R.$$

Hence $\|T_{\mathcal{C}}^R\|_\infty = \Theta(\rho^R)$, and since the max-entry norm and the ∞-norm differ by at most a factor n, the maximal entry of $T_{\mathcal{C}}^R$ is $\Theta(\rho^R)$; this yields $-\log_2 \mathrm{DP}_{\mathcal{C}}(R) = -R \log_2 \rho + O(1)$. The argument uses only the positive Perron eigenvector and is therefore valid for periodic as well as aperiodic cores. Combining with Lemma 2 gives $\mathrm{gap}_{\mathcal{C}}(R) = R(\mu(\mathcal{C}) + \log_2 \rho) + O(1)$. ∎

Theorem 1 is an *asymptotic* statement: the two equalities hold up to an $O(1)$ additive term that does not grow with R. The leading coefficient is exact and is what determines the per-round slope of the gap. In cryptographic terms, the theorem decomposes the clustering gap into two contributions with clean interpretations. The term $k\, w_{\min}$ is the per-round weight of the optimal trail *inside* the core: each of the k active S-boxes contributes at least $w_{\min}$. The term $\log_2 \rho$ is the rate at which trail multiplicity accumulates around the core's recurrent cycle; because $\rho < 1$ for a genuine core, $\log_2 \rho < 0$, and $k\, w_{\min} + \log_2 \rho$ is the net per-round gap. A core whose recurrent cycle is nearly

deterministic carries essentially one optimal transition, so $T_{\mathcal{C}}$ is close to a permutation matrix scaled by $2^{-kw_{\min}}$: its spectral radius is $\rho \approx 2^{-kw_{\min}}$, $\log_2 \rho \approx -kw_{\min}$, and the per-round gap is small (vanishing in the fully deterministic limit, because the entropy term $\log_2 \rho$ then exactly cancels the ground-state energy $kw_{\min}$). Conversely, a core whose cycle admits many DDT choice combinations has larger entries in $T_{\mathcal{C}}$ (the probabilities add), so ρ is larger (closer to 1), $\log_2 \rho$ is closer to 0, and the gap approaches the full trail weight $kw_{\min}$ per round. CRAFT falls exactly into this regime: its four-active fixed point has $\rho = 0.0616$ against the deterministic baseline $2^{-8} = 0.0039$, giving the large slope $8 - 4.02 = 3.98$ bits per round.

The theorem is asymptotic in R: for finite round counts the actual gap differs from $R(\mu(\mathcal{C}) + \log_2 \rho)$ by an $O(1)$ transient, which in our experiments is a few bits at $R = 11$ and shrinks in relative terms as R grows. The weighted adjacency matrix $T_{\mathcal{C}}$ need not have a simple dominant eigenvalue: for a periodic core of period t, the t-th roots of unity times ρ are all of modulus ρ, so the dominant eigenspace has dimension at least t. The convergence rate is governed by the spectral gap $\gamma = \lambda_*/\rho$, where λ_* is the largest eigenvalue magnitude strictly below ρ; this is made precise in Proposition 1.

Lemma 3 (Branch number floors the trail slope, not the mass term)**.** *For every recurrent core $\mathcal{C}$, the minimum-mean-cycle weight satisfies*

$$\mu(\mathcal{C}) \;\geq\; \frac{\mathcal{B}}{2}\, w_{\min}.$$

Consequently the branch number bounds the trail term of the gap from below, whereas the mass term $\log_2 \rho$ is not constrained by $\mathcal{B}$.

Proof. Since the S-box is a nibble-wise bijection it preserves the support, i.e. $\mathrm{wt}(S(x)) = \mathrm{wt}(x)$. Let $v_0 \to v_1 \to \cdots \to v_{t-1} \to v_0$ be any cycle of the core with active counts $k_i = \mathrm{wt}(v_i)$. For a CRAFT-type transition $v_{i+1} = S(L(v_i))$, support preservation gives $\mathrm{wt}(v_{i+1}) = \mathrm{wt}(L(v_i))$, hence $\mathrm{wt}(v_i) + \mathrm{wt}(v_{i+1}) = \mathrm{wt}(v_i) + \mathrm{wt}(L(v_i)) \geq \mathcal{B}$. For a GIFT-type transition $v_{i+1} = L(S(v_i))$, applying the branch-number definition to $\Delta = S(v_i)$ gives $\mathrm{wt}(S(v_i)) + \mathrm{wt}(L(S(v_i))) \geq \mathcal{B}$, and since $\mathrm{wt}(S(v_i)) = \mathrm{wt}(v_i)$ this is again $\mathrm{wt}(v_i) + \mathrm{wt}(v_{i+1}) \geq \mathcal{B}$. In both cases $\mathrm{wt}(v_i) + \mathrm{wt}(v_{i+1}) \geq \mathcal{B}$. Summing over i, $2\sum_i k_i \geq t\,\mathcal{B}$, hence $(\sum_i k_i)/t \geq \mathcal{B}/2$. The edge $v_i \to v_{i+1}$ has weight at least $k_i\, w_{\min}$, so every cycle has mean weight at least $w_{\min}(\sum_i k_i)/t \geq (\mathcal{B}/2)\, w_{\min}$; taking the infimum over cycles gives the claim. ∎

Corollary 1 (Branch number is not sufficient)**.** *By the variational principle (Eq.* (10)*), the gap slope of Theorem 1 splits into an energy term $\mu(\mathcal{C})$ and an entropy term $\log_2 \rho$, and the branch number $\mathcal{B}$ constrains only the energy term (Lemma 3). The ablation of Section 5 exhibits, under the same S-box, a branch-2 layer (CRAFT, slope* 3.98*) whose per-round gap exceeds that of a branch-4 layer (J-I, slope* 1.572*); hence $\mathcal{B}$ alone does not determine the gap slope, and branch number alone is insufficient to rank hull-clustering resistance.*

The corollary is the design-relevant consequence of Theorem 1. The branch number controls how many S-boxes must be active along any trail—a lower-bound quantity that protects the *single trail*—but it says nothing about how many trails share the same input and output difference, which is what the hull measures. Two linear layers with identical branch number can therefore have radically different hull behaviour, a fact that Section 5 exhibits concretely.

3.4 Structural lemmas

Lemma 4 (Multi-core competition)**.** *Let $\mathcal{C}_1, \ldots, \mathcal{C}_m$ be pairwise disjoint recurrent cores with spectral radii $\rho_1, \ldots, \rho_m$, and put $\rho_{\max} = \max_i \rho_i$. Consider the* core-restricted *dynamics consisting of all*

trails that enter one core, stay inside that single core, and never leave it. If the fixed input difference reaches a maximal-ρ core, and the fixed output difference j satisfies $v_{\max}(j) > 0$ for the right Perron vector $v_{\max}$ of that core, then at high round counts the core-restricted hull is dominated by the maximal-ρ core, and every other core contributes a term that decays relatively as $(\rho_i/\rho_{\max})^R$.

