

Equivalence Classes of BOGI-Based Ciphers for Differential and Linear Cryptanalysis

Insung Kim¹ , Seonggyeom Kim² , Sunyeop Kim^{1,3} ,
Donggeun Kwon⁴ , Byoungjin Seok⁵ , Deukjo Hong⁶ ,
Jaechul Sung⁷ , Seokhie Hong⁸ , Sangjin Lee¹ and Dongjae Lee⁹ 

¹ Korea University, Seoul, Republic of Korea

² Samsung Electronics, Hwaseong, Republic of Korea

³ Nanyang Technological University, Singapore, Singapore

⁴ Kunsan National University, Gunsan, Republic of Korea

⁵ Hansung University, Seoul, Republic of Korea

⁶ Jeonbuk National University, Jeonju, Republic of Korea

⁷ University of Seoul, Seoul, Republic of Korea

⁸ SmartM2M, Busan, Republic of Korea

⁹ Kangwon National University, Chuncheon, Republic of Korea

Abstract

BOGI-based ciphers extend the design space of **GIFT** by combining 4-bit S-boxes with bit permutations satisfying the “Bad Output must go to Good Input” principle. Prior work reduced this space to 41,472 parameter representatives, but did not determine whether their complete differential and linear trail spaces were distinct. We define DC/LC-equivalence in terms of weight-preserving bijections between the differential and linear trail sets of two ciphers for an arbitrary number of rounds, and give sufficient conditions based on permutation characteristics and trail reversal. For each pair of mixing permutations and each 4-bit permutation, we decide exactly whether the required initial word permutation exists. The relations generated by these transformations and trail reversal partition the 41,472 ciphers into 864 classes of size 48 for BOGI-64 and 5,184 classes of size 8 for BOGI-128. Using the BOGI-128 classification, we perform differential and linear best-trail searches through round 20, completing both searches for 5,081 classes. These results and additional threshold decisions show that the earliest round at which both best-trail weights reach 128 bits is round 19, attained by at least 59 classes. The class containing **GIFT**-128 reaches both thresholds at round 22. Thus, at least 59 classes reach both thresholds three rounds earlier than the **GIFT**-128 class. We also compare selected BOGI-64 and BOGI-128 instances with **GIFT** in hardware and software. The resulting security and implementation data can support component selection in future **GIFT**-based primitives.

Keywords: GIFT, BOGI-based ciphers, equivalence relations, differential cryptanalysis, linear cryptanalysis

Emails: cmcom35@korea.ac.kr (Insung Kim), jeffgyeom@gmail.com (Seonggyeom Kim), kin3548@gmail.com (Sunyeop Kim), dgkwon@kunsan.ac.kr (Donggeun Kwon), bjseok@hansung.ac.kr (Byoungjin Seok), deukjo.hong@jbnu.ac.kr (Deukjo Hong), jcsung@uos.ac.kr (Jaechul Sung), shhong@smartm2m.co.kr (Seokhie Hong), sangjin@korea.ac.kr (Sangjin Lee), dongjae.lee@kangwon.ac.kr (Dongjae Lee)

This work is licensed under a “CC BY 4.0” license.

Date of this document: 2026-09-03.



1 Introduction

GIFT is a lightweight block cipher whose round function combines a 4-bit S-box with a bit permutation [BPP⁺17]. The cipher has since served as a building block for several lightweight constructions. GIFT-128 is the underlying primitive of the authenticated encryption designs SUNDAE-GIFT [BBP⁺19] and GIFT-COFB [BCI⁺20], and GIFT-COFB was a finalist in the NIST Lightweight Cryptography standardization process [TMC⁺23]. The tweakable block ciphers TGIF-TBC [IKM⁺19] and TweGIFT [CDJ⁺21] likewise derive their internal primitives from GIFT.

The use of GIFT in these constructions has been accompanied by continued research on its implementation and security. Fixslicing provides fast constant-time implementations of GIFT and GIFT-COFB on ARM Cortex-M [ANP20]. Recent cryptanalysis of reduced-round GIFT includes automated differential-linear attacks on GIFT-64 and GIFT-128, improved single-key and related-key differential attacks on GIFT-64, and a closer examination of the gap between trail probabilities and differential probabilities under related keys [WLHL24, DGMNP25, Sun26].

Resistance to differential cryptanalysis [BS91] and to linear cryptanalysis [Mat94] has a central role in the design of GIFT [BPP⁺17]. The S-box and bit permutation were selected jointly according to their one-bit DDT and LAT behavior so that the BOGI principle excludes consecutive single-active-bit transitions in both differential and linear characteristics. The designers also used estimates for differential characteristics and linear hulls to support their choices of 28 rounds for GIFT-64 and 40 rounds for GIFT-128 [BPP⁺17]. These arguments do not replace the analysis of other attack classes, but they motivate the focus on DC/LC-equivalence in this paper.

Preserving the same principle while varying both components gives the broader family of BOGI-based ciphers studied here [KHSH20, KHSH22]. In this family, the S-box determines the available local transitions and their weights, while the bit permutation determines how these transitions are connected across rounds. A search for stronger ciphers must therefore consider S-box and bit-permutation pairs, producing a much larger search space than either component family alone. Proving that two resulting ciphers have weight-preserving corresponding trails allows one representative to replace both in a best-trail search without omitting distinct trail behavior.

Prior work on GIFT and BOGI-based ciphers. Sun et al. fixed the GIFT S-box and varied the 16-bit group mapping of GIFT-64 [SPWW22]. Their GM-equivalence relation partitions the 2,304 group mappings into 168 classes. By searching one representative from each class, they identified a class containing 24 variants with improved differential resistance and comparable linear resistance to GIFT-64.

Kim et al. studied the broader BOGI family obtained from optimal BOGI-applicable S-boxes and LS-based bit permutations [KHSH20, KHSH22]. They reduced approximately $2^{37.09}$ parameter combinations to 41,472 representatives and searched all representatives through 13 rounds for BOGI-64 and through 11 rounds for BOGI-128. For BOGI-128, only the eight representatives attaining the best 11-round weights were extended through 22 rounds.

Kim’s thesis subsequently studied implementation-oriented selection in the same design space, including S-box cost and bitsliced mixing layers for the 24 representative LS permutations [Kim23]. These studies reduce the cost of exploring the design space, but they do not determine whether the complete differential and linear trail spaces of the 41,472 representatives are distinct. Best-trail searches may therefore be repeated for ciphers whose trails correspond with the same weights.

We address this gap by defining DC/LC-equivalence for complete BOGI- N ciphers and constructing the required trail bijections from permutation characteristics and trail

reversal. The resulting classification reduces the number of representatives required for best-trail search.

Our contributions. Our theoretical contributions are summarized as follows.

- We define DC/LC-equivalence between SbPN instances in terms of weight-preserving bijections between their complete differential and linear trail sets at any number of rounds. We give sufficient conditions for constructing these bijections using permutation characteristics and trail reversal.
- For each block size, we decide exactly whether the required initial word permutation exists over $\mathcal{W}_{16}$ or $\mathcal{W}_{32}$ for all 13,824 LS-pair and bit-permutation triples. The resulting relations partition the 41,472 BOGI ciphers into 864 classes of size 48 for BOGI-64 and 5,184 classes of size 8 for BOGI-128. The corresponding maps explicitly transport differential and linear trails between members of each class.

We apply these classifications to best-trail search and implementation-oriented selection.

- We search differential and linear best trails through round 20 for BOGI-128. Both searches were completed for 5,081 classes. Together with additional threshold decisions, these results show that the minimum round at which both best-trail weights reach 128 bits is round 19. At least 59 classes attain this value, while the class containing GIFT-128 reaches both thresholds at round 22.
- For BOGI-64, we evaluate the 24,576 implementations obtained from the two classes attaining the best reported 12-round weights. Under a common SKY130 synthesis flow, 97 are both smaller and faster than GIFT-64, and six remain after the joint area-delay comparison. The same six remain when the comparison includes the XE-translated variants derived from the GIFT S-box. We also evaluate the selected candidates using bitsliced and fixed-plane software implementations.
- For BOGI-128, we evaluate implementations drawn from the classes verified to attain the minimum joint threshold round. The LUT-based flow identifies implementations that improve both area and delay over GIFT-128. A separate PEIGEN-expression evaluation shows the tradeoff obtained when the S-box equations are fixed before synthesis. In a scalar software evaluation, the fastest tested BOGI-128 configuration outperforms GIFT-128 in both bitslicing and a fixslicing-style representation.

Organization. Section 2 gives the notation and the BOGI construction. Section 3 defines DC/LC-equivalence and gives the criteria used in the classification. Section 4 presents the results for BOGI-64 and BOGI-128. The remaining sections discuss applications and conclude the paper. Our artifact repository: <https://anonymous.4open.science/r/Equivalence-Classes-of-BOGI-Based-Ciphers>.

2 Preliminaries

This section fixes the notation for SbPN instances and their differential and linear trails, describes the BOGI construction, and states the reduced parameter set from which our classification starts.

2.1 Notation and the SbPN model

For a positive integer k , let $[k] := \{0, 1, \dots, k-1\}$, and let $\text{Sym}(k)$ denote the symmetric group on $[k]$. Unless stated otherwise, additions, inner products, and matrix operations are over $\mathbb{F}_2$. For an equivalence relation $\sim$ on a set $\mathcal{X}$, write

$$[X]_\sim := \{X' \in \mathcal{X} : X' \sim X\}, \quad \mathcal{X}/\sim := \{[X]_\sim : X \in \mathcal{X}\}.$$

For positive integers m and w , let $\mathcal{V}_{m,w} := (\mathbb{F}_2^w)^m$ and $N := mw$. A state $x \in \mathcal{V}_{m,w}$ is written as $x = (x[0], \dots, x[m-1])$, where $x[i] \in \mathbb{F}_2^w$ is its i th word. For $b \in [w]$, the notation $x[i]_b$ denotes bit b of word i . The pair (i, b) corresponds to flat bit position $wi + b \in [N]$, and x_c denotes the bit at flat position c .

For $p \in \text{Sym}(w)$ and $z \in \mathbb{F}_2^w$, we use the same symbol p for the induced bit permutation of a word, defined by $(p(z))_{p(b)} = z_b$. Its parallel application to q words is denoted by $\bar{p}^{(q)}$ and satisfies

$$(\bar{p}^{(q)}(x))[i] := p(x[i]), \quad i \in [q]. \quad (1)$$

Definition 1 (Bit and word permutations). For $\pi \in \text{Sym}(N)$, the bit permutation P_π moves the bit at flat position c to position $\pi(c)$, so $(P_\pi(x))_{\pi(c)} = x_c$. For $\omega \in \text{Sym}(m)$, the word permutation W_ω moves word i to word position $\omega(i)$ without changing its internal bit positions, so $(W_\omega(x))[\omega(i)] = x[i]$. We write

$$\mathcal{W}_m := \{W_\omega : \omega \in \text{Sym}(m)\}$$

for the set of all word permutations.

Every bit permutation is represented by a permutation matrix. Hence $P^{-1} = P^T$ and $P^{-T} = P$. The same bit permutation therefore connects consecutive rounds in both differential and linear trails.

Definition 2 (SbPN instance). Let $S : \mathbb{F}_2^w \rightarrow \mathbb{F}_2^w$ be a bijective S-box. Its parallel substitution layer is $\text{Sub}_S(x)[i] := S(x[i])$ for $i \in [m]$. An SbPN instance is a pair

$$X := (S, \Pi),$$

where Π is a bit permutation of $\mathcal{V}_{m,w}$. Its unkeyed round function is

$$R_X := \Pi \circ \text{Sub}_S.$$

For a round key $k \in \mathbb{F}_2^N$, let $\text{XOR}_k(x) := x + k$. Key additions do not change differences and affect only the signs, not the magnitudes, of linear correlations. We therefore describe trails using the unkeyed pair $X = (S, \Pi)$ under the Markov model with independent round keys. No claim is made about related-key trails or a fixed key schedule.

2.2 Differential and linear trails

Definition 3 (Propagation tables). For an S-box $S : \mathbb{F}_2^w \rightarrow \mathbb{F}_2^w$, define

$$\begin{aligned} \text{DDT}_S(a, b) &:= |\{x \in \mathbb{F}_2^w : S(x) + S(x+a) = b\}|, \\ \text{LAT}_S(u, v) &:= |\{x \in \mathbb{F}_2^w : \langle u, x \rangle = \langle v, S(x) \rangle\}| - 2^{w-1}, \\ \text{SQLAT}_S(u, v) &:= \text{LAT}_S(u, v)^2. \end{aligned}$$

Thus $\text{LAT}_S(u, v)/2^{w-1}$ is the correlation of the linear approximation (u, v) .

For a supported differential transition and a supported linear transition, define

$$\omega_D^S(a, b) := -\log_2 \left(\frac{\text{DDT}_S(a, b)}{2^w} \right), \quad (2)$$

$$\omega_L^S(u, v) := -\log_2 \left(\frac{\text{SQLAT}_S(u, v)}{2^{2w-2}} \right). \quad (3)$$

These are the negative logarithms of the local differential probabilities and linear potentials used in trail analysis [LMM91].

Definition 4 (Differential and linear trails). Let $X = (S, \Pi)$ be an SbPN instance and let $r \geq 1$. An r -round differential trail is a tuple

$$\tau = (a_0, b_0, \dots, a_{r-1}, b_{r-1}, a_r)$$

such that, for every $0 \leq t < r$, $a_{t+1} = \Pi(b_t)$ and

$$\text{DDT}_S(a_t[i], b_t[i]) > 0 \quad \text{for every } i \in [m].$$

Its set is denoted by $\mathcal{T}_D^r(X)$. An r -round linear trail is a tuple

$$\lambda = (u_0, v_0, \dots, u_{r-1}, v_{r-1}, u_r)$$

such that, for every $0 \leq t < r$, $u_{t+1} = \Pi(v_t)$ and

$$\text{SQLAT}_S(u_t[i], v_t[i]) > 0 \quad \text{for every } i \in [m].$$

Its set is denoted by $\mathcal{T}_L^r(X)$. The subsets with nonzero initial difference or mask are denoted by $\mathcal{T}_{D, \neq 0}^r(X)$ and $\mathcal{T}_{L, \neq 0}^r(X)$.

The corresponding trail weights are

$$\begin{aligned} \text{wt}_D(\tau) &:= \sum_{t=0}^{r-1} \sum_{i=0}^{m-1} \omega_D^S(a_t[i], b_t[i]), \\ \text{wt}_L(\lambda) &:= \sum_{t=0}^{r-1} \sum_{i=0}^{m-1} \omega_L^S(u_t[i], v_t[i]). \end{aligned}$$

The minimum nonzero-trail weights are

$$\begin{aligned} \text{MW}_D^r(X) &:= \min_{\tau \in \mathcal{T}_{D, \neq 0}^r(X)} \text{wt}_D(\tau), \\ \text{MW}_L^r(X) &:= \min_{\lambda \in \mathcal{T}_{L, \neq 0}^r(X)} \text{wt}_L(\lambda). \end{aligned}$$

2.3 BOGI-based ciphers

The “Bad Output must go to Good Input” principle was introduced in the design of GIFT [BPP⁺17]. For a bijective 4-bit S-box S , let e_i denote the unit vector at bit position i . An input position i is bad if $\text{DDT}_S(e_i, e_j) > 0$ or $\text{LAT}_S(e_i, e_j) \neq 0$ for some j , and an output position j is bad if the same condition holds for some i . We denote the resulting sets by $\mathcal{B}_{\text{in}}(S)$ and $\mathcal{B}_{\text{out}}(S)$, and their complements by $\mathcal{G}_{\text{in}}(S)$ and $\mathcal{G}_{\text{out}}(S)$.

