AICR · OPEN PROBLEM

Logarithmic Independence Implies Pseudorandomness for Local Permutations

Constant-locality round permutations at independence order $\log n$

Status. Statement: AI-written from Jesko Dujmovic, Angelos Pelecanos, Stefano Tessaro, *When Simple Permutations Mix Poorly: Limited Independence Does Not Imply Pseudorandomness*, Cryptology ePrint Archive 2026, not yet checked by a human. The original constant-order conjecture (order 4, and in fact every constant order k) is refuted by the source paper's explicit construction, while the logarithmic-order version stated here is untouched: the paper offers neither a proof nor a counterexample, and states that an unconditional proof would imply that one-way functions exist.

Category. *Symmetric Cryptography & Pseudorandomness.*

Conjecture (informal). *Many block cipher security proofs do not use computational assumptions at all: they show that the cipher's outputs on any few distinct inputs, under one random key, are statistically indistinguishable from a random tuple of distinct strings. It has long been asked whether this kind of limited-independence guarantee, for round-based ciphers whose single round is very simple, is enough to conclude full pseudorandomness. "Very simple" here means local: every output bit of one round depends on only a bounded number of input bits, with the bound not growing with the block length. The paper this problem comes from settles the question negatively when the number of inputs is a fixed constant: it exhibits a local round permutation whose repeated composition is statistically close to k-wise independent for whatever constant k you like, and yet is broken by an efficient attack that queries only $k+1$ points. The construction works by smuggling a small number of predictable invariants through the rounds and reading them off from a bounded number of wires. The conjecture recorded here is what survives that attack: raise the independence order from a constant to logarithmic in the block length, keep everything else the same, and the implication is conjectured to hold again. The authors' reason for believing it is that their counterexample would now have to sustain a growing number of invariants at once, which the technique cannot do. An unconditional proof is out of reach for the usual reasons — it would imply that one-way functions exist — so the realistic targets are a proof under standard assumptions or a genuinely new counterexample.*

1 The setting

Security proofs for block ciphers that avoid computational assumptions usually target *k-wise independence*: one shows that the cipher's outputs on any k distinct inputs, under a single random key, are statistically close to a uniformly random tuple of k distinct strings. Already $k = 2$ rules out linear and differential cryptanalysis, and results of this shape are known for random local circuits [Gow96, BH08, GHP25, GHKO25] and for substitution–permutation networks [LTV21].

Hoory, Magen, Myers and Rackoff [HMMR04, HMMR05] proposed reading limited independence as evidence of full pseudoran-

domness. Let $\mathcal{P}$ be a randomized permutation of $\{0,1\}^n$ that is *local*, meaning each output bit depends on only a constant number of input bits, and let $\mathcal{P}^T$ denote the composition of T independent copies of $\mathcal{P}$. Their conjecture: if $\mathcal{P}^T$ is negligibly far from 4-wise independent, then $\mathcal{P}^T$ is a pseudorandom permutation. The constant 4 is close to minimal, since counterexamples are known at 3.

The source paper [DPT26] refutes that conjecture. For every constant k, it constructs a local randomized permutation over a constant-size alphabet (of size depending on k) whose $(n-2)$-fold composition is negligibly far from k-wise independent, and yet is distinguished from a random permutation, with constant probability, by an efficient adversary querying only $k+1$ points. The construction carries a unary round counter inside the state, drives $k+1$ families of "special" strings deterministically around cycles of length $n-2$ (one cycle per special value), while a conditional brickwork circuit randomizes everything else, and in the last round sends the $k+1$ strings whose counter has run out to $k+1$ values obeying one linear relation that a random permutation would not obey.

What that counterexample consumes is a constant amount of structure: a constant number of invariants, each detectable from a constant number of wires. The conjecture below is the variant the authors single out as untouched by their attack. Raising the order of independence from a constant to $\log n$ would force $\omega(1)$ invariants to be maintained at once, which they say their approach cannot do. The authors also record that a proof is not to be expected unconditionally, since it would imply the existence of one-way functions and that P $\neq$ NP; they suggest proofs under non-trivial assumptions, or a refutation, as the realistic targets. They further note that Gowers' conjecture [Gow96], made explicit by Barak [Bar17] and connected to obfuscation candidates by Canetti et al. [CCMR24], is weaker, in the sense that refuting it would also refute the conjecture stated here.