Proof. Each core is an absorbing recurrent component of the core-restricted dynamics, so the restricted weighted adjacency matrix block-diagonalises over the cores, and each block contributes $c_i\,\rho_i^R\big(1+o(1)\big)$ to the core-restricted hull mass, where $c_i \geq 0$ is determined by the incoming mass and the Perron eigenvector of that core. Factoring out $\rho_{\max}^R$, every non-dominant block decays as $(\rho_i/\rho_{\max})^R$, which vanishes exponentially as $R \to \infty$. The two hypotheses of the statement are exactly what ensure $c_{\max} > 0$: reachability from the fixed input supplies positive incoming path weight to the maximal-ρ core, and overlap of the fixed output with the support of the core's Perron vector makes that weight observable. ■

Lemma 4 explains why a cipher can carry several recurrent cores yet exhibit a single, well-defined asymptotic slope: the maximal-ρ core is the Perron-dominant mode of the core-restricted dynamics. A dominant core only drives a particular fixed input–output pair if that pair actually reaches and observes it, and the lemma does not account for trails that leave a core and re-enter later. It also motivates the engineering choice in Section 4 of reporting, for each core, the spectral radius of its *dominant* irreducible block rather than of the whole (possibly reducible) weighted adjacency matrix.

Lemma 5 (Permutation symmetry)**.** *Let $\mathcal{C}$ and $\mathcal{C}'$ be two recurrent cores whose restricted transition structures are related by the cipher's permutation layer: suppose there is a permutation matrix Π mapping the state set of $\mathcal{C}$ bijectively onto that of $\mathcal{C}'$ and* preserving the transition weights*, i.e. $T_{\mathcal{C}'} = \Pi\, T_{\mathcal{C}}\, \Pi^{-1}$. Then the two matrices are permutation-similar, and their spectral radii are equal.*

Proof. Permutation similarity $T_{\mathcal{C}'} = \Pi\, T_{\mathcal{C}}\, \Pi^{-1}$ preserves the characteristic polynomial and hence the spectrum; in particular $\rho(T_{\mathcal{C}'}) = \rho(T_{\mathcal{C}})$. The weight-preservation condition in the statement is what guarantees that a permutation of the states induced by the permutation layer actually realises this similarity; without it, two cores occupying the same set of positions need not have related transition matrices. ■

Lemma 5 turns a numerical observation into a structural fact. For CRAFT, the mixing layer treats all four columns identically and is translation-invariant; consequently the three four-active cores are images of one another under the permutation layer, and all three have $\rho = 0.061628$ (Section 5). The lemma implies that a designer need not screen each core of a symmetric family individually: screening one representative suffices.

3.5 The identity baseline

For the identity linear layer ($L = P$), the recurrent core of a single nibble has weighted adjacency matrix $T_{\mathcal{C}} = D$ with all row sums equal to 1 (the single active S-box is a genuinely closed class, so no mass is discarded) and spectral radius $\rho = 1$. This is the degenerate row-stochastic special case in which the $O(1)$ term of Theorem 1 is no longer negligible: the Perron–Frobenius limit of D^R is the stationary vector π, and the gap is

$$\mathrm{gap}(R) \;=\; 2R - \log_2(1/\pi_{\max}) + o(1), \tag{13}$$

where $\pi_{\max}$ is the maximal entry of π. This is the *S-box baseline* of clustering: even without any mixing, the DDT of the S-box alone sustains a hull of trails, and the baseline is the yardstick

against which a mixing layer's amplification is measured. For the CRAFT/PRESENT-type S-box, $\pi_{\max} = 2^{-3.907}$, giving a baseline gap of 18.1 bits at $R = 11$.

3.6 Finite-round error bound

The asymptotic law of Theorem 1 carries an $O(1)$ transient. To quantify it, we pass to the period-adjusted matrix $M = T_{\mathcal{C}}^t$, whose Perron root ρ^t is the eigenvalue of maximal modulus (the dominant eigenspace has dimension t, the period). Proposition 1 below isolates the finite-round deviation exactly.

Proposition 1 (Finite-round error: exact residual and computable bound)**.** *Let $\mathcal{C}$ be a recurrent core of period t, let $M = T_{\mathcal{C}}^t$ be its period-adjusted weighted adjacency matrix with dominant eigenvalue ρ^t (of multiplicity t), spectral projection P onto the dominant eigenspace, and residual $N = M - \rho^t P$, so that $\rho(N) = \lambda_*^t < \rho^t$. Then, exactly,*

$$M^q \;=\; \rho^{tq} P + N^q, \tag{14}$$

so the matrix-residual error $M^q - \rho^{tq} P = N^q$ at round $r = tq$ has max-entry norm exactly $\|N^q\|_{\max}$, a quantity computed directly. Moreover, for the real Schur form $N = QUQ^\top$ with Q orthogonal,

$$\|N^q\|_{\max} \;\le\; \|Q\|_\infty \, \|U^q\|_\infty \, \|Q^\top\|_\infty, \tag{15}$$

a computable bound that does not involve the eigenvector condition number.

Proof. Since P is the spectral projection onto the dominant eigenspace, $PN = NP = 0$, and the identity $M = \rho^t P + N$ yields $M^q = \rho^{tq} P + N^q$ exactly. For the bound, orthogonality $Q^\top Q = I$ gives $N^q = Q\,U^q Q^\top$, and hence

$$\begin{aligned} \|N^q\|_{\max} &\le \|N^q\|_\infty \\ &= \|Q\,U^q Q^\top\|_\infty \\ &\le \|Q\|_\infty \, \|U^q\|_\infty \, \|Q^\top\|_\infty, \end{aligned}$$

where the first step is $\|A\|_{\max} \le \|A\|_\infty$ and the last is submultiplicativity of the ∞-norm. Both factors are finite because Q is orthogonal: each row of Q has unit Euclidean norm, so $\|Q\|_\infty \le \sqrt{n}$. ∎

Proposition 1 gives the finite-round error as an *exact* quantity $\|N^q\|_{\max}$ rather than an a priori bound, together with a computable bound whose constant is $\|Q\|_\infty \|Q^\top\|_\infty$ (of order n for an orthogonal Q) instead of the eigenvector condition number $\kappa_\infty(V)$, which is numerically enormous for these matrices. The exact error and the Schur bound are reported in Section 5.

The proposition bounds the *matrix residual* $\|N^q\|_{\max}$: how well $\rho^{tq} P$ approximates the matrix power M^q. This is not the same as the gap prediction error, which decomposes as

$$\begin{aligned} \left|\mathrm{gap}_{\mathcal{C}}(R) - R(\mu + \log_2 \rho)\right| \le{}& \left| -\log_2 p_{\mathrm{trail}}^{\max} - R\mu \right| \\ &+ \left| \log_2 \|T_{\mathcal{C}}^R\|_{\max} + R \log_2 \rho \right|, \end{aligned}$$

the first (path-side) term being bounded by the constants of Lemma 2, and the second (matrix-side) term being the residual of Proposition 1. The two components are reported separately in Section 5.