A 16-bit mixing permutation follows the BOGI principle if no bit in $\mathcal{B}_{\text{out}}(S)$ enters a position in $\mathcal{B}_{\text{in}}(S)$ in the next S-box layer. An S-box is BOGI-applicable if such a mixing permutation exists. We use the optimal BOGI-applicable S-boxes of Kim et al., which have differential uniformity 6, linearity 8, and two good and two bad positions on each side [KSH20]. There are 2,654,208 such S-boxes [KSH20].




Figure 1: Substitution, mixing, and word-shuffle layers of a BOGI-64 round.

A 4×4 Latin square $L = (\ell_{r,b})$ over $[4]$ induces a 16-bit LS permutation by sending bit b of source word r to bit b of destination word $\ell_{r,b}$. There are 576 LS permutations [KHS22]. For $ib, ob \in \text{Sym}(4)$, define

$$M(L, ib, ob) := \overline{ob}^{(4)} \circ L \circ \overline{ib}^{(4)}.$$

The pair (ib, ob) is a BOGI pair for S if

$$(ob \circ ib)(\mathcal{B}_{\text{out}}(S)) \subseteq \mathcal{G}_{\text{in}}(S).$$

We denote the set of such pairs by $\text{BP}(S)$. Every optimal BOGI-applicable S-box has 96 BOGI pairs [KHS22].

For $N = 16g$, put $m = 4g$ and regard the state as g consecutive groups of four 4-bit words. Let $\text{Mix}_M^{(N)}$ apply M independently to these groups, and define

$$\text{Shuf}_N := W_{\sigma_N}, \quad \sigma_N(i) := g(i \bmod 4) + \left\lfloor \frac{i}{4} \right\rfloor.$$

The BOGI- N bit permutation and the corresponding SbPN instance are

$$\Pi_M^{(N)} := \text{Shuf}_N \circ \text{Mix}_M^{(N)}, \quad X_{S,M}^{(N)} := (S, \Pi_M^{(N)}).$$

For each block size N , the shuffle Shuf_N is fixed throughout the BOGI- N family. Cipher instances vary only in the S-box S and the LS-based mixing permutation M . Figure 1 shows this round structure for $N = 64$.

2.4 The starting parameter set

Kim et al. reduce the local parameter triples $(S, L, (ib, ob))$ in three steps [KHS20, KHS22], summarized in Table 1. First, the optimal BOGI-applicable S-boxes form 10,368 XE classes whose representatives have pairwise distinct DDTs [KHS20]. Equality of DDTs is equivalent to equality of SQLATs [DH19]. The same S-box family forms 20 PXE classes [KHS20]. This coarser partition is not an intermediate step in the reduction to 41,472 representatives. Second, a BOGI pair is absorbed into the S-box by

$$S' := ib \circ S \circ ob, \tag{4}$$

for which $(\text{id}, \text{id}) \in \text{BP}(S')$. Among the fixed XE representatives, 1,728 admit the identity pair. The complete trail correspondence associated with this step is stated in Appendix A. Third, for $p \in \text{Sym}(4)$, put

$$L' := \overline{p^{-1}}^{(4)} \circ L \circ \overline{p}^{(4)}, \quad ib' := p^{-1} \circ ib, \quad ob' := ob \circ p. \tag{5}$$

Table 1: Prior reductions leading to the starting parameter set.

Selected representatives	Number
XE representatives of optimal BOGI-applicable S-boxes	10,368
Representatives admitting the identity BOGI pair	1,728
LS representatives	24
Final S-box and LS pairs	$1,728 \cdot 24 = 41,472$

Then $M(L, ib, ob) = M(L', ib', ob')$, and the 576 LS permutations are represented by 24 choices [KHS22].

For $L = (\ell_{r,b})$ and $p \in \text{Sym}(4)$, the permutation

$$(\bar{p}^{(4)})^{-1} \circ L \circ \bar{p}^{(4)}$$

is induced by the Latin square whose (r, b) entry is $\ell_{r,p(b)}$; thus it only changes the order of the four columns. Since every row of a Latin square contains four distinct values, this permutation equals L only when $p = \text{id}$. We choose $\mathcal{L}_{\text{rep}}$ to contain one representative from each set of 24 column reorderings. Consequently, every LS permutation T determines unique $L_j \in \mathcal{L}_{\text{rep}}$ and $p \in \text{Sym}(4)$ such that

$$T = (\bar{p}^{(4)})^{-1} \circ L_j \circ \bar{p}^{(4)}.$$

Denote the 1,728 fixed S-box representatives by $S_0, \dots, S_{1727}$ and the 24 LS representatives by

$$\mathcal{L}_{\text{rep}} := \{L_0, \dots, L_{22}, L_{\text{GIFT}}\}.$$

For a supported block size N , define

$$\Pi_j^{(N)} := \text{Shuf}_N \circ \text{Mix}_{L_j}^{(N)}, \quad X_{s,j}^{(N)} := (S_s, \Pi_j^{(N)}),$$

and let

$$\mathcal{X}_N := \left\{ X_{s,j}^{(N)} : s \in [1728], L_j \in \mathcal{L}_{\text{rep}} \right\}.$$

Thus $|\mathcal{X}_N| = 41,472$. This is the starting set for the equivalence relations constructed in the following sections.

3 Trail equivalence via PXE and permutation characteristics

The set $\mathcal{X}_N$ is the starting point of our classification. We first recall the S-box correspondences and permutation characteristics used below, and then combine them to construct weight-preserving bijections between complete differential and linear trail sets.

3.1 S-box correspondences

Definition 5 (XE and PXE [Saa12, KHS20]). Let $S, S' : \mathbb{F}_2^w \rightarrow \mathbb{F}_2^w$ be bijective S-boxes. They are XOR-equivalent (XE) if there exist $c_{\text{in}}, c_{\text{out}} \in \mathbb{F}_2^w$ such that

$$S'(x) = S(x + c_{\text{in}}) + c_{\text{out}} \quad \text{for every } x \in \mathbb{F}_2^w.$$

They are permutation-XOR-equivalent (PXE) if there additionally exist $p_{\text{in}}, p_{\text{out}} \in \text{Sym}(w)$ such that

$$S'(x) = p_{\text{out}}(S(p_{\text{in}}(x + c_{\text{in}}))) + c_{\text{out}} \quad \text{for every } x \in \mathbb{F}_2^w.$$

Lemma 1 (Propagation identities induced by PXE). *Under the PXE relation above, for all $a, b, u, v \in \mathbb{F}_2^w$,*

$$\text{DDT}_{S'}(p_{\text{in}}^{-1}(a), p_{\text{out}}(b)) = \text{DDT}_S(a, b),$$

and

$$\text{SQLAT}_{S'}(p_{\text{in}}^{-1}(u), p_{\text{out}}(v)) = \text{SQLAT}_S(u, v).$$

Proof. For differences, the input translation cancels. The input difference $p_{\text{in}}^{-1}(a)$ becomes a at the input of S , and an output difference b becomes $p_{\text{out}}(b)$. For linear masks, the corresponding mask pair for S' is $(p_{\text{in}}^{-1}(u), p_{\text{out}}(v))$ because permutation matrices satisfy $p_{\text{in}}^T = p_{\text{in}}^{-1}$ and $p_{\text{out}}^{-T} = p_{\text{out}}$. The XOR translations may change the sign of a correlation but not its square. $\square$

3.2 Permutation characteristics

Permutation characteristics were introduced by Leander et al. for invariant-subspace attacks [LMR15]. Kim et al. used round-dependent word permutations to identify differential and linear trails that differ only by a relabeling of the parallel S-box positions [KHSH22].

Definition 6 (Permutation characteristic [LMR15, KHSH22]). Let Π be a bit permutation of $\mathcal{V}_{m,w}$. An r -round permutation characteristic of Π is a sequence

$$(D_0, D_1, \dots, D_r) \in \mathcal{W}_m^{r+1}$$

satisfying

$$\Pi \circ D_t = D_{t+1} \circ \Pi, \quad 0 \leq t < r.$$

Equivalently, $D_{t+1} = \Pi \circ D_t \circ \Pi^{-1}$. The characteristic is iterative if $D_r = D_0$.

Each D_t only reorders complete words. It therefore commutes with the parallel substitution layer and preserves the multiset of local differential and linear transition weights. Since Π and every D_t are bit permutations, they act in the same way on differences and linear masks. An iterative characteristic can be repeated for an arbitrary number of rounds.

3.3 DC/LC-Equivalence

Definition 7 (DC/LC-equivalence). Two SbpN instances X and X' are *DC/LC-equivalent* if, for every $r \geq 1$, there is a bijection from $\mathcal{T}_D^r(X)$ to $\mathcal{T}_D^r(X')$ that preserves the differential weight, and there is a bijection from $\mathcal{T}_L^r(X)$ to $\mathcal{T}_L^r(X')$ that preserves the linear weight. When this holds, we write

$$X \sim_{\text{DC/LC}} X'.$$

Remark 1 (Trail-weight multisets). For a fixed number of rounds, the differential and linear trail sets are finite. A weight-preserving bijection between two such sets exists if and only if the two sets contain the same number of trails of each weight. Thus, Definition 7 compares the complete differential and linear trail-weight multisets, including multiplicities, for every number of rounds. It is stronger than equality of the minimum trail weights or equality of the sets of weights that occur. The equivalences established below are proved by constructing explicit bijections between the trail sets.

The identity maps on trail sets give reflexivity, inverse bijections give symmetry, and composition of the bijections gives transitivity. Thus, $\sim_{\text{DC/LC}}$ is an equivalence relation.

For each supported block size N , let $\sim_{\text{DC/LC}}^{(N)}$ denote the restriction of $\sim_{\text{DC/LC}}$ to $\mathcal{X}_N$. Unlike the reductions used to obtain $\mathcal{X}_N$, this relation applies to complete cipher instances at each supported block size. It requires a correspondence between their complete trail sets for every number of rounds.

3.4 Permutation Characteristics Between Two Layers

Definition 8 (Permutation characteristic between two permutation layers). Let Π and $\tilde{\Pi}$ be bit permutations of $\mathcal{V}_{m,w}$. A sequence

$$W_0, W_1, \dots \in \mathcal{W}_m$$

is a *permutation characteristic from Π to $\tilde{\Pi}$* if

$$\tilde{\Pi} \circ W_t = W_{t+1} \circ \Pi \quad \text{for every } t \geq 0. \quad (6)$$

Lemma 2 (Inversion and composition of permutation characteristics). *Let $(W_t)_{t \geq 0}$ be a permutation characteristic from Π_0 to Π_1 . Then $(W_t^{-1})_{t \geq 0}$ is a permutation characteristic from Π_1 to Π_0 . If $(Z_t)_{t \geq 0}$ is a permutation characteristic from Π_1 to Π_2 , then $(Z_t \circ W_t)_{t \geq 0}$ is a permutation characteristic from Π_0 to Π_2 .*

Proof. The defining identity $\Pi_1 \circ W_t = W_{t+1} \circ \Pi_0$ is equivalent to

$$\Pi_0 \circ W_t^{-1} = W_{t+1}^{-1} \circ \Pi_1.$$

Moreover,

$$\Pi_2 \circ Z_t \circ W_t = Z_{t+1} \circ \Pi_1 \circ W_t = Z_{t+1} \circ W_{t+1} \circ \Pi_0.$$

Both conclusions follow from Definition 8. $\square$

Unlike Definition 6, the two permutation layers in Definition 8 need not be the same. Once W_0 is fixed, Equation (6) uniquely determines the entire sequence:

$$W_t = \tilde{\Pi}^t \circ W_0 \circ \Pi^{-t}, \quad t \geq 0. \quad (7)$$

The condition to be checked is that every permutation obtained from Equation (7) remains word-wise.

Proposition 1 (Trail equivalence from PXE and permutation characteristics). *Let $X := (S, \Pi)$ and $X' := (S', \Pi')$ be SbPN instances. Suppose that there exist*

$$p_{\text{in}}, p_{\text{out}} \in \text{Sym}(w), \quad c_{\text{in}}, c_{\text{out}} \in \mathbb{F}_2^w$$

such that

$$S'(x) = p_{\text{out}}(S(p_{\text{in}}(x + c_{\text{in}}))) + c_{\text{out}} \quad \text{for every } x \in \mathbb{F}_2^w.$$

Put

$$P_{\text{in}} := \bar{p}_{\text{in}}^{(m)}, \quad P_{\text{out}} := \bar{p}_{\text{out}}^{(m)},$$

and put

$$\tilde{\Pi} := P_{\text{in}} \circ \Pi' \circ P_{\text{out}}. \quad (8)$$

Suppose that there exists a sequence

$$W_0, W_1, \dots \in \mathcal{W}_m$$

satisfying

$$\tilde{\Pi} \circ W_t = W_{t+1} \circ \Pi \quad \text{for every } t \geq 0.$$

Then X and X' are DC/LC-equivalent. In particular, suppose that

$$p_{\text{in}} = p^{-1}, \quad p_{\text{out}} = p,$$

and put $P := \bar{p}^{(m)}$. Then

$$\tilde{\Pi} = P^{-1} \circ \Pi' \circ P.$$

For an r -round differential trail, the transformation is

$$a'_t := (W_t \circ P)(a_t), \quad b'_t := (W_t \circ P)(b_t), \quad 0 \leq t < r,$$

with final state difference

$$a'_r := (W_r \circ P)(a_r).$$

Proof. Let $(W_t)_{t \geq 0}$ be a permutation characteristic from Π to $\tilde{\Pi}$. The parallel within-word permutations P_{in} and P_{out} commute with every $W_t \in \mathcal{W}_m$. Equations (6) and (8) therefore give

$$\begin{aligned} \Pi' \circ P_{\text{out}} \circ W_t &= P_{\text{in}}^{-1} \circ \tilde{\Pi} \circ W_t \\ &= P_{\text{in}}^{-1} \circ W_{t+1} \circ \Pi. \end{aligned} \quad (9)$$

Let

$$\tau := (a_0, b_0, \dots, a_{r-1}, b_{r-1}, a_r) \in \mathcal{T}_D^r(X).$$

For $0 \leq t < r$, set

$$a'_t := (P_{\text{in}}^{-1} \circ W_t)(a_t), \quad b'_t := (P_{\text{out}} \circ W_t)(b_t),$$

and set

$$a'_r := (P_{\text{in}}^{-1} \circ W_r)(a_r).$$

By Equation (9),

$$a'_{t+1} = (P_{\text{in}}^{-1} \circ W_{t+1})(\Pi(b_t)) = \Pi'((P_{\text{out}} \circ W_t)(b_t)) = \Pi'(b'_t).$$

Lemma 1 preserves each local differential transition, while W_t only permutes the m word positions. Hence the resulting tuple is an r -round differential trail of X' with the same differential weight as τ .

For linear trails, use the same formulas with (u_t, v_t) in place of (a_t, b_t) . If F is the permutation matrix of any bit permutation in Equation (9), then $F^{-T} = F$, so the same calculation gives the linear connection condition. The SQLAT identity in Lemma 1, together with the word permutation W_t , preserves every local linear weight.