Progress to date. For each constant k, the paper's main theorem constructs a constant-locality randomized permutation over a constant-size alphabet (whose size grows with k) whose $(n-2)$-fold composition is negligibly far from k-wise independent while admitting an efficient $(k+1)$-query distinguisher, which refutes the constant-order form of the implication. Towards the logarithmic-order form it supplies only an obstruction: the counterexample

depends on distinguishing a constant number of invariants from a constant number of wires, and the authors state that this approach is not viable once $\omega(1)$ invariants must be maintained. The paper additionally observes that no unconditional proof of this or its sibling conjectures should be expected, since such a proof would imply the existence of one-way functions and P $\neq$ NP.

Why it matters. The whole point of proving k-wise independence for a cipher is that one hopes it is evidence of something stronger. The source paper shows that at any constant order the hope is formally unfounded, which leaves the field without a threshold at which the evidence becomes conclusive. This conjecture proposes $\log n$ as that threshold. A proof, even under standard assumptions, would turn existing $\log n$-wise independence results for random local circuits into pseudorandomness results and would give the statistical-proof programme a genuine endpoint. A refutation would be at least as interesting: it would need a counterexample maintaining a growing number of locally checkable invariants, which is exactly the technique the paper says it does not have, and the same object would likely say something about how much structure can survive a mixing process.

2 Notation and parameters

- n — the block length in bits. All asymptotics are in n; a function ε is *negligible*, written $\varepsilon(n) = \mathrm{negl}(n)$, if for every constant $c > 0$ we have $\varepsilon(n) < n^{-c}$ for all sufficiently large n.
- $[n] := \{1, \dots, n\}$.
- ℓ — the locality parameter, a constant independent of n.
- $\mathcal{P} = \{\mathcal{P}_n\}_{n \in \mathbb{N}}$ — the one-round randomized permutation family (Definition 1).
- $T = T(n)$ — the number of rounds; $\mathcal{P}^T$ is the T-fold sequential composition with fresh independent randomness in every round (Definition 2).
- $k(n) = \lfloor \log_2 n \rfloor$ — the order of independence at which the implication is conjectured.

- $A^{(k)}$ — the set of ordered k-tuples of pairwise distinct elements of a set A; $U_{n,k}$ — the uniform distribution on $(\{0,1\}^n)^{(k)}$.
- d_{TV} — total variation distance, $d_{\mathrm{TV}}(p,q) = \frac{1}{2}\sum_\omega |p(\omega) - q(\omega)|$.
- Π — a uniformly random permutation of $\{0,1\}^n$.

3 The conjecture

Definition 1 (ℓ-local randomized permutation family). A *randomized permutation* on $\{0,1\}^n$ is an algorithm $\mathcal{P}_n$ taking an input $v \in \{0,1\}^n$ and randomness r, and returning $\mathcal{P}_n(v;\mathsf{r}) \in \{0,1\}^n$, such that for every fixed value of r the map $v \mapsto \mathcal{P}_n(v;\mathsf{r})$ is a permutation of $\{0,1\}^n$. Equivalently, sampling r uniformly makes $\mathcal{P}_n$ a distribution over permutations of $\{0,1\}^n$; we write $\mathcal{P}_n(v)$ for $\mathcal{P}_n(v;\mathsf{r})$ with r uniform.
For $\ell \in \mathbb{N}$, such a $\mathcal{P}_n$ is *ℓ-local* if for every value of r and every $j \in [n]$ there is a set $S \subseteq [n]$ with $|S| \le \ell$ such that the j-th output bit of $\mathcal{P}_n(\cdot;\mathsf{r})$ is a function of the input bits indexed by S alone. A family $\mathcal{P} = \{\mathcal{P}_n\}_{n\in\mathbb{N}}$ is ℓ-local if every $\mathcal{P}_n$ is ℓ-local for the same constant ℓ.

Definition 2 (sequential composition). For $T \in \mathbb{N}$ and a randomized permutation $\mathcal{P}_n$, the T-fold sequential composition $\mathcal{P}^T$ is the randomized permutation

$$\mathcal{P}^T : x \mapsto \big(\mathcal{P}_n(\cdot;\mathsf{r}_T) \circ \cdots \circ \mathcal{P}_n(\cdot;\mathsf{r}_1)\big)(x),$$

where $\mathsf{r}_1,\dots,\mathsf{r}_T$ are independent and uniform. For a function $T = T(n)$ this defines a family $\{\mathcal{P}^{T(n)}\}_{n\in\mathbb{N}}$ of randomized permutations, again written $\mathcal{P}^T$.