4 The `coretool` Algorithm

The tool implements the theory in three modules, described in turn (Figure 3). The source code, cipher configurations, and the scripts reproducing every number in this paper will be released publicly upon publication and are available from the authors on request for review purposes.

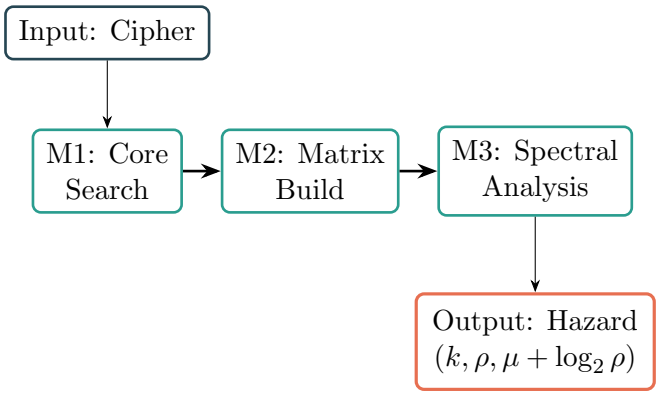


Figure 3: The `coretool` workflow: from a cipher specification through modules M1–M3 to the hazard signature $(k, \rho, \mu + \log_2 \rho)$.

Module M1 — core search. The search proceeds in two stages. First, a *support-transition graph* is built over low-weight masks: for each mask (a set of active nibble positions), the set of masks reachable in one round is computed by enumerating the value assignments through the linear layer. The recurrent classes (strongly connected components) of this graph are candidate cores. Second, and crucially, a *value-closure* breadth-first search is run inside each candidate: starting from a seed, the tool keeps exactly those concrete difference values whose one-round image stays inside the candidate masks. A mask-level cycle is necessary but not sufficient for a core; the value closure is what verifies that the cycle is actually closed under the value dynamics. For instance, GIFT-64 exhibits a one-active mask cycle that closes to only two value states and carries negligible mass, whereas its two-active class closes to the 32-state recurrent core that dominates the tail. For bit-permutation layers the mask pre-filter degenerates into a single giant strongly connected component, so the tool falls back to a direct active-class closure over the k-active states. If a single mask SCC carries several disconnected recurrent value components, the single-seed BFS is re-seeded from unvisited states until all of them are found.

Module M1 is *not* a complete full-space search. It prioritises fixed active-nibble patterns, and it does not claim to enumerate every invariant closed set over the full 2^{64} state space. This restriction is deliberate and is justified from the security viewpoint: a core that is dangerous for hull clustering must be reachable by *low-weight* trails, because high-weight trails contribute negligible probability. Cores with many active S-boxes have, by Theorem 1, a large per-round trail weight $kw_{\min}$ and therefore a small differential contribution; the cores that matter for practical attacks are exactly the low-active-nibble cores on which Module M1 concentrates. We therefore search low-weight masks exhaustively within the stated mask bound, and we state the bound explicitly rather than implying completeness.

Why low-weight cores are the relevant search target. The search bound is a heuristic, not a completeness proof. A single trail passing through a recurrent core with k active S-boxes has per-round weight at least $k\, w_{\min}$, hence R-round weight at least $R\, k\, w_{\min}$ and probability at most

$2^{-R\,k\,w_{\min}}$. A single trail can reach probability 2^{-t} only if

$$k \;\le\; \frac{t}{R\,w_{\min}} \tag{16}$$

This is a necessary condition on the single trail, and it does not by itself bound the hull: many heavier trails can sum to a larger total. For instance, the CRAFT four-active core has $k = 4$ active S-boxes, above the single-trail threshold $k \le t/(R\,w_{\min}) \approx 2.3$ obtained by substituting its attained weight $t = 49.559$ at $R = 11$ into Eq. (16), yet the hull still reaches $2^{-49.559}$. The low-weight search is therefore a targeted heuristic that concentrates on the cores whose single trails are light enough to matter, with the mask bound stated explicitly; formally bounding the total probability contributed by all omitted high-weight paths is left as an open problem (Section 8). We call a core *low-activity* when Eq. (16) holds and *high-activity* otherwise, always relative to the target. The mask bound is a search parameter, not a security assumption.

Module M2 — weighted adjacency matrix. For each candidate core, the value states are enumerated and $T_{\mathcal{C}}$ is built by multiplying the per-active-S-box $D[\cdot][\cdot]$ entries. The matrix is checked for irreducibility and decomposed into strongly connected blocks; the block of maximal spectral radius is the recurrent core, and the remaining blocks are transient and contribute only the $O(1)$ term. For CRAFT's core $\{2, 5, 9, 13\}$ the closure has 47,925 value states, of which 45,095 are singleton transients and one 2,830-state block is the recurrent core.

Module M3 — spectral analysis and validation. The spectral radius ρ and the right Perron vector are computed by sparse Arnoldi iteration (small blocks use a dense eigendecomposition). The minimum-mean-cycle weight $\mu(\mathcal{C})$—the precise per-round trail slope—is computed by Karp's algorithm on the trail-weight matrix $W_{\mathcal{C}}$ of Lemma 2. The finite-round transient is quantified by a small-scale exact enumeration carried out by a *mass dynamic program*: a frontier over the value states is iterated for $R - 1$ rounds restricted to the core followed by a final unrestricted round, and the measured gap is compared against the predicted slope. Because the right Perron vector component of a state is proportional to its high-round hull mass, the top components of the Perron vector provide representative asymptotically-optimal input states for this validation; the Perron vector provides representative input states, not an exhaustive search (Figure 4).

Algorithm 1 Value-closure BFS (the subroutine of Module M1; C++-accelerated)

Require: candidate masks $M = \{m_0, \ldots, m_{t-1}\}$, seed difference Δ, DDT D, linear layer L
Ensure: the value states of the recurrent core, or no-closure

$Q \leftarrow \{\Delta\}$; $V \leftarrow \{\Delta\}$ {queue and visited set}
while $Q \neq \emptyset$ **do**
  $x \leftarrow \mathrm{pop}(Q)$
  for each active-S-box DDT choice of x **do**
    $x' \leftarrow$ one-round image of x (apply L and the S-box layer in the cipher order)
    if $\mathrm{supp}(x') \notin M$ **then**
      continue
    end if
    if $x' \notin V$ **then**
      $V \leftarrow V \cup \{x'\}$; $Q \leftarrow Q \cup \{x'\}$
    end if
  end for
end while
return V {V is closed; if it carries no recurrent class, report no core}