All transformations defining the two trail maps are permutations. Applying their inverses recovers the source trail, so both constructions are bijections for every $r \geq 1$.

When $p_{\text{in}} = p^{-1}$ and $p_{\text{out}} = p$, we have $P_{\text{in}} = P^{-1}$ and $P_{\text{out}} = P$. Substitution into Equation (8) and the differential trail map gives the stated formulas. $\square$

Connection with GM-equivalence. Sun et al. [SPWW22] establish GM-equivalence for GIFT-64 group mappings using row and column transformations. In their nibble-matrix representation, let R_ρ and C_ρ denote the row and column transformations associated with $\rho \in \text{Sym}(4)$, and put $Z_\rho := C_\rho \circ R_\rho$. Their relations

$$R_\rho \circ \text{TN} = \text{TN} \circ C_\rho, \quad C_\rho \circ \text{TN} = \text{TN} \circ R_\rho, \quad C_\rho \circ \text{GM}_g = \text{GM}_g \circ C_\rho$$

together with

$$\text{GM}_{g_2} = R_\rho \circ \text{GM}_{g_1} \circ R_\rho^{-1}$$

imply

$$\Pi_{g_2} \circ Z_\rho = Z_\rho \circ \Pi_{g_1}, \quad \Pi_g := \text{TN} \circ \text{GM}_g.$$

Hence the constant sequence $W_t := Z_\rho$ is a permutation characteristic from Π_{g_1} to Π_{g_2} . This identifies the fixed relabeling used in GM-equivalence with the permutation-characteristic condition used here. Proposition 1 also allows PXE-related S-boxes and round-dependent word relabelings. Sun et al. additionally establish the corresponding keyed equivalence by relabeling the round keys.

3.5 Checking a Permutation Characteristic

Equation (7) shows that the choice of W_0 determines the entire permutation characteristic. Since Π and $\tilde{\Pi}$ are permutations of a finite state space, both have finite order. It is therefore enough to examine finitely many terms of the resulting sequence.

Lemma 3 (Finite criteria for a permutation characteristic). *Fix $W_0 \in \mathcal{W}_m$, let*

$$W_t := \tilde{\Pi}^t \circ W_0 \circ \Pi^{-t},$$

and put

$$d := \text{lcm}(\text{ord}(\Pi), \text{ord}(\tilde{\Pi})).$$

Also put

$$V := W_0 \circ \Pi \circ W_0^{-1}.$$

Then the following conditions are equivalent.

(i) $W_t \in \mathcal{W}_m$ for every $t \geq 0$.

(ii) $W_t \in \mathcal{W}_m$ for $0 \leq t < d$.

(iii)

$$V^k \circ (\tilde{\Pi} \circ V^{-1}) \circ V^{-k} \in \mathcal{W}_m \quad \text{for } 0 \leq k < \text{ord}(V). \quad (10)$$

Proof. By the definition of d , $\Pi^d = \tilde{\Pi}^d = \text{id}$. Consequently,

$$W_{t+d} = \tilde{\Pi}^{t+d} \circ W_0 \circ \Pi^{-(t+d)} = W_t$$

for every $t \geq 0$. This proves the equivalence of (i) and (ii).

Since $W_0 \circ \Pi^{-t} = V^{-t} \circ W_0$ and $\tilde{\Pi} = (\tilde{\Pi} \circ V^{-1}) \circ V$, we have

$$W_t \circ W_0^{-1} = \prod_{k=0}^{t-1} (V^k \circ (\tilde{\Pi} \circ V^{-1}) \circ V^{-k}), \quad (11)$$

where the factor with $k = 0$ is placed on the left. The conjugates in Equation (10) have period $\text{ord}(V)$. If (iii) holds, every factor in Equation (11) is therefore word-wise. The right-hand side is word-wise because $\mathcal{W}_m$ is a group, which proves (i).

Conversely, suppose that (i) holds, and put

$$Q_t := W_t \circ W_0^{-1}.$$

The consecutive products in Equation (11) are then word-wise, and

$$Q_{k+1} = Q_k \circ (V^k \circ (\tilde{\Pi} \circ V^{-1}) \circ V^{-k}).$$

It follows that

$$Q_k^{-1} \circ Q_{k+1} = V^k \circ (\tilde{\Pi} \circ V^{-1}) \circ V^{-k} \in \mathcal{W}_m \quad \text{for every } k \geq 0.$$

This proves (iii). □

Thus we do not search independently for $W_1, W_2, \dots$. Once W_0 is fixed, the unique sequence determined by Equation (7) can be checked by either finite condition in Lemma 3.

4 Further Reduction of the BOGI Parameter Set

The prior reductions leave $\mathcal{X}_N$ as a set of 41,472 ciphers to be considered. Distinct elements of this set may still be DC/LC-equivalent after the shuffle for block size N is included. We now construct permutation-characteristic and trail-reversal transformations. The permutation-characteristic transformations are justified by Proposition 1, whereas the trail-reversal transformations are established separately for each block size. The connected components of the resulting graphs further reduce the number of representatives needed for trail analysis. We begin with $N = 64$, for which the state contains sixteen 4-bit words arranged as four groups of four words.

Lemma 4 (Closure of the identity BOGI pair). *Let S be an optimal BOGI-applicable S-box satisfying $(\text{id}, \text{id}) \in \text{BP}(S)$, and let $p \in \text{Sym}(4)$. Then both*

$$p \circ S \circ p^{-1} \quad \text{and} \quad p \circ S^{-1} \circ p^{-1}$$

are optimal BOGI-applicable S-boxes for which the identity pair is valid. Consequently, the XE classes containing these two S-boxes have representatives in the fixed list $S_0, \dots, S_{1727}$.

Proof. Every optimal BOGI-applicable S-box has two good and two bad input positions and two good and two bad output positions. The condition $(\text{id}, \text{id}) \in \text{BP}(S)$ therefore gives

$$\mathcal{B}_{\text{out}}(S) = \mathcal{G}_{\text{in}}(S), \quad \mathcal{B}_{\text{in}}(S) = \mathcal{G}_{\text{out}}(S).$$

For the inverse S-box,

$$\text{DDT}_{S^{-1}}(a, b) = \text{DDT}_S(b, a), \quad \text{LAT}_{S^{-1}}(u, v) = \text{LAT}_S(v, u).$$

Hence

$$\mathcal{B}_{\text{out}}(S^{-1}) = \mathcal{B}_{\text{in}}(S) = \mathcal{G}_{\text{out}}(S) = \mathcal{G}_{\text{in}}(S^{-1}),$$

which proves $(\text{id}, \text{id}) \in \text{BP}(S^{-1})$. Conjugation by p relabels all four good and bad position sets by the same permutation p , so it preserves the identity-pair condition. Inversion and bit-position relabeling also preserve differential uniformity and linearity. An XE transformation leaves the DDT unchanged and preserves the absolute values of all LAT entries. It therefore leaves the four good and bad position sets unchanged, so the identity-pair condition holds throughout each of the two resulting XE classes. Since the fixed list contains one representative from every such XE class, both classes have representatives in $S_0, \dots, S_{1727}$. $\square$

4.1 Permutation-characteristic equivalences for BOGI-64

The first construction combines a permutation characteristic of a local LS permutation with a word permutation and a PE transformation.

Permutation-characteristic transformations. Let $X_{s,i}^{(64)} := (S_s, \Pi_i^{(64)}) \in \mathcal{X}_{64}$. Choose $A \in \mathcal{W}_4$, an iterative permutation characteristic

$$(D_0, D_1, \dots, D_h)$$

of L_i , where $D_h = D_0$, and choose D_ℓ for some $\ell \in [h]$. Put

$$T := A \circ L_i \circ D_\ell \circ A^{-1}. \tag{12}$$

If $A = W_a$ and $D_\ell = W_d$, then T is induced by the Latin square with entries

$$t_{r,b} = a((L_i)_{d(a^{-1}(r)),b}).$$

Hence T is an LS permutation. Let $L_j \in \mathcal{L}_{\text{rep}}$ and $p \in \text{Sym}(4)$ be the unique pair satisfying

$$T = (\bar{p}^{(4)})^{-1} \circ L_j \circ \bar{p}^{(4)}. \quad (13)$$

By Lemma 4, there is a unique index $s' \in [1728]$ such that $S_{s'}$ is XE to

$$p \circ S_s \circ p^{-1}.$$

The resulting map

$$X_{s,i}^{(64)} \mapsto X_{s',j}^{(64)}$$

is called a *permutation-characteristic transformation*.

The existence of W_0 is decided over the full set $\mathcal{W}_{16}$ as described in Appendix B. The following structured family is retained because it contains an initial word permutation for every triple satisfying the condition and gives a compact description of those permutations. Index the sixteen words by $(r, c) \in [4] \times [4]$. For $a, b \in \text{Sym}(4)$, define

$$G_{a,b}(r, c) := (a(r), b(c)), \quad H_{64}(r, c) := (c, r),$$

and let

$$\mathcal{F}_{64} := \{G_{a,b} : a, b \in \text{Sym}(4)\} \cup \{H_{64} \circ G_{a,b} : a, b \in \text{Sym}(4)\}. \quad (14)$$

Thus

$$|\mathcal{F}_{64}| = 2(4!)(4!) = 1,152.$$

Proposition 2 (A criterion for DC/LC-equivalence in BOGI-64). *Let a permutation-characteristic transformation map*

$$X_{s,i}^{(64)} \mapsto X_{s',j}^{(64)},$$

and let p be the bit permutation in Equation (13). Put

$$P := \bar{p}^{(16)}, \quad \tilde{\Pi} := P^{-1} \circ \Pi_j^{(64)} \circ P.$$

Assume that there exists $W_0 \in \mathcal{W}_{16}$ for which the permutation

$$V := W_0 \circ \Pi_i^{(64)} \circ W_0^{-1},$$

satisfies

$$V^k \circ (\tilde{\Pi} \circ V^{-1}) \circ V^{-k} \in \mathcal{W}_{16} \quad \text{for } 0 \leq k < \text{ord}(V). \quad (15)$$

Then

$$X_{s,i}^{(64)} \sim_{\text{DC/LC}}^{(64)} X_{s',j}^{(64)}.$$

Proof. Equation (15) and Lemma 3 imply that

$$W_t := \tilde{\Pi}^t \circ W_0 \circ (\Pi_i^{(64)})^{-t}$$

belongs to $\mathcal{W}_{16}$ for every $t \geq 0$.

The target S-box $S_{s'}$ is PXE to S_s with $p_{\text{in}} = p^{-1}$ and $p_{\text{out}} = p$. The XOR constants arise only from the choice of the XE representative. Therefore Proposition 1 gives the required weight-preserving bijections for differential and linear trails. $\square$

The exact calculation in Appendix B considers every $(i, j, p) \in [24] \times [24] \times \text{Sym}(4)$ and decides whether such a W_0 exists anywhere in $\mathcal{W}_{16}$. Exactly 576 triples admit such a W_0 . They are precisely the triples obtained from the transformations above, and each has a representative $W_0 \in \mathcal{F}_{64}$.

4.2 Trail reversal

The second construction uses the inverse S-box and the inverse LS permutation. Let $K_{64} := \text{Shuf}_{64}$. The BOGI-64 shuffle is an involution, so $K_{64}^{-1} = K_{64}$.

Trail-reversal transformation. Let $X_{s,i}^{(64)} \in \mathcal{X}_{64}$. The inverse of an LS permutation is again an LS permutation: for each bit position b , its b th column is the inverse permutation of the b th column of L_i . Let $L_j \in \mathcal{L}_{\text{rep}}$ and $p_0 \in \text{Sym}(4)$ be the unique pair such that

$$(\overline{p_0}^{(4)})^{-1} \circ L_j \circ \overline{p_0}^{(4)} = L_i^{-1}. \quad (16)$$

By Lemma 4, there is a unique index $s' \in [1728]$ such that $S_{s'}$ is XE to

$$p_0 \circ S_s^{-1} \circ p_0^{-1}.$$

The map from $X_{s,i}^{(64)}$ to $X_{s',j}^{(64)}$ is called its *trail-reversal transformation*.

Proposition 3 (DC/LC-equivalence under trail-reversal transformations). *Every BOGI-64 cipher is DC/LC-equivalent to the cipher obtained by its trail-reversal transformation for every round number.*

Proof. Let $P_0 := \overline{p_0}^{(16)}$. Parallel application of Equation (16), together with $K_{64}^{-1} = K_{64}$, gives

$$P_0^{-1} \circ \Pi_j^{(64)} \circ P_0 = K_{64} \circ (\Pi_i^{(64)})^{-1} \circ K_{64}^{-1}. \quad (17)$$

For a differential trail

$$\tau = (a_0, b_0, \dots, a_{r-1}, b_{r-1}, a_r)$$

of $X_{s,i}^{(64)}$, define

$$a'_t := (P_0 \circ K_{64})(b_{r-1-t}), \quad b'_t := (P_0 \circ K_{64})(a_{r-1-t})$$

for $0 \leq t < r$. Define the final state difference by

$$a'_r := (P_0 \circ K_{64} \circ (\Pi_i^{(64)})^{-1})(a_0).$$

Indeed, Equation (17) implies

$$\Pi_j^{(64)} \circ P_0 \circ K_{64} = P_0 \circ K_{64} \circ (\Pi_i^{(64)})^{-1}. \quad (18)$$

which gives $a'_{t+1} = \Pi_j^{(64)}(b'_t)$ for every $0 \leq t < r$. Moreover,

$$\text{DDT}_{S_s^{-1}}(b, a) = \text{DDT}_{S_s}(a, b),$$

so the differential weight is unchanged. The corresponding identity

$$\text{SQLAT}_{S_s^{-1}}(v, u) = \text{SQLAT}_{S_s}(u, v)$$

gives the same conclusion for local linear weights. If F is the permutation matrix of any bit permutation in Equation (18), then $F^{-T} = F$. Hence the same reversed formulas, with (u_t, v_t) in place of (a_t, b_t) , satisfy the target linear connection. The remaining XE transformation is handled by Lemma 1. Its translations may change correlation signs but preserve SQLAT entries. Thus the construction is a weight-preserving bijection in both cases. $\square$

Let the permutation-characteristic graph on $\mathcal{X}_{64}$ have an edge between two ciphers whenever a permutation-characteristic transformation between them satisfies Equation (15). Let $\sim_{\text{perm}}^{(64)}$ denote the equivalence relation whose classes are the connected components of this graph, and let $\sim_{\text{rev}}^{(64)}$ be defined after adding the trail-reversal edges. Propositions 2 and 3 give

$$\sim_{\text{perm}}^{(64)} \subseteq \sim_{\text{rev}}^{(64)} \subseteq \sim_{\text{DC/LC}}^{(64)}$$

on $\mathcal{X}_{64}$.

Theorem 1 (BOGI-64 partition under permutation characteristics and trail reversal). *The quotient sets satisfy*

$$\begin{aligned} |\mathcal{X}_{64} / \sim_{\text{perm}}^{(64)}| &= 1,728, & |[X]_{\sim_{\text{perm}}^{(64)}}| &= 24, \\ |\mathcal{X}_{64} / \sim_{\text{rev}}^{(64)}| &= 864, & |[X]_{\sim_{\text{rev}}^{(64)}}| &= 48 \end{aligned}$$

for every $X \in \mathcal{X}_{64}$. Moreover, every $\sim_{\text{rev}}^{(64)}$ -class is contained in a $\sim_{\text{DC/LC}}^{(64)}$ -class.