Definition 3 (negligibly far from $k(n)$-wise independence). A family $Q = \{Q_n\}_{n\in\mathbb{N}}$ of randomized permutations on $\{0,1\}^n$ is *negligibly far from $k(n)$-wise independent* if there is a negligible ε such that for every n and every tuple of pairwise distinct

inputs $(x^1,\dots,x^{k(n)}) \in (\{0,1\}^n)^{(k(n))}$,

$$d_{\mathrm{TV}}\big((Q_n(x^1),\dots,Q_n(x^{k(n)})),\ U_{n,k(n)} \big) \ \le\ \varepsilon(n),$$

where the $k(n)$ outputs are produced by one single sample of Q_n, so that the tuple lies in $(\{0,1\}^n)^{(k(n))}$.

Definition 4 (pseudorandom permutation). A family $Q = \{Q_n\}_{n\in\mathbb{N}}$ of randomized permutations on $\{0,1\}^n$ is a *pseudorandom permutation* if for every probabilistic oracle algorithm D running in time polynomial in n,

$$\Big|\ \Pr\big[D^{Q_n}(1^n)=1\big] \ -\ \Pr\big[D^{\Pi}(1^n)=1\big]\ \Big| \ =\ \mathrm{negl}(n),$$

where in the first probability a single sample of Q_n is fixed and D may query it on inputs of its choice, in the second Π is a uniformly random permutation of $\{0,1\}^n$, and in both cases D is given forward evaluation queries only.

Conjecture 1 (logarithmic independence). *Let $\ell \in \mathbb{N}$ be a constant, let $\mathcal{P} = \{\mathcal{P}_n\}_{n\in\mathbb{N}}$ be an ℓ-local randomized permutation family in the sense of Definition 1, and let $T : \mathbb{N} \to \mathbb{N}$ be arbitrary, with $\mathcal{P}^T$ the $T(n)$-fold sequential composition of Definition 2. Set $k(n) := \lfloor \log_2 n \rfloor$.*
If $\mathcal{P}^T$ is negligibly far from $k(n)$-wise independent in the sense of Definition 3, then $\mathcal{P}^T$ is a pseudorandom permutation in the sense of Definition 4.

4 Bibliography

[Bar17] B. Barak. *The complexity of public-key cryptography.* Tutorials on the Foundations of Cryptography: Dedicated to Oded Goldreich, pages 45–77, Springer, 2017.

[BH08] A. Brodsky, S. Hoory. *Simple permutations mix even better.* Random Structures and Algorithms, 32(3):274–289, 2008.

[CCMR24] R. Canetti, C. Chamon, E. R. Mucciolo, A. E. Ruckenstein. *Towards general-purpose program obfuscation via local mixing.* TCC 2024, Part IV, volume 15367 of LNCS, pages 37–70, Springer, 2024.

[DPT26] J. Dujmovic, A. Pelecanos, S. Tessaro. *When Simple Permutations Mix Poorly: Limited Independence Does Not Imply Pseudorandomness.* Cryptology ePrint Archive, Paper 2025/2282, 2026. [UNVERIFIED]

[GHKO25] W. Gay, W. He, N. Kocurek, R. O'Donnell. *Pseudorandomness properties of random reversible circuits.* CRYPTO 2025, Part I, volume 16000 of LNCS, pages 651–678, Springer, 2025.

[GHP25] L. Gretta, W. He, A. Pelecanos. *More efficient approximate k-wise independent permutations from random reversible circuits via log-Sobolev inequalities.* 36th SODA, pages 5582–5598, ACM-SIAM, 2025.

[Gow96] W. T. Gowers. *An almost m-wise independent random permutation of the cube.* Combinatorics, Probability and Computing, 5:119–130, 1996.

[HMMR04] S. Hoory, A. Magen, S. Myers, C. Rackoff. *Simple permutations mix well.* ICALP 2004, volume 3142 of LNCS, pages 770–781, Springer, 2004.

[HMMR05] S. Hoory, A. Magen, S. A. Myers, C. Rackoff. *Simple permutations mix well.* Theoretical Computer Science, 348(2–3):251–261, 2005.

[LTV21] T. Liu, S. Tessaro, V. Vaikuntanathan. *The t-wise independence of substitution-permutation networks.* CRYPTO 2021, Part IV, volume 12828 of LNCS, pages 454–483, Springer, 2021.