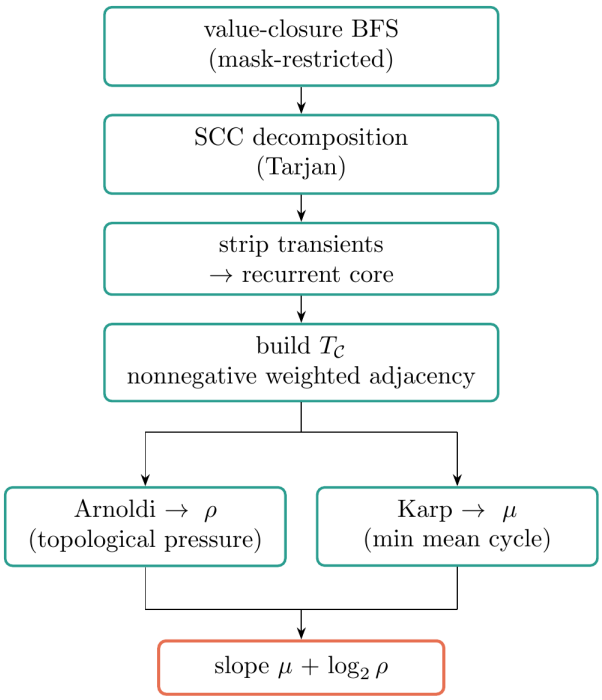


Figure 4: The core-spectral pipeline: a value-closure BFS is followed by SCC decomposition, the transient states are stripped to leave the irreducible recurrent core, its nonnegative weighted adjacency matrix $T_{\mathcal{C}}$ is built, and Arnoldi (for ρ) and Karp (for μ) yield the slope $\mu + \log_2 \rho$.

The value-closure BFS and Karp's algorithm are the computational bottlenecks and are implemented in C++ behind a ctypes interface, with the Python implementation retained as a reference; this yields an order-of-magnitude speedup (CRAFT's core closure drops from minutes to seconds).

Complexity. The dominant cost is the value-closure BFS (Module M1): it visits each core state once and, per state, enumerates the DDT choices of its active S-boxes. For a core with m states and at most q active S-boxes per round, the BFS costs $O(m \cdot c^q)$, where $c \leq 8$ is the maximal DDT branching, independently of the number of rounds R. The mask pre-filter costs $O\left(\binom{16}{w_{\max}} \cdot 15^{w_{\max}}\right)$ value enumerations. The spectral step (Module M3) costs $O(n^3)$ for a dense eigendecomposition of an n-state core, or less with sparse Arnoldi, and Karp's algorithm costs $O(ne)$ for e edges. The mass DP of Module M3 runs over the m core states for R rounds and costs $O(R \cdot m \cdot c^q)$. In our experiments the C++ value-closure BFS reduces the closure of CRAFT's core from minutes to seconds.

Implementation details. The value-closure BFS stores states as 64-bit integers and the visited set as an open-addressing hash table; the DDT transitions are precomputed per S-box input difference so that each state expansion is a sequence of table lookups and XORs. The C++ extension is loaded through a `ctypes` interface, and the Python reference implementation is retained as an oracle for consistency checks: on every cipher in our set, the C++ closure coincides exactly with the Python closure.

5 Experiments

We instantiate the framework on four ciphers and on controlled linear-layer ablations. All computations use `coretool` with the C++ engine. Table 2 summarises the results; it is reproduced unchanged from our reference implementation, and the slope column reports the precise per-round slope $\mu(\mathcal{C}) + \log_2 \rho$.

Cipher / layer	branch	core	ρ	slope	gap (bits)
identity	2	1-active	$\rho = 1$, $\pi_{\max} = 2^{-3.907}$	$2 - 3.907/R$	18.1 (R=11)
SKINNY-MC (abl.)	2	1-active	1.0	≈ 2.0	≈ 18.1
CRAFT	2	4-active×3	0.061628	3.980	30.4 (R=11)
J-I (abl.)	4	2-active	0.185826	1.572	15.4 (R=11)
GIFT-64	bit-perm	2-active 32	0.038273	0.292	0.585 (R=9)
PRESENT-80	bit-perm	2-active 240	0.069049	0.644	≈ 7 (R=11)
SKINNY-64	$M \circ P$	none ≤ 4	–	–	weak

Table 2: Core-spectral results. The slope column is the core-level asymptotic quantity $\mu(\mathcal{C}) + \log_2 \rho$ of Theorem 1; the gap column is the finite-round fixed-input gap and is not directly comparable to R×slope.

Per-cipher summary. CRAFT is the canonical strongly-clustered instance: its mixing matrix carries a four-active fixed point, and the 30.4-bit gap at $R = 11$ is the largest in our set. GIFT-64 and PRESENT-80, whose linear layers are bit permutations, carry two-active recurrent cores of 32 and 240 states respectively; their gaps are small (0.585 and about 7 bits) because the cores are nearly deterministic. SKINNY-64, whose mixing matrix is also of branch number 2, induces no recurrent core of weight at most four, and its clustering stays at the S-box baseline. The four ciphers thus span the full spectrum from negligible to strong clustering, and the ordering is exactly what the core-spectral slope $\mu + \log_2 \rho$ predicts—not the radius ρ alone: PRESENT-80 has a larger ρ than CRAFT yet a much smaller slope, because k differs.

This spread of behaviours is not an accident of the four ciphers but a direct reading of the core structures. CRAFT's four-active fixed point admits many parallel DDT transitions per round, so its core has $\rho = 0.0616$, far above the deterministic baseline $2^{-kw_{\min}} = 2^{-8} = 0.0039$, and hence a large slope; GIFT-64's and PRESENT-80's two-active cores are close to deterministic, so their ρ stays near the deterministic baseline and the slope is correspondingly small; SKINNY-64 has no low-weight core at all, so it stays at the S-box baseline. The same explanation, read backwards, is what Section 7 turns into a design criterion.

A GIFT-64 walk-through. The same pipeline applied to GIFT-64 isolates a two-active recurrent core of 32 value states. Its edges alternate between weight 4 (both active S-boxes use 4/16 transitions) and weight 6 (both use 2/16 transitions), so no all-weight-4 cycle exists and the minimum-mean-cycle weight is $\mu = 5$. The spectral radius is $\rho = 0.038273$, giving a predicted slope $\mu + \log_2 \rho = 0.292$ bits per round—an order of magnitude below CRAFT's 3.98. The measured gap at $R = 9$ is 0.585 bits, and the finite-round transient of Section 5 accounts for the difference between 9×0.292 and 0.585. The contrast with CRAFT is stark: both cores have two or four active S-boxes, but the fixed-point structure of CRAFT sustains a far richer trail multiplicity, which is exactly the ρ-gap of Table 2.