Proof. Appendix B decides the existence of W_0 over all of $\mathcal{W}_{16}$ and shows that exactly 576 triples satisfy the condition. These are precisely the triples used in the construction. The resulting permutation-characteristic graph has 1,728 connected components of size 24. Adding every trail-reversal edge joins these components in pairs and gives 864 connected components of size 48. Proposition 2 shows that every permutation-characteristic edge joins DC/LC-equivalent ciphers, and Proposition 3 gives the same conclusion for every trail-reversal edge. Transitivity of $\sim_{\text{DC/LC}}^{(64)}$ therefore shows that each $\sim_{\text{rev}}^{(64)}$ -class is contained in one $\sim_{\text{DC/LC}}^{(64)}$ -class. $\square$

Consequently, choosing one representative from each class in $\mathcal{X}_{64} / \sim_{\text{rev}}^{(64)}$ is sufficient for trail analysis whenever the quantity of interest is preserved by DC/LC-equivalence. These transformations therefore reduce the number of representatives required for $\mathcal{X}_{64}$ from 41,472 to 864.

Theorem 1 gives the complete component structure of the permutation-characteristic graph and the graph obtained by adding trail-reversal edges. It does not assert that $\sim_{\text{rev}}^{(64)} = \sim_{\text{DC/LC}}^{(64)}$. A weight-preserving trail bijection outside the permutation-characteristic and trail-reversal transformations may join different $\sim_{\text{rev}}^{(64)}$ -classes.

4.3 Permutation-characteristic equivalences for BOGI-128

For $N = 128$, we have $g = 8$ and $m = 32$. The word shuffle is induced by

$$\sigma_{128}(j) = 8(j \bmod 4) + \left\lfloor \frac{j}{4} \right\rfloor, \quad j \in [32],$$

and has order 5. To specify the initial word permutations used in the calculation, we view the 32 word positions through the three placements of the binary coordinate in a $4 \times 4 \times 2$ coordinate system.

Candidate word permutations. Let

$$\mathcal{R}_{128} := \{(4, 4, 2), (4, 2, 4), (2, 4, 4)\}.$$

For $\mathbf{r} = (r_0, r_1, r_2) \in \mathcal{R}_{128}$, encode a word position by

$$\text{enc}_{\mathbf{r}}(t_0, t_1, t_2) := (t_0 r_1 + t_1) r_2 + t_2.$$

For

$$\mathbf{g} = (g_0, g_1, g_2) \in \text{Sym}(r_0) \times \text{Sym}(r_1) \times \text{Sym}(r_2),$$

define $G_{\mathbf{r}, \mathbf{g}} \in \text{Sym}(32)$ by

$$G_{\mathbf{r}, \mathbf{g}}(\text{enc}_{\mathbf{r}}(t_0, t_1, t_2)) := \text{enc}_{\mathbf{r}}(g_0(t_0), g_1(t_1), g_2(t_2)).$$

Let

$$\mathcal{F}_{128} := \{W_{G_{\mathbf{r}, \mathbf{g}}} : \mathbf{r} \in \mathcal{R}_{128}, g_k \in \text{Sym}(r_k)\}, \quad (19)$$

where identical word permutations arising from different coordinate descriptions are identified.

Thus, a member of $\mathcal{F}_{128}$ independently relabels two four-valued coordinates and one binary coordinate. The three elements of $\mathcal{R}_{128}$ place the binary coordinate in each possible position. Each coordinate description gives $4! 4! 2! = 1,152$ permutations, and

$$|\mathcal{F}_{128}| = 3,072.$$

The intersections among the three coordinate descriptions are calculated in Appendix B. This family provides a compact description of the initial word permutations used in the construction. It is not a restriction on the existence test below, which ranges over all of $\mathcal{W}_{32}$.

For $L_i, L_j \in \mathcal{L}_{\text{rep}}$ and $p \in \text{Sym}(4)$, put

$$P := \bar{p}^{(32)}, \quad \tilde{\Pi}_{j,p} := P^{-1} \circ \Pi_j^{(128)} \circ P.$$

Permutation-characteristic triples. Let $\mathcal{G}_{128}$ be the set of triples (i, j, p) with $L_i, L_j \in \mathcal{L}_{\text{rep}}$ and $p \in \text{Sym}(4)$ for which there exists $W_0 \in \mathcal{W}_{32}$ such that, with

$$V := W_0 \circ \Pi_i^{(128)} \circ W_0^{-1},$$

we have

$$V^k \circ (\tilde{\Pi}_{j,p} \circ V^{-1}) \circ V^{-k} \in \mathcal{W}_{32} \quad \text{for } 0 \leq k < \text{ord}(V). \quad (20)$$

Permutation-characteristic transformations. Let $X_{s,i}^{(128)} := (S_s, \Pi_i^{(128)}) \in \mathcal{X}_{128}$ and let $(i, j, p) \in \mathcal{G}_{128}$. By Lemma 4, there is a unique index $s' \in [1728]$ such that $S_{s'}$ is XE to

$$p \circ S_s \circ p^{-1}.$$

The map

$$X_{s,i}^{(128)} \mapsto X_{s',j}^{(128)}$$

is called the permutation-characteristic transformation determined by (i, j, p) .

Proposition 4 (DC/LC-equivalence under permutation-characteristic transformations for BOGI-128). *Every permutation-characteristic transformation maps a BOGI-128 cipher to a DC/LC-equivalent cipher for every round number.*

Proof. Since $(i, j, p) \in \mathcal{G}_{128}$, there is a $W_0 \in \mathcal{W}_{32}$ satisfying Equation (20). Lemma 3 therefore shows that

$$W_t := \tilde{\Pi}_{j,p}^t \circ W_0 \circ (\Pi_i^{(128)})^{-t}$$

belongs to $\mathcal{W}_{32}$ for every $t \geq 0$. The target S-box is PXE to S_s with $p_{\text{in}} = p^{-1}$ and $p_{\text{out}} = p$. Proposition 1 therefore gives weight-preserving bijections between the differential trail sets and between the linear trail sets for every round number. $\square$

Appendix B decides the existence of W_0 over all of $\mathcal{W}_{32}$ for the 13,824 triples. Exactly 96 triples admit such a W_0 , and each has one in the structured family $\mathcal{F}_{128}$.

The smaller reduction factor for BOGI-128 is not caused by restricting W_0 to $\mathcal{F}_{128}$. The exact decision over $\mathcal{W}_{32}$ finds only four pairs (j, p) for each source index i , whereas the corresponding BOGI-64 decision finds 24. The BOGI-128 shuffle has order 5, so the word-wise condition on W_t must hold throughout five distinct positions of its cycle. Each resulting triple has an initial permutation with one of the $4 \times 4 \times 2$ coordinate descriptions used in $\mathcal{F}_{128}$, and the exact decision shows that no other word permutation produces an additional triple. These transformations give components of size 4, which trail reversal subsequently joins in pairs.

4.4 Trail reversal for BOGI-128

Let $d := (0, 2, 1, 3) \in \text{Sym}(4)$ and put $D := W_d$. Define

$$L_i^{\text{rev}} := D \circ L_i^{-1} \circ D. \quad (21)$$

The inverse L_i^{-1} is an LS permutation, and composition with D on both sides only relabels its rows and symbols. Hence L_i^{rev} is also an LS permutation. Let $h := (0, 4, 2, 6, 1, 5, 3, 7) \in \text{Sym}(8)$, define $\eta_h \in \text{Sym}(32)$ by

$$\eta_h(4c + r) := 4h(c) + r, \quad c \in [8], \quad r \in [4],$$

and put

$$H_{128} := W_{\eta_h}, \quad \hat{D} := \text{Mix}_D^{(128)}, \quad K_{128} := \text{Shuf}_{128} \circ \hat{D} \circ H_{128}.$$

At word level, K_{128} reverses the five binary digits of a word index, and hence $K_{128}^{-1} = K_{128}$.

Lemma 5 (BOGI-128 reversal identity). *For every $L_i \in \mathcal{L}_{\text{rep}}$,*

$$K_{128} \circ (\Pi_i^{(128)})^{-1} \circ K_{128}^{-1} = \Pi_{L_i^{\text{rev}}}^{(128)}. \quad (22)$$

Proof. Put $Q := \text{Shuf}_{128}$ and $\hat{L}_i := \text{Mix}_{L_i}^{(128)}$. The permutation H_{128} permutes the eight 16-bit blocks and therefore commutes with $\hat{L}_i$ and $\hat{D}$. Moreover, H_{128} and $\hat{D}$ are involutions. Since $K_{128}^{-1} = K_{128}$,

$$K_{128} \circ (\Pi_i^{(128)})^{-1} \circ K_{128}^{-1} = Q \circ \hat{D} \circ H_{128} \circ \hat{L}_i^{-1} \circ \hat{D} \circ H_{128} = Q \circ \hat{D} \circ \hat{L}_i^{-1} \circ \hat{D} = \Pi_{L_i^{\text{rev}}}^{(128)}.$$

□

Trail-reversal transformation. Let $X_{s,i}^{(128)} \in \mathcal{X}_{128}$. Let $L_j \in \mathcal{L}_{\text{rep}}$ and $p_0 \in \text{Sym}(4)$ be the unique pair such that

$$L_i^{\text{rev}} = (\overline{p_0}^{(4)})^{-1} \circ L_j \circ \overline{p_0}^{(4)}, \quad (23)$$

By Lemma 4, let $s' \in [1728]$ be the unique index such that $S_{s'}$ is XE to

$$p_0 \circ S_s^{-1} \circ p_0^{-1}.$$

The map from $X_{s,i}^{(128)}$ to $X_{s',j}^{(128)}$ is called its *trail-reversal transformation*.

Proposition 5 (DC/LC-equivalence under trail-reversal transformations for BOGI-128). *Every BOGI-128 cipher is DC/LC-equivalent to the cipher obtained by its trail-reversal transformation for every round number.*

Proof. Put

$$P_0 := \overline{p_0}^{(32)}, \quad B := P_0 \circ K_{128}.$$

Lemma 5 and Equation (23) give

$$\Pi_j^{(128)} \circ B = B \circ (\Pi_i^{(128)})^{-1}. \quad (24)$$

For a differential trail

$$\tau = (a_0, b_0, \dots, a_{r-1}, b_{r-1}, a_r)$$

of $X_{s,i}^{(128)}$, define

$$a'_t := B(b_{r-1-t}), \quad b'_t := B(a_{r-1-t})$$

for $0 \leq t < r$, and put

$$a'_r := (B \circ (\Pi_i^{(128)})^{-1})(a_0).$$

Equation (24) gives $a'_{t+1} = \Pi_j^{(128)}(b'_t)$ for every $0 \leq t < r$. Moreover,

$$\text{DDT}_{S_s^{-1}}(b, a) = \text{DDT}_{S_s}(a, b), \quad \text{SQLAT}_{S_s^{-1}}(v, u) = \text{SQLAT}_{S_s}(u, v),$$

so reversing a local transition preserves its differential and linear weights. If F is the permutation matrix of B or either permutation layer, then $F^{-T} = F$. Hence the same reversed formulas with (u_t, v_t) in place of (a_t, b_t) satisfy the target linear connection. The PE transformation applies p_0 to both entries, K_{128} only permutes word positions, and the remaining XOR constants leave DDT and SQLAT entries unchanged. Thus the construction gives weight-preserving bijections for both trail sets. Applying the inverse maps and reversing the order again recovers the source trail. $\square$

Let the permutation-characteristic graph on $\mathcal{X}_{128}$ have an edge between two ciphers whenever a transformation determined by an element of $\mathcal{G}_{128}$ maps one to the other. Let $\sim_{\text{perm}}^{(128)}$ denote the equivalence relation whose classes are the connected components of this graph, and let $\sim_{\text{rev}}^{(128)}$ be defined after adding the trail-reversal edges. Propositions 4 and 5 give

$$\sim_{\text{perm}}^{(128)} \subseteq \sim_{\text{rev}}^{(128)} \subseteq \sim_{\text{DC/LC}}^{(128)}$$

on $\mathcal{X}_{128}$.

Theorem 2 (BOGI-128 partition under permutation characteristics and trail reversal). *The quotient sets satisfy*

$$\begin{aligned} |\mathcal{X}_{128} / \sim_{\text{perm}}^{(128)}| &= 10,368, & |[X]_{\sim_{\text{perm}}^{(128)}}| &= 4, \\ |\mathcal{X}_{128} / \sim_{\text{rev}}^{(128)}| &= 5,184, & |[X]_{\sim_{\text{rev}}^{(128)}}| &= 8 \end{aligned}$$

for every $X \in \mathcal{X}_{128}$. Moreover, every $\sim_{\text{rev}}^{(128)}$ -class is contained in a $\sim_{\text{DC/LC}}^{(128)}$ -class.

Proof. Appendix B decides the existence of W_0 over all of $\mathcal{W}_{32}$ and shows that exactly 96 triples satisfy the condition. The resulting permutation-characteristic graph has 10,368 connected components of size 4. Adding every trail-reversal edge joins these components in pairs and gives 5,184 connected components of size 8. Proposition 4 shows that every permutation-characteristic edge joins DC/LC-equivalent ciphers, and Proposition 5 gives the same conclusion for every trail-reversal edge. Transitivity of $\sim_{\text{DC/LC}}^{(128)}$ therefore shows that each $\sim_{\text{rev}}^{(128)}$ -class is contained in one $\sim_{\text{DC/LC}}^{(128)}$ -class. $\square$

Table 2: Reduction in the number of best-trail searches.

N	$ \mathcal{X}_N $	$ \mathcal{X}_N / \sim_{\text{rev}}^{(N)} $	Reduction factor
64	41,472	864	48
128	41,472	5,184	8

Consequently, choosing one representative from each class in $\mathcal{X}_{128} / \sim_{\text{rev}}^{(128)}$ is sufficient for trail analysis whenever the quantity of interest is preserved by DC/LC-equivalence. These transformations therefore reduce the number of representatives required for $\mathcal{X}_{128}$ from 41,472 to 5,184.

Theorem 2 gives the complete component structure of the permutation-characteristic graph and the graph obtained by adding trail-reversal edges. It does not assert that $\sim_{\text{rev}}^{(128)} = \sim_{\text{DC/LC}}^{(128)}$. A weight-preserving trail bijection outside the permutation-characteristic and trail-reversal transformations may join different $\sim_{\text{rev}}^{(128)}$ -classes.

Corollary 1 (Minimum trail weights within a constructed class). *For $N \in \{64, 128\}$, $r \geq 1$, and $X, X' \in \mathcal{X}_N$ with $X \sim_{\text{rev}}^{(N)} X'$, we have*

$$\text{MW}_{\text{D}}^r(X) = \text{MW}_{\text{D}}^r(X'), \quad \text{MW}_{\text{L}}^r(X) = \text{MW}_{\text{L}}^r(X').$$

Proof. For a bijective S-box, a supported differential transition has zero input difference if and only if its output difference is zero. The analogous statement holds for a supported linear transition, since a nonzero mask has zero correlation with the zero mask. Therefore a supported trail has nonzero initial state if and only if every round boundary, including the final state, is nonzero. The trail maps in Propositions 2, 3, 4, and 5 consist of invertible bit permutations and, for trail reversal, an exchange of the two ends. They consequently restrict to weight-preserving bijections between the nonzero-initial trail sets. Taking the minimum weight gives the result. $\square$

5 Applications of the Classification

5.1 Reduction of best-trail searches

Corollary 1 allows a best differential- and linear-trail search to be performed once for each constructed class. Table 2 gives the resulting search counts. The number of searches decreases from 41,472 to 864 for BOGI-64 and to 5,184 for BOGI-128.

An explicit best trail can also be recovered for every member without another search. Starting from the searched representative, apply the trail map attached to each edge of a path in the class graph. The resulting trail has the same weight, and the inverse edge maps recover the original trail.

Table 3 distinguishes the use of the earlier search results from the new search conducted in this paper.

5.2 Reinterpretation of the BOGI-64 results

The BOGI-64 results provide a consistency check for the classification. Kim et al. searched all 41,472 representatives through round 13 and found 96 representatives attaining the largest reported 12-round differential and linear weight pair (68, 64) [KHS22]. Under $\sim_{\text{rev}}^{(64)}$, these 96 representatives form exactly two constructed classes of size 48. The classification does not improve the previously reported best weights. It shows that the 96 combinations reduce to two cases for every trail property preserved by the bijections established in this paper.

Table 3: Best-trail searches in the previous work and in this paper.

Block size	Kim et al. [KSH22]	This paper
BOGI-64	All 41,472 representatives were searched through round 13. At round 12, 96 representatives attained the weight pair (68, 64).	The 96 representatives are shown to form exactly two constructed classes of size 48. No new best weight is claimed.
BOGI-128	All 41,472 representatives were searched through round 11. Only the eight representatives attaining the reported optimum were extended through round 22.	Best-trail searches through round 20 are conducted for the reduced set of 5,184 class representatives. Both results are complete for 5,081 representatives.

5.3 Extended search for BOGI-128

The benefit of the reduction is more substantial for BOGI-128. Kim et al. searched all 41,472 representatives through round 11, but extended only the eight representatives attaining the reported 11-round optimum through round 22 [KSH22]. These eight representatives belong to one class under $\sim_{\text{rev}}^{(128)}$. We instead applied the same best-trail search to the 5,184 representatives of $\mathcal{X}_{128}/\sim_{\text{rev}}^{(128)}$ through round 20.

For a representative X , define

$$R_D(X) := \min\{r \geq 1 : \text{MW}_D^r(X) \geq 128\}, \quad R_L(X) := \min\{r \geq 1 : \text{MW}_L^r(X) \geq 128\},$$

and put

$$J(X) := \max\{R_D(X), R_L(X)\}.$$

Here the minimum of the empty set is defined as ∞ . Smaller values of J mean that both best-trail weights reach the 128-bit threshold after fewer rounds.

The searches ran for more than two years on a server with an 80-core CPU. Both differential and linear searches completed through round 20 for 5,081 representatives. One of the two searches is incomplete for each of the other 103 representatives. Among the paired-complete representatives, none has $J \leq 18$ and 47 have $J = 19$. For 85 of the incomplete cases, the completed side has $R_D \geq 20$ or $R_L \geq 20$, and therefore $J \geq 20$. Each of the other 18 has $R_D = 19$ and an incomplete linear search. For these representatives, we decide whether a 19-round linear trail of correlation weight below 64 exists. The exact decision search verifies $\text{MW}_L^{19} \geq 128$ for another 12 representatives, so $R_L \leq 19$ and hence $J = 19$. The remaining six cases still satisfy $J \geq 19$, although their exact values are unresolved. Therefore,

$$\min_{X \in \mathcal{X}_{128}/\sim_{\text{rev}}^{(128)}} J(X) = 19,$$

and at least 59 constructed classes attain this value. The completion counts, candidate lists, and threshold decisions are given in Appendix D.

Figure 2 shows the joint distribution of the differential and linear best-trail weights at rounds 19 and 20. The shaded upper-right region contains the representatives for which both weights are at least 128. Its annotations (s, p) use the S-box and bit-permutation indices of Kim et al. [KSH22].

Figure 3 shows the best differential and linear trail weights through round 20 for the 5,081 paired-complete representatives. It also identifies the class containing GIFT-128 and the class containing the eight variants extended in [KSH22]. The GIFT-128 class has $R_L = 20$ and $R_D = 22$, and hence $J = 22$. Thus, every one of the 59 verified minimizers reaches both thresholds three rounds earlier than the GIFT-128 class.




Figure 2: Differential and linear best-trail weights of the 5,081 paired-complete BOGI-128 representatives at rounds 19 and 20.

5.4 Implementation selection within the BOGI-64 classes

DC/LC-equivalence does not require two members of a constructed class to have the same implementation cost. After selecting classes by their trail weights, we therefore expand their concrete S-box and bit-permutation pairs and evaluate those implementations separately. The evaluation below uses the two BOGI-64 classes identified in Section 5.2. Detailed construction, measurement, and validation procedures are given in Appendix C.

Kim’s thesis evaluated the same two security classes using S-box costs and generated operation sequences for their bitsliced mixing layers [Kim23]. We extend that evaluation to complete cipher implementations. The two classes contain 96 normalized S-box and bit-permutation pairs. Expanding every S-box through its 256 XE forms gives 24,576 candidate labels and 12,288 distinct S-box tables.

Hardware evaluation. We use a common iterative architecture with a 64-bit state, a 128-bit key, and 28 rounds. The key schedule, round-key placement, round constants, and controller are fixed to those of GIFT-64. Only the S-box and bit permutation vary. All 24,576 candidates and the GIFT-64 reference are synthesized from LUT RTL with Yosys and ABC [WG13, BM10] using the sky130_fd_sc_hd library [Sky20]. We additionally synthesize GIFT-64 and the six candidates selected from this scan using PEIGEN-expression RTL. The reported values are standard-cell area and the ABC critical-path estimate before placement and routing.

The GIFT-64 reference has area $7827.51 \mu\text{m}^2$ and estimated delay 705.48 ps. Among the BOGI-64 candidates, 97 have both smaller area and lower delay. Six implementations remain after removing every candidate for which another candidate is no larger and no slower, with at least one strict improvement. Figure 4 and Table 4 give these results.

For comparison, we also synthesize all 256 XE forms of the GIFT S-box with the same permutation and the same flow. When these implementations are considered together with the 24,576 BOGI-64 candidates, the six remaining area-delay tradeoffs are precisely the six BOGI-64 candidates in Table 4.

The PEIGEN-expression experiment is a targeted re-evaluation of these six candidates rather than a second full scan. Relative to the PEIGEN-expression GIFT-64 reference, none of the six candidates improves both area and delay. Candidate C has lower delay, while candidates D–F have smaller area. This difference shows that the synthesis comparison




Figure 3: Best differential and linear trail weights through round 20 for the 5,081 paired-complete BOGI-128 representatives.

depends on the RTL representation.

Software evaluation. We generate Boolean S-box expressions with PEIGEN [BGLS19] and word-operation sequences for the bit permutations with the Sirrida generator [Sir]. We evaluate both a classical bitsliced representation and the fixed-plane representation derived from GIFT-64 fixslicing [ANP20]. Both process two blocks in parallel. For each candidate, all four fixed-plane choices are generated and the fastest complete message path is reported.

No candidate is faster than GIFT-64 in either representation. Candidate C is the closest classical alternative at 9.2% above the GIFT-64 time, while candidate F is the closest fixed-plane alternative at 9.8% above it. Under the hardware flow, candidate F is 0.10% smaller than GIFT-64 and has 21.9% lower estimated critical-path delay. The isolated S-box cost is therefore insufficient to predict the complete software ranking because the bit-permutation cost also contributes to every round.

5.5 Implementation selection within the BOGI-128 classes

We next evaluate implementations from the 59 classes verified to attain $J = 19$. Their class members give 472 normalized S-box and bit-permutation specifications. Expanding every S-box through its 256 XE forms gives 120,832 LUT-RTL candidates. The six classes whose $J = 19$ status remains unresolved are not included in this evaluation. Detailed construction, measurement, and validation procedures are given in Appendix C.

Hardware evaluation. We use a common iterative architecture with a 128-bit state, a 128-bit key, and 40 rounds. The key schedule, round-key placement, round constants, and controller are fixed to those of GIFT-128. Only the S-box and bit permutation vary. The synthesis flow and reported area and delay metrics are the same as in the BOGI-64 evaluation.

All 120,832 LUT-RTL candidates were synthesized successfully. The corresponding GIFT-128 implementation has area $11300.84 \mu\text{m}^2$ and estimated delay 342.22 ps. Among the BOGI-128 candidates, 3,798 are smaller, 69,375 are faster, and 3,004 improve both




Figure 4: Area and delay of the 24,576 BOGI-64 LUT-RTL candidates under the common SKY130 synthesis flow.

quantities. At the faster of the two remaining area-delay points, the area is $10940.49 \mu\text{m}^2$ and the delay is 260.32 ps. These values improve the GIFT-128 baseline by 3.2% and 23.9%, respectively.

We also perform a separate synthesis using PEIGEN-generated Boolean expressions rather than LUT descriptions. The functionally validated set contains 712 implementations derived from 18 S-box expressions and covers 160 normalized specifications in 20 verified classes. This is a targeted expression-level evaluation, not a complete PEIGEN enumeration of all 120,832 LUT-RTL candidates. Relative to the GIFT-128 implementation generated in the same expression-based flow, 72 candidates are smaller and 44 are faster, but none improves both quantities. Thus the LUT and PEIGEN-expression results are reported separately. Figure 5 and Table 6 summarize both evaluations.

Software evaluation. We evaluate 32 PEIGEN-derived BOGI-128 configurations from the verified $J = 19$ classes and the corresponding GIFT-128 implementation as generated scalar x86 C. Every configuration is evaluated with bitslicing and a period-5 fixed-plane representation in the style of fixslicing. The same BOGI-128 configuration is fastest in both representations. In the bitsliced evaluation, 16 of the 32 configurations are faster than GIFT-128, and the minimum time is 9.8% lower. In the fixed-plane evaluation, 19 are faster, and the minimum time is 7.5% lower. Table 7 reports the complete time per encrypted block, including packing and unpacking. This targeted evaluation does not establish a software optimum over every implementation in the verified $J = 19$ classes.

6 Conclusion

We showed that the 41,472 reduced BOGI representatives from prior work still contain systematic repetitions at the level of complete differential and linear trail sets. Permutation characteristics and trail reversal generate 864 classes of size 48 for BOGI-64 and 5,184 classes of size 8 for BOGI-128. A best-trail search therefore needs only one representative from each constructed class. Whether additional weight-preserving bijections join different

Table 4: Hardware results for GIFT-64 and six selected BOGI-64 implementations. Each RTL block uses the corresponding GIFT-64 reference. The PEIGEN-expression rows re-evaluate the designs selected from the LUT scan.

RTL input	Point	Area (μm^2)	Delay (ps)
LUT	GIFT-64	7827.51	705.48
LUT	A	7753.69	695.66
LUT	B	7777.46	682.57
LUT	C	7782.46	667.46
LUT	D	7796.23	593.68
LUT	E	7809.99	553.65
LUT	F	7820.00	550.81
PEIGEN expression	GIFT-64	7880.06	1057.65
PEIGEN expression	A	8000.17	1246.93
PEIGEN expression	B	8115.28	1432.81
PEIGEN expression	C	8073.99	917.94
PEIGEN expression	D	7837.52	1220.88
PEIGEN expression	E	7669.86	1304.29
PEIGEN expression	F	7674.86	1272.97

constructed classes remains open.

This reduction enabled a systematic BOGI-128 best-trail search through round 20. We established $J_{\min} = 19$ and verified at least 59 classes attaining this value, whereas the class containing GIFT-128 has $J = 22$. Thus the classification does more than remove repeated computation. It exposes many component choices that reach both 128-bit best-trail-weight thresholds three rounds earlier than GIFT-128.

Unfolding a selected class also permits implementation costs to be compared without repeating the security search. Our hardware and software experiments show that the resulting ranking depends on the RTL or software representation. Consequently, the representative used for cryptanalytic evaluation need not be the best implementation choice within its class. This separation between classification, best-trail search, and implementation selection provides a practical basis for exploring future GIFT-based primitives.

AI disclosure. The authors supplied the initial research ideas and methods, including the definitions, proof ideas, and experimental designs. OpenAI’s GPT-5.6 Sol was used to help develop these ideas into detailed proofs and computational procedures, refine their presentation, assist with language editing and research-code development, check proofs and computational results, and prepare tables and figures. All AI-assisted output was reviewed by the authors, who verified the correctness and originality of the final content and take full responsibility for it.

A Trail Correspondence for the Identity BOGI Pair

The reduction to representatives with the identity BOGI pair is stated in prior work in terms of equal best differential and linear weights [KHS22]. The same calculation gives an explicit correspondence between the complete trail sets.

Table 5: Software measurements for GIFT-64 and the six hardware candidates. The fixslicing columns report the fastest of the four plane choices. Code size is the sum of the .text sizes of the setup and complete message-path functions.

Label	Bitslicing		Fixed plane i	Fixslicing	
	Time per encryption (ns)	Code size (bytes)		Time per encryption (ns)	Code size (bytes)
GIFT-64	200.897	1834	2	70.903	1961
A	277.508	2226	3	114.668	2102
B	279.770	2185	3	119.245	2133
C	219.459	1834	0	80.505	1955
D	236.480	2053	3	120.220	2228
E	239.944	2056	3	125.370	2183
F	223.722	1831	3	77.825	1986




Figure 5: Area and delay of the verified- $J = 19$ BOGI-128 implementations. Each panel uses the GIFT-128 implementation generated from the same RTL representation.

Lemma 6 (Identity BOGI-pair representation). *Let L be an LS permutation, let $(ib, ob) \in \text{BP}(S)$, and put $S' := ib \circ S \circ ob$. Define*

$$I := \overline{ib}^{(m)}, \quad O := \overline{ob}^{(m)},$$

and consider

$$X := (S, \Pi_{M(L, ib, ob)}^{(N)}), \quad X' := (S', \Pi_{M(L, id, id)}^{(N)}).$$

For every $r \geq 1$, the map

$$(a_0, b_0, \dots, a_{r-1}, b_{r-1}, a_r) \mapsto (O^{-1}(a_0), I(b_0), \dots, O^{-1}(a_{r-1}), I(b_{r-1}), O^{-1}(a_r))$$

is a weight-preserving bijection from $\mathcal{T}_D^r(X)$ to $\mathcal{T}_D^r(X')$. Replacing (a_t, b_t) by (u_t, v_t) gives a weight-preserving bijection from $\mathcal{T}_L^r(X)$ to $\mathcal{T}_L^r(X')$.

Proof. Write $\Pi := \Pi_{M(L, ib, ob)}^{(N)}$ and $\Pi' := \Pi_{M(L, id, id)}^{(N)}$. Parallel application of the mixing permutation gives

$$\text{Mix}_{M(L, ib, ob)}^{(N)} = O \circ \text{Mix}_{M(L, id, id)}^{(N)} \circ I.$$

Table 6: Hardware results for GIFT-128 and selected area-delay points in Figure 5. Each RTL block uses the corresponding GIFT-128 reference.

RTL input	Point	Area (μm^2)	Delay (ps)
LUT	GIFT-128	11300.84	342.22
LUT	A	10900.45	426.83
LUT	B	10940.49	260.32
PEIGEN expression	GIFT-128	11541.07	703.00
PEIGEN expression	C	11460.99	821.63
PEIGEN expression	D	11861.38	590.46

Table 7: Software measurements for GIFT-128 and the fastest evaluated BOGI-128 configuration.