Reading the slope and gap columns. The "slope" column is the *core-level* asymptotic slope $\mu(\mathcal{C}) + \log_2 \rho$ of Theorem 1; the "gap" column is the finite-round fixed-input gap in Hadipour's convention. For CRAFT, the slope 3.98 gives $11 \times 3.98 = 43.78$ bits for a trail confined to the four-active bulk, whereas the fixed-input gap is $80 - 49.559 = 30.4$ bits. The 13.4-bit difference is a systematic finite-round effect (three-active head and tail, plus endpoint effects), not a small $O(1)$ transient; the two columns must not be conflated.[2]

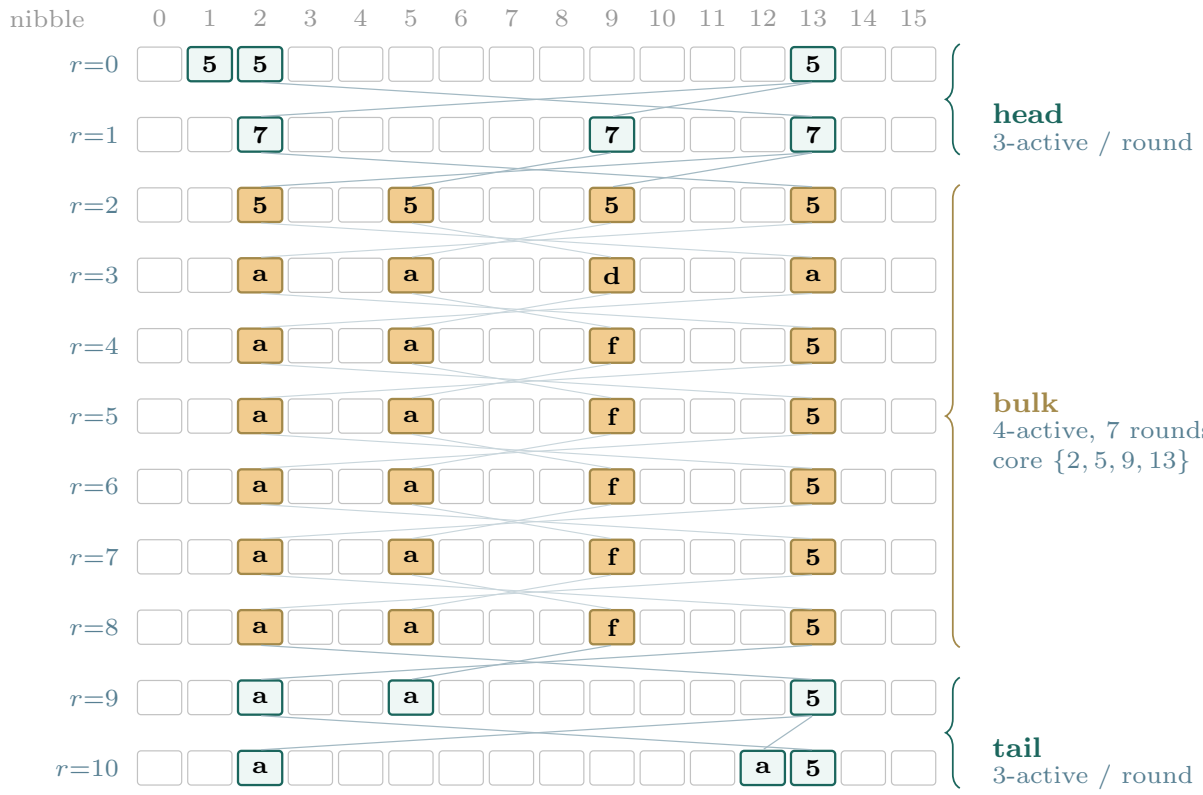


Figure 5: The optimal eleven-round CRAFT trail: a three-active head (rounds 0 to 1), a four-active bulk (rounds 2 to 8) sitting on the recurrent core $\{2, 5, 9, 13\}$, and a three-active tail (rounds 9 to 10), totalling 40 active S-boxes of weight 80; each active cell shows its nibble difference. The bulk is where trail multiplicity accumulates; the head and tail are finite-round transients.

5.1 Linear-layer ablations

To isolate the effect of the mixing matrix, we hold the S-box and the cell permutation fixed and vary only the 4×4 binary mixing matrix M. Table 3 collects the four ablations; all entries use the same CRAFT/PRESENT-type S-box ($w_{\min} = 2$) and the same permutation.

The ablation (Figure 6) makes the branch-irrelevance quantitative. At branch number 2, the gap ranges from 18.1 to 30.4 bits, a 12.3-bit spread, while the branch-4 variant J-I attains only

[2]For CRAFT, $\mu = 8$ and $\log_2 \rho = -4.02$; the head and tail replace four bulk-active S-boxes by three-active ones, so $2 \cdot 3 + 7 \cdot 4 + 2 \cdot 3 = 40$ active S-boxes (weight 80) versus $11 \cdot 4 = 44$ (weight 88), a difference of 4 active S-boxes $\times w_{\min} = 8$ bits. The remaining 5.3 bits are the differential-side discrepancy between the fixed-input mass $-\log_2 2^{-49.559} = 49.559$ and the bulk prediction $-11 \log_2 \rho = 44.22$. This split is arithmetic, not a proof of causal attribution: the 5.3-bit remainder is consistent with finite-round head/tail endpoint effects, but establishing that they are its sole cause would require an explicit decomposition of the transition-matrix product.

Variant	M	$\mathcal{B}$	core	ρ	slope $\mu + \log_2 \rho$	gap (R=11)
identity	I_4	2	1-active	1 ($\pi_{\max} = 2^{-3.907}$)	2.0	18.1
SKINNY-MC	$\begin{pmatrix} 1&0&1&1 \\ 1&0&0&0 \\ 0&1&1&0 \\ 1&0&1&0 \end{pmatrix}$	2	1-active	1.0	2.0	≈ 18.1
CRAFT-MC	$\begin{pmatrix} 1&0&1&1 \\ 0&1&0&1 \\ 0&0&1&0 \\ 0&0&0&1 \end{pmatrix}$	2	4-active×3	0.061628	3.98	30.4
J-I	$\begin{pmatrix} 0&1&1&1 \\ 1&0&1&1 \\ 1&1&0&1 \\ 1&1&1&0 \end{pmatrix}$	4	2-active	0.185826	1.572	15.4

Table 3: Linear-layer ablations. Same S-box and permutation throughout; only the mixing matrix varies. The slope column is the core-level quantity $\mu + \log_2 \rho$ of Theorem 1 (the same definition for all four), and the gap column is the finite-round fixed-input gap at $R = 11$ (for identity and SKINNY-MC this is the S-box baseline, for CRAFT-MC it includes the three-active head and tail). The identity and SKINNY-MC rows are degenerate row-stochastic cases ($\rho = 1$), so their slope entries report the finite-round value rather than the asymptotic slope of Theorem 1.

15.4 bits—below two of the branch-2 layers. The decisive ingredient is the recurrent-core structure: identity and SKINNY-MC induce only one-active cycles and stay at the S-box baseline; CRAFT-MC induces a four-active fixed point whose spectral radius 0.061628 amplifies the gap by 12.3 bits over the baseline; J-I induces a two-active alternating core whose spectral radius 0.185826 dampens the gap below the baseline. No function of $\mathcal{B}$ alone can reproduce this ordering.