Label	Bitslicing	Fixslicing-style	
	Time per encryption (ns)	Fixed plane i	Time per encryption (ns)
GIFT-128	649.577	3	174.431
Selected BOGI-128	585.664	0	161.273

The permutation O applies the same bit permutation independently to every word and therefore commutes with Shuf_N . Hence

$$\Pi' \circ I = O^{-1} \circ \Pi. \quad (25)$$

Lemma 1, with $p_{\text{in}} = ob$ and $p_{\text{out}} = ib$, gives

$$\text{DDT}_{S'}(ob^{-1}(a), ib(b)) = \text{DDT}_S(a, b), \quad \text{SQLAT}_{S'}(ob^{-1}(u), ib(v)) = \text{SQLAT}_S(u, v).$$

Thus the maps preserve every local weight. If $a_{t+1} = \Pi(b_t)$, Equation (25) gives

$$O^{-1}(a_{t+1}) = \Pi'(I(b_t)),$$

which is the connection condition for X' . The same argument applies to linear masks because all maps are bit permutations. Finally, I and O are invertible, so both trail maps are bijections. $\square$

B Calculation of the Transformation Graphs

This appendix gives the finite calculations used in Theorems 1 and 2. The code, input encodings, initial word permutations satisfying the criterion, and machine-readable graph summaries are available from our artifact repository: <https://anonymous.4open.science/r/Equivalence-Classes-of-BOGI-Based-Ciphers>. A permutation map Q on n coordinates is represented by its source tuple $(q_0, \dots, q_{n-1})$, defined by

$$(Qx)_j = x_{q_j}, \quad j \in [n].$$

In particular, if P_π moves input coordinate c to output coordinate $\pi(c)$, then its source tuple is

$$(\pi^{-1}(0), \dots, \pi^{-1}(n-1)).$$

Composition, inversion, membership in a word-permutation set, and equality of two permutations are evaluated in this representation. Two descriptions are identified exactly when their source tuples agree. For each block size, the graph has vertex set $\mathcal{X}_N$. Every transformation satisfying its stated criterion contributes an undirected edge between its source and target, repeated edges are identified, and the equivalence classes are the connected components of the resulting graph.

For a fixed source permutation, target permutation, and W_0 , the permutation V in Equations (15) and (20) has finite order. The conjugates in these equations are periodic with period dividing $\text{ord}(V)$, so the range $0 \leq k < \text{ord}(V)$ contains every distinct condition that can occur. Lemma 3 then gives a word permutation W_t for every $t \geq 0$. Hence the calculation has a finite domain while the resulting trail bijections apply to every number of rounds.

B.1 Exact decision of the initial word permutation

Let $n \in \{16, 32\}$, let R be a source permutation, and let U be the corresponding effective target permutation. Write every flat position uniquely as $4u + b$ with $u \in [n]$ and $b \in [4]$, and define

$$\lambda(4u + b) := u, \quad \beta(4u + b) := b.$$

Every element of $\mathcal{W}_n$ is W_σ for a unique $\sigma \in \text{Sym}(n)$, where

$$W_\sigma(4u + b) = 4\sigma(u) + b.$$

Put

$$d := \text{lcm}(\text{ord}(R), \text{ord}(U)), \quad Z_{t,\sigma} := U^t \circ W_\sigma \circ R^{-t}.$$

By Lemma 3, W_σ is a valid initial word permutation if and only if $Z_{t,\sigma} \in \mathcal{W}_n$ for $0 \leq t < d$. This condition is equivalent to

$$\beta(Z_{t,\sigma}(4u + b)) = b, \tag{26}$$

$$\lambda(Z_{t,\sigma}(4u + b)) = \lambda(Z_{t,\sigma}(4u)) \tag{27}$$

for every $u \in [n]$, $b \in [4]$, and $0 \leq t < d$. Equation (26) requires preservation of the position inside each word, while Equation (27) requires the four bits of a word to be moved to one word. Since $Z_{t,\sigma}$ is a permutation, the induced map on word positions is then also a permutation.

Algorithm 1 gives the exact decision procedure. For each variable u , its domain $D_u \subseteq [n]$ contains the possible values of $\sigma(u)$. Substituting each possible value into Equation (26) gives unary restrictions on these domains. Each instance of Equation (27) involving two variables gives a binary relation containing exactly their compatible value pairs. An equality involving the same variable gives another unary restriction.

During the search, an assigned value is removed from every other domain. For each binary relation, a value is removed if it has no compatible value in the other domain. These reductions are repeated until no domain changes. The remaining all-distinct condition is then tested by bipartite matching between the variables and their domains. Failure of this test proves that the partial assignment has no completion. Otherwise, the search chooses a variable with the smallest nonsingleton domain, tries every value in that domain, and backtracks when a recursive call fails. Ties are resolved first by the number of incident binary relations and then by the variable index. The ordering affects only which valid W_0 is returned.

The recursion does not enumerate all $n!$ completed permutations. Nevertheless, it is exhaustive: every permutation satisfying Equations (26) and (27) determines a root-to-leaf sequence that is never removed, because pruning is applied only when the current partial

Table 8: Exact counts of triples admitting an initial word permutation. The last column restricts W_0 to the structured family used to present the results.

Block size	Triples examined	Admitting $W_0 \in \mathcal{W}_n$	Admitting $W_0 \in \mathcal{F}_N$
64	13,824	576	576
128	13,824	96	96

assignment admits no completion satisfying these equations and the all-distinct condition. Every returned assignment is checked independently over the full period $0 \leq t < d$.

Algorithm 1: Exact decision of an initial word permutation.

Input: Permutations R, U on $4n$ bits.
Output: A permutation $\sigma \in \text{Sym}(n)$ such that $U^t \circ W_\sigma \circ R^{-t} \in \mathcal{W}_n$ for every $t \geq 0$, or NONE.

- 1 $d \leftarrow \text{lcm}(\text{ord}(R), \text{ord}(U))$;
- 2 construct the domains $(D_u)_{u \in [n]}$ and binary relations $\mathcal{C}$ from Equations (26) and (27) for $0 \leq t < d$;
- 3 **Function** $\text{Search}((D_u)_{u \in [n]})$:
- 4 **repeat**
- 5 remove every singleton value from all other domains;
- 6 remove every value having no support in an adjacent domain under the corresponding relation in $\mathcal{C}$;
- 7 **until** *no domain changes*;
- 8 **if** *some domain is empty* **then**
- 9 **return** NONE;
- 10 **if** *the variable-domain graph has no matching covering all variables* **then**
- 11 **return** NONE;
- 12 **if** *every domain is a singleton* **then**
- 13 let $\sigma(u)$ be the unique element of D_u for every u ;
- 14 **if** $U^t \circ W_\sigma \circ R^{-t} \in \mathcal{W}_n$ *for every* $0 \leq t < d$ **then**
- 15 **return** σ ;
- 16 **return** NONE;
- 17 choose u with the smallest nonsingleton domain, breaking ties as specified above;
- 18 **foreach** $v \in D_u$ *in a fixed order* **do**
- 19 $(D'_x)_{x \in [n]} \leftarrow (D_x)_{x \in [n]}$ and $D'_u \leftarrow \{v\}$;
- 20 $\sigma \leftarrow \text{Search}((D'_x)_{x \in [n]})$;
- 21 **if** $\sigma \neq \text{NONE}$ **then**
- 22 **return** σ ;
- 23 **return** NONE;
- 24 **return** $\text{Search}((D_u)_{u \in [n]})$;

For each block size, the calculation considers all

$$24 \cdot 24 \cdot 24 = 13,824$$

triples (i, j, p) . The result is shown in Table 8.

Algorithm 2 gives the complete calculation from the 13,824 triples to the two graph partitions. In this algorithm, EXACTINITIALWORDPERMUTATION denotes Algorithm 1. The trail-reversal image is defined for the corresponding block size in Section 4.

Algorithm 2: Transformation graph for block size N .

Input: The 24 permutations $\Pi_i^{(N)}$, the 1,728 S-box representatives, and $n = N/4$.
Output: The permutation-characteristic and enlarged component partitions.

```

1  $\mathcal{G}_N \leftarrow \emptyset$  and initialize an empty map  $\omega_N$ ;
2 foreach  $i, j \in [24]$  and  $p \in \text{Sym}(4)$  do
3    $R \leftarrow \Pi_i^{(N)}$  and  $P \leftarrow \bar{p}^{(n)}$ ;
4    $U \leftarrow P^{-1} \circ \Pi_j^{(N)} \circ P$ ;
5    $\sigma \leftarrow \text{EXACTINITIALWORDPERMUTATION}(R, U, n)$ ;
6   if  $\sigma \neq \text{NONE}$  then
7      $\text{add } (i, j, p) \text{ to } \mathcal{G}_N \text{ and set } \omega_N(i, j, p) \leftarrow \sigma$ ;
8 initialize an undirected graph  $H$  with vertex set  $\mathcal{X}_N$ ;
9 foreach  $X_{s,i}^{(N)} \in \mathcal{X}_N$  and  $(i, j, p) \in \mathcal{G}_N$  do
10   let  $s'$  be the unique index such that  $S_{s'}$  is XE to  $p \circ S_s \circ p^{-1}$ ;
11   add the edge  $\{X_{s,i}^{(N)}, X_{s',j}^{(N)}\}$  to  $H$ ;
12  $\mathcal{P}_{\text{perm}} \leftarrow$  the connected components of  $H$ ;
13 foreach  $X \in \mathcal{X}_N$  do
14   add the edge joining  $X$  to its trail-reversal image;
15  $\mathcal{P}_{\text{rev}} \leftarrow$  the connected components of  $H$ ;
16 return  $(\mathcal{G}_N, \omega_N, \mathcal{P}_{\text{perm}}, \mathcal{P}_{\text{rev}})$ ;
```

For $N = 64$, Algorithm 2 gives $|\mathcal{G}_{64}| = 576$, 1,728 components before the trail-reversal edges, and 864 afterward. For $N = 128$, it gives $|\mathcal{G}_{128}| = 96$, 10,368 components before the trail-reversal edges, and 5,184 afterward.

Thus the structured families $\mathcal{F}_{64}$ and $\mathcal{F}_{128}$ contain an initial word permutation for every triple admitting one in the corresponding full word-permutation set. Restricting the presentation to these families therefore omits no triple satisfying the condition. In particular, the exact calculation over $\mathcal{W}_{16}$ and $\mathcal{W}_{32}$ produces the same permutation-characteristic graphs used in the classifications below.

B.2 The BOGI-64 graph

For $L_i \in \mathcal{L}_{\text{rep}}$, define

$$\mathcal{D}_i := \{D \in \mathcal{W}_4 : L_i^t \circ D \circ L_i^{-t} \in \mathcal{W}_4 \text{ for every } t \geq 0\}.$$

Since L_i has finite order, it is sufficient to test $0 \leq t < \text{ord}(L_i)$. Hence $\mathcal{D}_i$ is obtained by a finite calculation over the $4! = 24$ elements of $\mathcal{W}_4$. Its elements are exactly the choices that can occur as D_ℓ in an iterative permutation characteristic of L_i .

When $D_\ell = \text{id}$, Equation (12) reduces to

$$T = A \circ L_i \circ A^{-1}.$$

Define

$$\mathcal{C}_i := \{C \in \mathcal{W}_4 : C \circ L_i \circ C^{-1} = L_i\}.$$

Two choices $A, A' \in \mathcal{W}_4$ give the same T if and only if $A' = A \circ C$ for some $C \in \mathcal{C}_i$. The 24 choices of A therefore give $24/|\mathcal{C}_i|$ distinct values of T . For each value, Equation (13) gives the corresponding triple (i, j, p) . Table 9 gives $|\mathcal{C}_i|$, $|\mathcal{D}_i|$, and the resulting component sizes for every LS representative.

The 12 representatives with $|\mathcal{D}_i| = 1$ and the 12 representatives with $|\mathcal{D}_i| = 4$ give

$$\sum_i |\mathcal{W}_4| |\mathcal{D}_i| = 12 \cdot 24 + 12 \cdot 24 \cdot 4 = 1,440$$

Table 9: Finite calculation for the BOGI-64 iterative permutation characteristics.

LS representatives	$ \mathcal{C}_i $	$ \mathcal{D}_i $	Component size	
			$D_\ell = \text{id}$	$D_\ell \in \mathcal{D}_i$
L_0, L_1, L_2, L_3	4	4	6	24
$L_4, L_5, L_8, L_9, L_{22}, L_{\text{GIFT}}$	2	4	12	24
L_{14}, L_{18}	1	4	24	24
$L_6, L_7, L_{10}, L_{11}, L_{12}, L_{13}, L_{15}, L_{16}, L_{17}, L_{19}, L_{20}, L_{21}$	1	1	24	24

Table 10: Distribution of the BOGI-64 triples by the number of distinct values generated from the selected W_0 .

Number of distinct values	1	2	4	6
Number of triples	360	108	72	36

triples (i, A, D_ℓ) before equal values of T are identified. For each i , conjugation by L_i permutes $\mathcal{D}_i$. Over all 60 elements in these sets, the resulting cycle-length distribution consists of 42 cycles of length 1, six cycles of length 2, and two cycles of length 3.

The number of distinct triples (i, j, p) obtained with $D_\ell = \text{id}$ is

$$4 \cdot 6 + 6 \cdot 12 + 2 \cdot 24 + 12 \cdot 24 = 432.$$

To make the choice of W_0 reproducible, order $\text{Sym}(4)$ lexicographically by image tuple. Generate

$$H_{64}^u \circ G_{a,b} \circ H_{64}^v, \quad (u, a, b, v) \in [2] \times \text{Sym}(4) \times \text{Sym}(4) \times [2],$$

in lexicographic order of (u, a, b, v) and retain the first occurrence of each image tuple. The resulting sequence lists every element of $\mathcal{F}_{64}$ exactly once. For each triple, select the first W_0 in this sequence that satisfies Equation (15). Such a W_0 exists in every case. Applying the corresponding transformations to every vertex of $\mathcal{X}_{64}$ gives

$$1,152 \cdot 6 + 864 \cdot 12 + 1,008 \cdot 24 = 41,472.$$

Thus the graph restricted to $D_\ell = \text{id}$ has 1,152 components of size 6, 864 components of size 12, and 1,008 components of size 24.

The nonidentity elements of $\mathcal{D}_i$ give 144 additional triples. Thus $432 + 144 = 576$ distinct triples are obtained, with exactly 24 triples for each source i . The structured-family cross-check has

$$576 \cdot 1,152 = 663,552$$

pairs of a triple and an initial word permutation. Across these pairs, 4,896 initial word permutations satisfy Equation (15), and every one of the 576 triples has at least one such element of $\mathcal{F}_{64}$. This independently reconstructs the same 576 triples obtained by the exact decision over $\mathcal{W}_{16}$. For the first valid W_0 in the ordering above, Table 10 gives the number of distinct values in the sequence

$$V^k \circ (\tilde{\Pi} \circ V^{-1}) \circ V^{-k}, \quad 0 \leq k < \text{ord}(V).$$

For $i \in \{0, 1, 2, 3\}$, the transformations arising from nonidentity elements of $\mathcal{D}_i$ join four of the size-6 components into one component of size 24. For $i \in \{4, 5, 8, 9, 22, \text{GIFT}\}$, they join two of the size-12 components. For $i \in \{14, 18\}$, they remain inside the size-24 components already obtained from the identity. The remaining indices have $\mathcal{D}_i = \{\text{id}\}$. Consequently, the full permutation-characteristic graph has

$$1,728 \cdot 24 = 41,472$$

vertices in 1,728 components of size 24.

Finally, add the trail-reversal edge of every vertex and calculate the connected components again. The resulting component-size distribution is

$$\#\{C : |C| = 48\} = 864.$$

Every permutation-characteristic component has size 24, and adding edges cannot split such a component. Each component of size 48 in the enlarged graph therefore contains exactly two permutation-characteristic components. This proves the pairwise merger used in Theorem 1.

B.3 The BOGI-128 graph

For each $\mathbf{r} \in \mathcal{R}_{128}$, let

$$\mathcal{F}_{\mathbf{r}} := \{W_{G_{\mathbf{r},\mathbf{g}}} : g_k \in \text{Sym}(r_k)\}.$$

Each of the three sets has

$$|\mathcal{F}_{\mathbf{r}}| = 4! 4! 2! = 1,152.$$

Write $\mathcal{F}_{442}$, $\mathcal{F}_{424}$, and $\mathcal{F}_{244}$ for the three sets in the order displayed in $\mathcal{R}_{128}$. Their intersections have sizes

$$\begin{aligned} |\mathcal{F}_{442} \cap \mathcal{F}_{424}| &= 192, & |\mathcal{F}_{442} \cap \mathcal{F}_{244}| &= 32, \\ |\mathcal{F}_{424} \cap \mathcal{F}_{244}| &= 192, & |\mathcal{F}_{442} \cap \mathcal{F}_{424} \cap \mathcal{F}_{244}| &= 32. \end{aligned}$$

The intersections of size 192 consist of one common $\text{Sym}(4)$ action and three independent binary swaps, giving $4! 2^3 = 192$ permutations. The intersection $\mathcal{F}_{442} \cap \mathcal{F}_{244}$ consists of five independent binary swaps, giving $2^5 = 32$ permutations, and it equals the triple intersection. Inclusion-exclusion now gives

$$\begin{aligned} |\mathcal{F}_{128}| &= 3 \cdot 1,152 - 192 - 32 - 192 + 32 \\ &= 3,072. \end{aligned}$$

There are

$$|\mathcal{L}_{\text{rep}}|^2 |\text{Sym}(4)| = 24 \cdot 24 \cdot 24 = 13,824$$

triples (i, j, p) . The exact decision described above evaluates the existence condition over all of $\mathcal{W}_{32}$. The triples for which at least one initial word permutation satisfies the criterion form $\mathcal{G}_{128}$, and

$$|\mathcal{G}_{128}| = 96.$$

Table 11 lists all 96 triples. In the table, a permutation $p \in \text{Sym}(4)$ is written as the four-digit string $p(0)p(1)p(2)p(3)$, and the label GIFT denotes L_{GIFT} . Exactly four elements of $\mathcal{G}_{128}$ have any fixed first coordinate i . One is (i, i, id) , and the other three are nonidentity triples. Thus $\mathcal{G}_{128}$ contains 24 identity triples and 72 nonidentity triples.

Within the structured family $\mathcal{F}_{128}$, there are 928 pairs (i, j, p, W_0) satisfying Equation (20). Hence every one of the 96 triples found over $\mathcal{W}_{32}$ also has an initial word permutation in $\mathcal{F}_{128}$. We order the coordinate descriptions as $(4, 4, 2)$, $(4, 2, 4)$, and $(2, 4, 4)$, order the axis permutations lexicographically within each description, and retain only the first occurrence of a repeated word permutation. Selecting the first W_0 satisfying the equation uses the six word permutations in Table 12. Their frequencies sum to 96.

Lemma 2, together with the fact that every parallel bit relabeling commutes with every word permutation, gives

$$(i, j, p) \in \mathcal{G}_{128} \implies (j, i, p^{-1}) \in \mathcal{G}_{128}$$

Table 11: The 96 triples in $\mathcal{G}_{128}$. For each source label i , the table gives the four pairs (j, p) such that $(i, j, p) \in \mathcal{G}_{128}$.

Source i	Pairs (j, p)
0	(0, 0123), (0, 1032), (0, 2301), (0, 3210)
1	(1, 0123), (1, 1023), (4, 2310), (4, 3210)
2	(2, 0123), (2, 2103), (GIFT, 1230), (GIFT, 3210)
3	(3, 0123), (3, 3120), (8, 1032), (8, 2031)
4	(1, 3201), (1, 3210), (4, 0123), (4, 0132)
5	(5, 0123), (5, 1032), (5, 2301), (5, 3210)
6	(6, 0123), (6, 0321), (16, 1032), (16, 1230)
7	(7, 0123), (7, 3120), (12, 1302), (12, 2301)
8	(3, 1032), (3, 1302), (8, 0123), (8, 0213)
9	(9, 0123), (9, 1032), (9, 2301), (9, 3210)
10	(10, 0123), (10, 0132), (10, 1023), (10, 1032)
11	(11, 0123), (11, 0321), (11, 2103), (11, 2301)
12	(7, 2031), (7, 2301), (12, 0123), (12, 0213)
13	(13, 0123), (13, 0321), (13, 2103), (13, 2301)
14	(14, 0123), (14, 1032), (14, 2301), (14, 3210)
15	(15, 0123), (15, 1023), (21, 2301), (21, 3201)
16	(6, 1032), (6, 3012), (16, 0123), (16, 2103)
17	(17, 0123), (17, 0132), (17, 1023), (17, 1032)
18	(18, 0123), (18, 1032), (18, 2301), (18, 3210)
19	(19, 0123), (19, 0213), (20, 1032), (20, 1302)
20	(19, 1032), (19, 2031), (20, 0123), (20, 3120)
21	(15, 2301), (15, 2310), (21, 0123), (21, 0132)
22	(22, 0123), (22, 1032), (22, 2301), (22, 3210)
GIFT	(2, 3012), (2, 3210), (GIFT, 0123), (GIFT, 0321)

Table 12: The six initial word permutations selected as the first W_0 satisfying Equation (20). All six belong to $\mathcal{F}_{442}$.

Index	g_0	g_1	g_2	Frequency
0	0123	0123	01	46
46	0123	3210	01	2
350	1032	1032	01	18
368	1032	2301	01	2
801	2301	2301	10	16
1151	3210	3210	10	12

and

$$(i, j, p), (j, k, q) \in \mathcal{G}_{128} \implies (i, k, q \circ p) \in \mathcal{G}_{128}.$$

Hence the transformations contain the identity and are closed under inversion and composition. For every $X \in \mathcal{X}_{128}$, the four images obtained from the triples with first coordinate i are distinct. Each connected component therefore contains four vertices, and

$$\frac{41,472}{4} = 10,368$$

gives the number of permutation-characteristic components.

Add the BOGI-128 trail-reversal edge of every vertex and calculate the connected components again. The resulting component-size distribution is

$$\#\{C : |C| = 8\} = 5,184.$$

Every permutation-characteristic component has size 4, so each component of size 8 in the enlarged graph contains exactly two permutation-characteristic components. In particular,

$$\frac{10,368}{2} = 5,184$$

gives the number of components of size 8.

The BOGI-128 calculation is exhaustive over all 13,824 triples and all initial word permutations in $\mathcal{W}_{32}$. The 96 triples satisfying the criterion are exactly those listed in Table 11. The structured family $\mathcal{F}_{128}$ is used only to give compact representatives of their initial word permutations.

C Implementation Evaluation Details

The source generators, validation programs, manifests, and result summaries used in this section are available from our artifact repository: <https://anonymous.4open.science/r/Equivalence-Classes-of-BOGI-Based-Ciphers>.

C.1 Candidate construction and hardware flow

The two selected BOGI-64 classes contain 96 normalized S-box and bit-permutation pairs, formed from 48 S-boxes and two bit permutations per S-box. For each S-box S and $c_{\text{in}}, c_{\text{out}} \in \mathbb{F}_2^4$, we evaluate the XE form

$$S_{c_{\text{in}}, c_{\text{out}}}(x) := S(x + c_{\text{in}}) + c_{\text{out}}.$$

Lemma 1 shows that this replacement preserves every differential and unsigned linear transition weight. The 256 XE forms give 24,576 candidate labels and 12,288 distinct S-box

tables. We regenerated PEIGEN expressions for all distinct tables and recomputed their Boolean gate counts.

We synthesized all 24,576 LUT-RTL candidates in a complete iterative core for the 28-round cipher, evaluating one round per clock cycle. The GIFT-64 key schedule, round-key placement, round constants, and control logic are fixed across all designs. Synthesis uses Yosys and ABC with the `sky130_fd_sc_hd` library. The reported area and critical-path delay exclude placement and routing. Relative to the GIFT-64 reference, 97 candidates improve both quantities. After including all 256 XE forms of the GIFT S-box under the same flow, the six remaining area-delay tradeoffs are the BOGI-64 implementations reported in Table 4.

We separately synthesized PEIGEN-generated Boolean expressions for GIFT-64 and these six implementations. Table 4 compares each implementation with the GIFT-64 reference generated using the same RTL representation. All reported designs passed functional validation. Detailed candidate construction, synthesis settings, and validation records are available in our artifact repository.

C.2 Software construction and measurement

The software implementations process two blocks B_0, B_1 in parallel in four 32-bit words $q_0, \dots, q_3$. For $b \in [2]$, $w \in [16]$, and $j \in [4]$, bit

$$8\lfloor w/4 \rfloor + 4b + (w \bmod 4)$$

of q_j stores bit j of word w in block B_b . Word w is the nibble at integer bit positions $4w, \dots, 4w + 3$.

Let P_j be the permutation induced on the 16 S-box positions of plane j . The classical representation applies P_j in every round. For a fixed-plane choice $i \in [4]$, put $a := P_i^{-1}$ and let $A_t := a^t$ act on the word positions of all four planes. If x_t is the state entering round t , the stored state is $\tilde{x}_t := A_t(x_t)$. The effective permutation on plane j is

$$F_{t,j} := A_{t+1} \circ P_j \circ A_t^{-1}, \quad F_{t,i} = \text{id}.$$

The round- t key is transformed by A_{t+1} , and A_{28}^{-1} is applied before unpacking. All four choices of i are generated for each candidate. The actions of every generated linear layer were checked on all basis vectors.

Software measurements were obtained on an Intel Xeon Gold 6230 at 2.10 GHz using GCC 13.3.0. We report the complete message-path time per encryption, including packing, 28-round encryption, the exit transformation, and unpacking with a precomputed key schedule. All implementations were checked against an independent reference. Detailed compiler options, benchmarking and code-size procedures, test vectors, and validation records are available in our artifact repository.

C.3 BOGI-128 hardware construction and validation

The BOGI-128 hardware evaluation considers the 59 classes verified to attain $J = 19$. Expanding their 472 normalized members through the 256 XE forms gives 120,832 LUT-RTL candidates. All candidates were synthesized using the same Yosys, ABC, and SKY130 flow as for BOGI-64. The evaluated design is a complete 40-round iterative core, including the key schedule and control logic, with one round evaluated per clock cycle. Placement and routing are not included.

The separate PEIGEN-expression evaluation covers 712 full-core implementations derived from 18 validated representative S-box expressions, spanning 160 normalized specifications in 20 verified classes. All implementations passed functional validation. The GIFT-128 reference was generated and evaluated using the same representation. Detailed construction and validation records are available in our artifact repository.

Table 13: BOGI-128 representatives reaching the 128-bit best-trail-weight threshold.

Round r	$W_D \geq 128$	$W_L \geq 128$	Both
18	31/5,177	5/5,088	0/5,081
19	448/5,177	181/5,088	47/5,081
20	1,703/5,177	1,040/5,088	554/5,081

C.4 BOGI-128 software measurement

The software evaluation compares 32 PEIGEN-derived configurations from verified $J = 19$ classes with GIFT-128. For each configuration, we generate scalar x86 C implementations using bitslicing and a period-5 fixed-plane representation. We report the complete message-path time per encryption, including packing, 40-round encryption, and unpacking. Key-schedule setup is excluded. All implementations were checked against an independent reference. Detailed compiler options, benchmarking procedures, source and binary checks, and validation records are available in our artifact repository.

D Detailed BOGI-128 Search Results

This appendix gives the search counts and threshold decisions supporting the BOGI-128 results in Section 5. A searched representative is indexed by (s, p) , where s identifies its S-box and p its bit permutation in the numbering of Kim et al. [KHS22]. Differential and linear weights are the negative base-2 logarithms of trail probability and squared trail correlation, respectively. Nonintegral values are rounded to two decimal places.

Both searches completed through round 20 for 5,081 of the 5,184 representatives. The remaining cases comprise seven incomplete differential searches and 96 incomplete linear searches, with no overlap. Table 13 counts the completed searches that reach the 128-bit threshold.

Among the 5,081 representatives for which both searches completed, 47 attain $J = 19$. Table 14 records their exact weights at rounds 18–20.

Table 15 lists the 103 representatives for which one search is incomplete.

The completed result excludes 85 of these representatives from attaining $J = 19$. Each of the other 18 has $R_D = 19$ and an incomplete linear search. It remains only to decide whether its 19-round linear trail set contains a trail with $W_L < 128$.

For the threshold decision, define the correlation weight by $\widehat{W}_L = -\log_2 |C|$, where C is the trail correlation. Thus $W_L = 2\widehat{W}_L$, and the required cutoff is $B = 64$. Let h_q be a certified lower bound on the correlation weight of every nonzero q -round trail of the same cipher, where $0 \leq q < r$ and $h_0 = 0$. Exact shorter-round minimum weights provide such bounds. A nonzero supported trail cannot reach the zero mask at an intermediate boundary, so h_q also applies to every remaining suffix. Algorithm 3 discards a partial trail only when no continuation can have correlation weight below B . An admissible continuation is a one-round mask transition with a nonzero LAT entry. The algorithm describes the exhaustive decision at trail level; its outer loop enumerates admissible first-round partial trails rather than all state masks independently.

Table 14: Exact weights of the 47 BOGI-128 class representatives attaining $J = 19$.