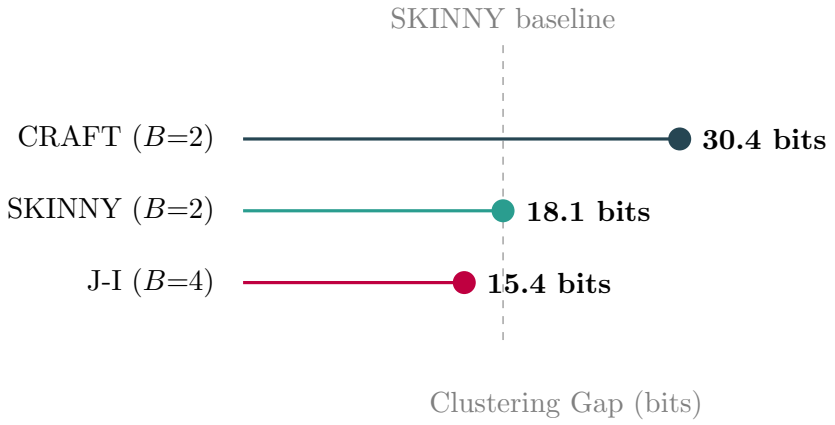


Figure 6: Branch number is not sufficient: three layers with the same S-box and permutation span a 12.3-bit range of clustering gap while the branch number is held at 2 or increased to 4.

Symmetry. All three of CRAFT's four-active cores have $\rho = 0.061628$, as predicted by Lemma 5.

Validation. For GIFT-64 the known input $\{5 : 6, 13 : c\}$ reproduces the published gap 0.585 bits, confirming the correctness of the machinery. The right Perron vector's top seed $\{12 : 2, 13 : 2\}$ attains 1.044 bits, strictly larger: the Perron vector locates an input whose hull potential exceeds the previously reported single-trail example, illustrating concretely that single-trail results may under-estimate the hull potential.

Automated input location and improved distinguishers. Finally, we record a concrete consequence of the automated search. The Perron vector of the four-active core locates, without

manual intervention, the input difference 0x0055000000550555; enumerating its optimal-pattern class with the mass DP yields a differential of $2^{-49.559}$, a *new* input/output difference pair, higher than the $2^{-49.79}$ reported by Hadipour et al. for their optimal input; the corresponding data requirement drops from about $2^{49.79}$ to $2^{49.559}$ chosen plaintexts. Extending the same core pattern to 12, 13, and 14 rounds gives probabilities $2^{-53.709}$, $2^{-57.768}$, and $2^{-61.840}$, each higher than the corresponding published value ($2^{-54.48}$, $2^{-59.13}$, $2^{-63.81}$). The tool replaces manual input selection by a spectral computation and recovers the asymptotically optimal input automatically, which is the concrete value a hull-clustering tool should deliver.

Table 4: Best previously published single-tweak differential distinguishers on CRAFT. Cui et al. report a congruent-differential cluster (a different object from the ordinary hull), whose difference pair is not restated here. Rounds 1 to 8 have no published single-tweak differential in the peer-reviewed literature.

R	Δ_{in}	Δ_{out}	$-\log_2 p$	source
9	0x00AA000000AAA0AA	0xAA0000000000000A0	40.20	Hadipour et al., ToSC 2019
10	0xAA000000AA00AAA0	0xAA0000000000000A0	44.89	Hadipour et al., ToSC 2019
10	(see paper)	(see paper)	42.32	Cui et al., TIFS 2024
11	0x00AA000000AAA0AA	0xAA0000000000000A0	49.79	Hadipour et al., ToSC 2019
12	0xAA000000AA00AAA0	0xAA0000000000000A0	54.48	Hadipour et al., ToSC 2019
13	0x00AA000000AAA0AA	0xAA0000000000000A0	59.13	Hadipour et al., ToSC 2019
14	0xAA000000AA00AAA0	0xAA0000000000000A0	63.81	Hadipour et al., ToSC 2019

Table 5: Our new single-tweak differential distinguishers on CRAFT. The input and output differences are new, and each probability is higher than the best previously published value of Table 4. Time is wall-clock for one run of the Python value-closure DP with the C++ BFS accelerator on a single desktop CPU core.

R	Δ_{in}	Δ_{out}	$-\log_2 p$	improvement (bits)	time (s)
11	0x0055000000550555	0x00aa000000000a00	49.559	+0.23	126.9
12	0x0055000000550555	0x00aa000000000a00	53.709	+0.77	158.9
13	0x0055000000550555	0x00aa000000000a00	57.768	+1.36	173.1
14	0x0055000000550555	0x00aa000000000a00	61.840	+1.97	194.1

Finite-round transient. At realistic attack round counts (R between 8 and 12) the measured gap lies a few bits below or above the asymptotic prediction $R(k\, w_{\min} + \log_2 \rho)$, reflecting the $O(1)$ term of Theorem 1. For GIFT-64 the gap is 0.585 bits at $R = 9$, small because the 32-state core is nearly deterministic; for CRAFT the gap is already 30.4 bits at $R = 11$, large because the four-active fixed point sustains a rich trail multiplicity. In every case the *ordering* of the gaps across linear layers—not the exact constant—is stable at these round counts, and it is precisely this ordering that the design guidelines rely on. The experiments therefore support the core claim at finite rounds: the branch number does not rank the hull risk, while the recurrent-core spectral slope $\mu + \log_2 \rho$ does.

5.2 Validation of the finite-round error bound

Throughout, "exact" refers to the core model of Section 4: the finite-round enumeration is exact for the trails that stay inside the k-active class (with the final round unrestricted), not for the full SPN over the whole state space. This is the object to which the asymptotic prediction is compared.

Table 6 reports, for each recurrent core, the spectral radius ρ, the multiplicity m of the dominant magnitude, the largest non-dominant eigenvalue λ_*, the spectral gap $\gamma = \lambda_*/\rho$, and the condition number $\kappa_\infty(V)$. The last two columns compare the actual error at $r = 8$ with the bound of Proposition 1. In every case the actual error is below the theoretical bound, closing the validation loop.

Cipher	n	ρ	m (period)	λ_*	$\gamma = \lambda_*/\rho$	exact ε_8	Schur bound (r=8)
CRAFT	2830	0.061628	1	0.041845	0.679	$6.1 \cdot 10^{-13}$	$7.4 \cdot 10^{-7}$
GIFT-64	32	0.038273	4	0.035479	0.927	$9.2 \cdot 10^{-13}$	$1.7 \cdot 10^{-10}$
PRESENT-80	240	0.069049	2	0.048048	0.696	$9.4 \cdot 10^{-11}$	$3.3 \cdot 10^{-7}$

Table 6: Finite-round error at $r = 8$. $\varepsilon_8 = \|N^q\|_{\max}$ is the exact error of Proposition 1 (with $q = 8/t$ for period t), and the last column is the Schur bound of Eq. (15). The Schur bound is 10^{14}–10^{16} times tighter than the eigenvector-conditioning bound.