Parameter indices (s, p)	Round 18: (W_D, W_L)	Round 19: (W_D, W_L)	Round 20: (W_D, W_L)
(4, 98)	(123, 120)	(129, 130)	(137, 138)
(4, 102)	(118.66, 120)	(129.42, 128)	(137.42, 136)
(10, 125)	(123, 122)	(129.83, 130)	(140.42, 138)
(14, 344)	(117.42, 120)	(128.42, 128)	(138.42, 136)
(27, 51)	(122, 122)	(130, 130)	(142, 138)
(27, 102)	(118.66, 120)	(129.42, 128)	(137.42, 138)
(39, 4)	(117.42, 120)	(128.42, 130)	(137.42, 136)
(39, 125)	(120, 120)	(129, 128)	(134, 138)
(192, 125)	(127.42, 124)	(137.42, 132)	(148.25, 140)
(200, 25)	(123.25, 120)	(130, 130)	(139.25, 140)
(200, 99)	(118.66, 126)	(130.32, 134)	(139.83, 140)
(384, 77)	(117, 120)	(128, 130)	(136.42, 140)
(389, 99)	(125, 130)	(131, 138)	(139, 146)
(389, 125)	(122, 130)	(130, 138)	(134, 144)
(391, 99)	(121.42, 124)	(133.42, 132)	(141.42, 140)
(391, 125)	(120, 126)	(130, 134)	(139, 142)
(394, 77)	(120.42, 118)	(128, 128)	(135.42, 138)
(394, 125)	(121, 122)	(130, 130)	(140, 142)
(399, 77)	(117, 120)	(128, 128)	(135, 136)
(403, 25)	(124, 124)	(134, 132)	(143, 142)
(406, 77)	(117, 118)	(128, 128)	(135, 138)
(407, 125)	(129.83, 118)	(135.83, 130)	(145.83, 140)
(414, 0)	(128, 124)	(134, 132)	(144, 140)
(414, 99)	(118.83, 122)	(128, 130)	(137.83, 138)
(418, 99)	(125, 124)	(133, 132)	(140, 140)
(418, 125)	(122, 124)	(132, 128)	(140, 136)
(420, 25)	(121, 126)	(128, 136)	(136, 144)
(420, 99)	(121, 118)	(128, 128)	(136, 136)
(864, 6)	(123, 122)	(131.42, 130)	(139.42, 140)
(864, 25)	(124, 120)	(132, 128)	(141.42, 138)
(866, 25)	(125.83, 122)	(134.83, 130)	(143.83, 138)
(868, 77)	(123.25, 118)	(130.25, 128)	(139.42, 138)
(1060, 99)	(124, 118)	(130.83, 128)	(136.83, 136)
(1070, 25)	(129.83, 122)	(139.83, 130)	(149.25, 140)
(1070, 99)	(123.83, 122)	(130.83, 128)	(137.83, 136)
(1072, 25)	(127.83, 120)	(135, 128)	(143, 136)
(1251, 25)	(126.25, 120)	(134, 128)	(141.83, 136)
(1443, 25)	(121.91, 120)	(129.42, 128)	(136.83, 138)
(1443, 99)	(122.83, 122)	(132.83, 130)	(138.83, 138)
(1445, 99)	(127.42, 124)	(136.42, 132)	(146.83, 140)
(1447, 25)	(125.25, 120)	(132.66, 128)	(140.66, 132)
(1447, 99)	(120.08, 122)	(128.83, 130)	(136.25, 138)
(1451, 99)	(120.08, 122)	(128.83, 130)	(133.83, 140)
(1462, 48)	(123.42, 122)	(130.42, 130)	(136.42, 138)
(1463, 125)	(123.25, 120)	(133.25, 128)	(141.83, 136)
(1472, 6)	(122.42, 120)	(132, 128)	(136.83, 136)
(1481, 125)	(131.25, 120)	(142, 132)	(146, 140)

Algorithm 3: Decision of whether a linear trail below a given correlation-weight cutoff exists.

Input: A cipher X , a round number r , a cutoff B , and certified suffix lower bounds $h_0, \dots, h_{r-1}$.

Output: A nonzero r -round linear trail of correlation weight below B , or NONE.

```

1 Function Extend( $t, a, w, T$ ):
2   if  $t > 0$  and  $w + h_{r-t} \geq B$  then
3     return NONE;
4   if  $t = r$  then
5     return  $T$ ;
6   foreach admissible one-round continuation  $a \mapsto a'$  of correlation weight  $\delta$  do
7      $T' \leftarrow T$  followed by this transition;
8      $T^* \leftarrow \text{Extend}(t + 1, a', w + \delta, T')$ ;
9     if  $T^* \neq \text{NONE}$  then
10      return  $T^*$ ;
11  return NONE;
12 foreach admissible first-round partial trail  $T_1 : a_0 \mapsto a_1$  of correlation weight  $\delta_0$ 
    with  $a_0 \neq 0$  and  $\delta_0 + h_{r-1} < B$  do
13    $T^* \leftarrow \text{Extend}(1, a_1, \delta_0, T_1)$ ;
14   if  $T^* \neq \text{NONE}$  then
15     return  $T^*$ ;
16 return NONE;

```

Every complete nonzero trail determines one branch of Algorithm 3. Hence a return value of NONE proves $W_L(s, p, 19) \geq 128$. The decision returned NONE for 12 representatives. Together with the 47 representatives identified by the completed searches, this verifies 59 classes with $J = 19$. The remaining six cases are unresolved and are not included in this count.

Kim et al. [KHS22] extended eight BOGI-128 variants through round 22. Under our classification, all eight belong to the single class represented by (S_{414}, L_0) . Thus, their long-round results describe one constructed class rather than eight distinct classes. For comparison, GIFT-128 belongs to the class represented by (S_{10}, L_2) . Table 17 compares the best-trail weights of these two classes.

Table 15: BOGI-128 parameter pairs with an incomplete search.

Missing differential results (7)					
(866, 96)	(873, 96)	(1057, 96)	(1078, 122)	(1083, 99)	(1083, 102)
(1093, 96)					
Missing linear results (96)					
(1, 25)	(1, 30)	(1, 48)	(1, 102)	(3, 48)	(3, 99)
(3, 102)	(4, 25)	(4, 48)	(4, 122)	(5, 1)	(5, 24)
(8, 1)	(10, 25)	(10, 96)	(10, 99)	(11, 1)	(11, 24)
(11, 124)	(13, 1)	(14, 6)	(14, 30)	(14, 72)	(14, 99)
(14, 102)	(14, 122)	(17, 30)	(18, 24)	(21, 48)	(22, 30)
(23, 25)	(23, 122)	(24, 48)	(24, 72)	(24, 102)	(27, 48)
(33, 1)	(33, 24)	(40, 48)	(40, 102)	(48, 28)	(48, 96)
(49, 72)	(49, 102)	(192, 30)	(192, 48)	(192, 99)	(192, 102)
(192, 122)	(199, 24)	(200, 6)	(200, 102)	(200, 122)	(200, 125)
(201, 1)	(203, 24)	(210, 48)	(210, 102)	(210, 122)	(213, 25)
(213, 30)	(213, 102)	(216, 6)	(216, 122)	(219, 6)	(219, 30)
(221, 30)	(224, 1)	(227, 25)	(237, 25)	(248, 72)	(594, 102)
(864, 30)	(865, 48)	(869, 28)	(869, 54)	(871, 48)	(879, 48)
(883, 25)	(884, 28)	(884, 48)	(884, 122)	(886, 99)	(886, 125)
(894, 96)	(901, 72)	(914, 54)	(914, 74)	(1070, 30)	(1072, 99)
(1073, 48)	(1073, 102)	(1076, 30)	(1102, 6)	(1102, 102)	(1102, 122)

Table 16: Status of the 19-round linear threshold decisions for the remaining 18 representatives.

Status	Parameter indices (s, p)
Verified $W_L(s, p, 19) \geq 128$ (12)	(4, 25), (10, 25), (23, 25), (192, 99), (200, 102), (200, 122), (210, 102), (213, 25), (227, 25), (883, 25), (886, 125), (1072, 99)
Unresolved (6)	(192, 102), (200, 125), (210, 122), (216, 122), (1070, 30), (1076, 30)

Table 17: Long-round best-trail weights for the previously extended class and the class containing GIFT-128.

Round r	(S_{414}, L_0)		(S_{10}, L_2)	
	W_D	W_L	W_D	W_L
18	128	124	103.42	112
19	134	132	110.83	118
20	144	140	121.42	128
21	150	148	126.42	136
22	158	156	132.42	148

References

- [ANP20] Alexandre Adomnicaï, Zakaria Najm, and Thomas Peyrin. Fixslicing: A new GIFT representation: Fast constant-time implementations of GIFT and GIFT-COFB on ARM cortex-m. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2020(3):402–427, 2020. doi:10.46586/tches.v2020.i3.402-427.
- [BBP⁺19] Subhadeep Banik, Andrey Bogdanov, Thomas Peyrin, Yu Sasaki, Siang Meng Sim, Elmar Tischhauser, and Yosuke Todo. SUNDABE-GIFT v1.0. Submission to Round 1 of the NIST Lightweight Cryptography Standardization Process, 2019. URL: <https://csrc.nist.gov/CSRC/media/Projects/Lightweight-Cryptography/documents/round-1/spec-doc/SUNDABE-GIFT-spec.pdf>.
- [BCI⁺20] Subhadeep Banik, Avik Chakraborti, Akiko Inoue, Tetsu Iwata, Kazuhiko Minematsu, Mridul Nandi, Thomas Peyrin, Yu Sasaki, Siang Meng Sim, and Yosuke Todo. GIFT-COFB. Cryptology ePrint Archive, Paper 2020/738, 2020. URL: <https://eprint.iacr.org/2020/738>.
- [BGLS19] Zhenzhen Bao, Jian Guo, San Ling, and Yu Sasaki. Peigen – a platform for evaluation, implementation, and generation of S-boxes. *IACR Transactions on Symmetric Cryptology*, 2019(1):330–394, 2019. doi:10.13154/tosc.v2019.i1.330-394.
- [BM10] Robert K. Brayton and Alan Mishchenko. ABC: An academic industrial-strength verification tool. In *Computer Aided Verification*, volume 6174 of *Lecture Notes in Computer Science*, pages 24–40. Springer, 2010. doi:10.1007/978-3-642-14295-6_5.
- [BPP⁺17] Subhadeep Banik, Sumit Kumar Pandey, Thomas Peyrin, Yu Sasaki, Siang Meng Sim, and Yosuke Todo. GIFT: A small present. In Wieland Fischer and Naofumi Homma, editors, *Cryptographic Hardware and Embedded Systems – CHES 2017*, pages 321–345, Cham, 2017. Springer International Publishing. doi:10.1007/978-3-319-66787-4_16.
- [BS91] Eli Biham and Adi Shamir. Differential cryptanalysis of DES-like cryptosystems. *Journal of Cryptology*, 4(1):3–72, 1991. doi:10.1007/BF00630563.
- [CDJ⁺21] Avik Chakraborti, Nilanjan Datta, Ashwin Jha, Cuauhtémoc Mancillas-López, Mridul Nandi, and Yu Sasaki. Elastic-tweak: A framework for short tweak tweakable block cipher. In *Progress in Cryptology – INDOCRYPT 2021*, volume 13143 of *Lecture Notes in Computer Science*, pages 114–137. Springer, 2021. doi:10.1007/978-3-030-92518-5_6.
- [DGMNP25] Patrick Derbez, Baptiste Germon, Bastien Michel, and María Naya-Plasencia. Improved cryptanalysis of GIFT-64. *IACR Transactions on Symmetric Cryptology*, 2025(4):284–307, 2025. doi:10.46586/tosc.v2025.i4.284-307.
- [DH19] Orr Dunkelman and Senyang Huang. Reconstructing an S-box from its difference distribution table. *IACR Transactions on Symmetric Cryptology*, 2019(2):193–217, 2019. doi:10.13154/tosc.v2019.i2.193-217.
- [IKM⁺19] Tetsu Iwata, Mustafa Khairallah, Kazuhiko Minematsu, Thomas Peyrin, Yu Sasaki, Siang Meng Sim, and Ling Sun. Thank goodness it’s friday

- (TGIF) v1.0. Submission to Round 1 of the NIST Lightweight Cryptography Standardization Process, 2019. URL: <https://csrc.nist.gov/CSRC/media/Projects/Lightweight-Cryptography/documents/round-1/spec-doc/TGIF-spec.pdf>.
- [KHS20] Seonggyeom Kim, Deukjo Hong, Jaechul Sung, and Seokhie Hong. Classification of 4-bit S-boxes for BOGI permutation. *IEEE Access*, 8:210935–210949, 2020. doi:10.1109/ACCESS.2020.3039273.
- [KHS22] Seonggyeom Kim, Deukjo Hong, Jaechul Sung, and Seokhie Hong. Accelerating the best trail search on AES-like ciphers. *IACR Transactions on Symmetric Cryptology*, 2022(2):201–252, 2022. doi:10.46586/tosc.v2022.i2.201–252.
- [Kim23] Seonggyeom Kim. *Best Trail Search on AES-like Ciphers*. PhD thesis, Korea University, 2023.
- [LMM91] Xuejia Lai, James L. Massey, and Sean Murphy. Markov ciphers and differential cryptanalysis. In *Advances in Cryptology – EUROCRYPT ’91, Workshop on the Theory and Application of Cryptographic Techniques, Brighton, UK, April 8–11, 1991, Proceedings*, pages 17–38, Berlin, Heidelberg, 1991. Springer. doi:10.1007/3-540-46416-6_2.
- [LMR15] Gregor Leander, Brice Minaud, and Sondre Rønjom. A generic approach to invariant subspace attacks: Cryptanalysis of robin, iSCREAM and zorro. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology – EUROCRYPT 2015, Part I*, volume 9056 of *Lecture Notes in Computer Science*, pages 254–283. Springer, 2015. doi:10.1007/978-3-662-46800-5_11.
- [Mat94] Mitsuru Matsui. Linear cryptanalysis method for DES cipher. In Tor Helleseth, editor, *Advances in Cryptology – EUROCRYPT ’93*, volume 765 of *Lecture Notes in Computer Science*, pages 386–397. Springer, 1994. doi:10.1007/3-540-48285-7_33.
- [Saa12] Markku-Juhani O. Saarinen. Cryptographic analysis of all 4×4 -bit S-boxes. In Ali Miri and Serge Vaudenay, editors, *Selected Areas in Cryptography – SAC 2011*, volume 7118 of *Lecture Notes in Computer Science*, pages 118–133. Springer, 2012. doi:10.1007/978-3-642-28496-0_7.
- [Sir] Sirrida. Bit permutations. Accessed 2026-09-01. URL: https://programming.sirrida.de/bit_perm.html.
- [Sky20] SkyWater PDK Authors. SkyWater SKY130 open-source PDK, 2020. URL: <https://skywater-pdk.readthedocs.io/en/main/>.
- [SPWW22] Ling Sun, Bart Preneel, Wei Wang, and Meiqin Wang. A greater GIFT: Strengthening GIFT against statistical cryptanalysis. In *Advances in Cryptology – EUROCRYPT 2022, Part III*, volume 13277 of *Lecture Notes in Computer Science*, pages 115–144. Springer, 2022. doi:10.1007/978-3-031-07082-2_5.
- [Sun26] Ling Sun. A deeper look at related-key differentials of GIFT-64: From trails to differentials. *IACR Transactions on Symmetric Cryptology*, 2026(2), 2026. doi:10.46586/tosc.av7ta6k5chw2.

-
- [TMC⁺23] Meltem Sönmez Turan, Kerry McKay, Donghoon Chang, Lawrence E. Bassham, Jinkeon Kang, Noah D. Waller, John M. Kelsey, and Deukjo Hong. Status report on the final round of the NIST lightweight cryptography standardization process. NIST Internal Report 8454, National Institute of Standards and Technology, 2023. doi:10.6028/NIST.IR.8454.
- [WG13] Clifford Wolf and Johann Glaser. Yosys: A free Verilog synthesis suite. In *Austrochip Workshop on Microelectronics*, pages 47–52, 2013. URL: <https://yosyshq.net/yosys/files/yosys-austrochip2013.pdf>.
- [WLHL24] Shichang Wang, Meicheng Liu, Shiqi Hou, and Dongdai Lin. Differential-linear cryptanalysis of GIFT family and GIFT-based ciphers. *IACR Communications in Cryptology*, 1(1):13, 2024. URL: <https://cic.iacr.org/p/1/1/13>, doi:10.62056/A6N5TXOL7.