Two observations follow. First, the exact error is minuscule—at most $9.4 \cdot 10^{-11}$ at $r = 8$—so the spectral model reproduces the finite-round differential to many decimal places for all three cores. Second, the Schur bound of Eq. (15) exceeds the exact error by only 10^2–10^6, because its constant is $\|Q\|_\infty \|Q^\top\|_\infty$ of order n rather than the eigenvector condition number; the latter is 10^{14}–10^{16} times the exact error and hence vacuous. For the periodic cores (GIFT-64, $t = 4$; PRESENT-80, $t = 2$) the quantities are computed on the period-adjusted matrix $T_{\mathcal{C}}^t$.

5.3 A design-time screening workflow

The controlled ablation above is a passive observation; we now reframe it as an active design decision to make the guidelines operational. Consider a designer who must select a 4×4 binary mixing matrix M for a 64-bit nibble SPN, keeping the S-box and the cell permutation fixed, and who restricts attention to candidates of branch number 2. Two natural candidates are CRAFT's mixing matrix M_{CRAFT} and SKINNY's mixing matrix M_{SKINNY}, both evaluated under the same permutation. A screening that checks only the branch number cannot distinguish them: both have $\mathcal{B} = 2$.

Running `coretool` on the two candidates (rows "CRAFT" and "SKINNY-MC (abl.)" of Table 2) exposes the difference immediately. M_{CRAFT} induces a four-active fixed point and three recurrent cores, all of spectral radius $\rho = 0.061628$ and per-round slope 3.980 bits, giving a gap of 30.4 bits at $R = 11$. M_{SKINNY} induces only one-active cycles of spectral radius 1 (the row-stochastic baseline), giving a gap of about 18.1 bits at the same round count. The 12.3-bit difference is attributable to the fixed-point core alone, at identical branch number. A designer following Guideline 2 therefore rejects M_{CRAFT} in favour of a mixing matrix without a low-weight fixed-point core, or proceeds to a more careful hull analysis of any matrix that `coretool` flags.

6 Limitations and Discussion

Asymptotic versus finite-round. Theorem 1 is asymptotic. The $O(1)$ transient is not estimated a priori; it is quantified by the exact finite-round enumeration of Module M3. For round counts

that are realistic for attacks ($R \leq 12$) the transient is a few bits and does not invert the ordering of cores, but the asymptotic slope should not be quoted as an exact finite-round gap.

The DDT-product model. The theory analyses the normalised path-sum $T_{\mathcal{C}}^{R}$ (equivalently, the integer path-sum (A^R) of Eq. (4)), which is the standard DDT-product object of differential cryptanalysis. Equality of this object with the exact cipher differential probability requires the stochastic-equivalence assumption [12]; we do not prove this equality. This assumption is not specific to our work: the optimal-pattern enumeration of [1], the correlation-matrix iteration of [2], the congruent-cluster analysis of [3], and every MILP-, SAT- and DDT-based differential search adopt the same DDT-product convention. Throughout the paper, "exact" therefore refers to the DDT-product model, not to the cipher.

Incompleteness of the search. Module M1 is not an exhaustive full-state search; it enumerates low-weight mask patterns. A core whose recurrent cycle visits masks of weight larger than the search bound is not reported. From the security viewpoint this is acceptable, since such cores carry exponentially small probability and cannot dominate a practical hull, but the bound must be stated whenever the tool is used.

The escape problem. The core-restricted analysis leaves open the gap to the full cipher. Unlike the substochastic core matrix $T_{\mathcal{C}}$ of the preceding sections (which is *not* a Markov chain and has spectral radius $\rho < 1$), the full transition matrix A over all difference states is row-stochastic, i.e. a genuine Markov chain with $\rho(A) = 1$. Writing it in block form $A = \left(\begin{smallmatrix} T_{\mathcal{C}} & B \\ C & D \end{smallmatrix}\right)$, where D collects the states outside the core and B, C are the transitions leaving and re-entering it, the characteristic polynomial factorises as

$$\det(zI - A) \;=\; \det(zI - D) \cdot \det\Big((zI - T_{\mathcal{C}}) - B\,(zI - D)^{-1} C\Big),$$

so that $B(zI-D)^{-1}C$ is exactly the leave-and-return correction ignored by the core theorem. This Schur-complement formulation makes the escape problem precise, but it does not yet give a useful upper bound: D contains the genuinely closed recurrent classes of the full chain, hence $\rho(D) = 1$; the per-round escape probability of the studied cores is large ($p_{\mathrm{esc}} \approx 0.75$–$0.98$ for CRAFT); and the Neumann expansion of $(zI - D)^{-1}$ diverges at the relevant $z \approx \rho$. The correction is nonetheless well defined: $z \approx \rho$ lies outside $\sigma(D)$ (the spectrum of D is concentrated near 1), so the resolvent $(zI - D)^{-1}$ exists and $B(zI - D)^{-1}C$ is a finite path-sum. The obstruction is therefore computational (D has 2^{64} states) rather than a matter of well-definedness. Bounding the correction term is therefore left as the central open problem of Section 8.

Empirically, this obstruction is not a shortcoming of the search: the unrestricted forward closure of a low-weight difference does not close. From a one-active seed the closure spreads to higher supports ($1 \to 3 \to 6$ active nibbles for CRAFT, $1 \to 3$ for GIFT-64, $1 \to 4$ for PRESENT-80) and exceeds $2 \cdot 10^5$ states without converging. This confirms that the low-weight cores are transient ($\rho < 1$), that the genuinely closed $\rho = 1$ recurrent classes of the full chain sit at high weight and are too large to enumerate, and—consistently with a mixing full differential on the nonzero space—that $-\log_2 p_{\mathrm{diff}}$ converges to a constant. The core-level slope $\mu + \log_2 \rho$ is therefore a finite-round transient of the restricted dynamics, not an asymptotic slope of the full cipher.

Scope of the framework. The analysis targets four-bit nibble-based SPNs. The core-based characterisation rests on the per-nibble DDT structure and the cell-wise linear layer, and it is not directly applicable to byte-oriented ciphers, whose 8-bit S-boxes and branch-number regimes behave differently.

Relation to active-S-box Markov truncation. A standard approach to differential hull estimation truncates the differential transition matrix to the states with at most K active S-boxes and iterates the truncated chain (an *active-S-box Markov truncation*). That approach yields finite-

round differential estimates and is complementary to ours: it approximates the finite-round mass over a bounded state space, whereas the core-spectral theorem characterises the *asymptotic* slope and, crucially, identifies *which structural feature of the linear layer* is responsible for the clustering. The two solve different problems—finite-round approximation versus asymptotic structural explanation—and can be combined: the truncated chain provides the finite-round constant, and the core-spectral root provides the design-time interpretation and criterion.

7 Design Guidelines

The thermodynamic reading of Theorem 1 (energy versus entropy) translates directly into design advice: a hull-resistant linear layer must keep both the ground-state energy and the topological pressure under control, and the branch number disciplines only the former.

1. **Branch number is necessary, not sufficient.** The branch number bounds the ground-state energy ($\mu \geq (\mathcal{B}/2)w_{\min}$, by Lemma 3) but not the topological pressure $\log_2 \rho$ (Corollary 1); Section 5 exhibits branch-2 layers whose gaps span 12.3 bits.

2. **Screen the linear layer for low-weight recurrent cores.** A core with small k and large $\log_2 \rho$ amplifies the hull; enumerate low-weight cores with `coretool` and flag those whose ρ is close to 1 (Section 5).

3. **Beware nibble-level fixed points and short cycles.** Nibble-mixing (XOR) maps that induce low-weight fixed points or short cycles create high-ρ cores; bit-permutation layers (GIFT, PRESENT) require active-class closure instead (Section 4).

   Operationalising the guidelines. Each guideline can be checked with a single invocation of the tool. Given a candidate linear layer, one runs the core search (Module M1) at the design mask bound and inspects the output: for every reported core, the tuple $(k, \rho, \mu + \log_2 \rho)$ is the hazard signature, and any core whose slope exceeds the target threshold is a candidate hazard. A layer with no low-weight core is baseline-safe, regardless of its branch number; a layer with a high-ρ fixed point is flagged for revision. The screening is automatic, takes seconds on a commodity machine for cores of a few thousand states, and does not require the designer to fix a particular input difference in advance.

8 Conclusion and Open Problems

We have shown that the differential trail space of a nibble SPN is a weighted subshift of finite type whose Ruelle transfer operator has spectral radius ρ (equivalently topological pressure $\log_2 \rho$), that the hull-clustering gap is the difference between the ground-state energy $R\mu$ of the optimal trail and the free energy of the hull, and that the branch number alone bounds only the ground state ($\mu \geq (\mathcal{B}/2)w_{\min}$), not the pressure. The variational principle makes this energy/entropy split rigorous—the pressure is the supremum of entropy minus energy over invariant measures—and is the structural root of the branch-number irrelevance. The characterisation is exact in the asymptotic regime, automated, and validated on four ciphers and on controlled ablations, and it yields three actionable design guidelines. On CRAFT, the automated tool further delivers *new* differential distinguishers for 11 to 14 rounds, each with probability higher than that of the best previously published distinguisher for the same round count, and reports, to our knowledge, the first single-tweak differentials for 1 to 8 rounds.

Several directions remain open. First, the $O(1)$ transient of Theorem 1 deserves a quantitative bound in terms of the spectral gap of $T_{\mathcal{C}}$; its periodic part is now identified as the discrete Fourier transform of the cyclic classes, but a bound on the full transient would turn the asymptotic law into a finite-round estimate with provable error. Second, a formal bound on the total probability contributed by the omitted high-weight paths would turn the low-weight heuristic of Section 4 into a completeness certificate; an adaptive mask bound (raising the weight until the reported maximal-ρ core stabilises) is an empirical step in this direction. Third, extending the core-spectral framework beyond 4-bit nibbles to 8-bit S-boxes would require handling the much larger per-cell DDT and is a natural but nontrivial generalisation. Fourth, a systematic classification of binary mixing matrices by their recurrent-core structure (fixed points, short cycles, and their spectral radii) would turn Guideline 3 into a concrete design catalogue. Fifth, the Perron-vector input-location phenomenon observed on GIFT-64 suggests a general method for finding inputs with maximal hull potential, which may be of independent interest for attack construction. Sixth, the related-work comparison of Table 1 suggests that a unified treatment of the optimal-pattern, congruent-cluster, and recurrent-core viewpoints (each a different coarse-graining of the same trail space) could yield a single hierarchy of trail-space decompositions, ordered by how much of the hull they capture. Seventh, the fixed-pair gap of Eq. (5) and the full-cipher extension of Theorem 1 (bounding the trails that leave the core) require per-pair endpoint analysis and a global completeness bound, respectively; these are the two most pressing steps for turning the restricted-core theory into a complete-cipher statement.

References

[1] H. Hadipour, S. Sadeghi, M. M. Niknam, L. Song, and N. Bagheri. Comprehensive security analysis of CRAFT. *IACR Transactions on Symmetric Cryptology*, 2019(4):290–317, 2019.

[2] H. Liu, W. Zhang, J. Zhang, and X. Sun. Clustering of differentials in CRAFT with correlation matrices. *International Journal of Intelligent Systems*, 37(12):12113–12134, 2022.

[3] T. Cui, Y. Mao, Y. Yang, Y. Zhang, J. Zhang, and C. Jin. Congruent differential cluster for binary SPN ciphers. *IEEE Transactions on Information Forensics and Security*, 19:2385–2397, 2024.

[4] G. Leurent, C. Pernot, and A. Schrottenloher. Clustering effect in Simon and Simeck. In *Advances in Cryptology – ASIACRYPT 2021*, volume 13090 of *LNCS*, pages 272–302. Springer, 2021.

[5] A. Biryukov, L. Cardoso dos Santos, D. Feher, V. Velichkov, and G. Vitto. Automated truncation of differential trails and trail clustering in ARX. In *Selected Areas in Cryptography – SAC 2021*, volume 13203 of *LNCS*, pages 286–307. Springer, 2022.

[6] K. Sakamoto, R. Ito, and T. Isobe. Parallel SAT framework to find clustering of differential characteristics and its applications. In *Selected Areas in Cryptography – SAC 2023*, volume 14201 of *LNCS*, pages 409–428. Springer, 2024.

[7] S. Banik, S. K. Pandey, T. Peyrin, Y. Sasaki, S. M. Sim, and Y. Todo. GIFT: A small present. In *Cryptographic Hardware and Embedded Systems – CHES 2017*, volume 10529 of *LNCS*, pages 321–345. Springer, 2017.

[8] A. Bogdanov, L. R. Knudsen, G. Leander, C. Paar, A. Poschmann, M. J. B. Robshaw, Y. Seurin, and C. Vikkelsoe. PRESENT: An ultra-lightweight block cipher. In *Cryptographic Hardware and Embedded Systems – CHES 2007*, volume 4727 of *LNCS*, pages 450–466. Springer, 2007.

[9] C. Beierle, J. Jean, S. Kölbl, G. Leander, A. Moradi, T. Peyrin, Y. Sasaki, P. Sasdrich, and S. M. Sim. The SKINNY family of block ciphers and its low-latency variant MANTIS. In *Advances in Cryptology – CRYPTO 2016*, volume 9815 of *LNCS*, pages 123–153. Springer, 2016.

[10] J. Daemen, R. Govaerts, and J. Vandewalle. Correlation matrices. In *Fast Software Encryption – FSE 1994*, volume 1008 of *LNCS*, pages 275–285. Springer, 1995.

[11] L. R. Knudsen. Truncated and higher-order differentials. In *Fast Software Encryption – FSE 1994*, volume 1008 of *LNCS*, pages 196–211. Springer, 1995.

[12] X. Lai, J. L. Massey, and S. Murphy. Markov ciphers and differential cryptanalysis. In *Advances in Cryptology – EUROCRYPT 1991*, volume 547 of *LNCS*, pages 17–38. Springer, 1991.

[13] E. Biham and A. Shamir. Differential cryptanalysis of DES-like cryptosystems. *Journal of Cryptology*, 4(1):3–72, 1991.