Nonlinear Invariant Attacks on Full Craft: Chosen-Tweak Key Recovery

Roberto Civino

Abstract

A nonlinear invariant is a Boolean function of the state that every round preserves, up to a constant, for a class of keys. It then relates plaintext and ciphertext with correlation one, whatever the number of rounds. Finding one is in general a problem about the whole state.

We study a family of invariants of the tweakable block cipher Craft for which this problem splits into independent pieces. Each function reads half of the state linearly and the other half through components of the S-box, one nibble at a time, so that the condition for a round to preserve it becomes one small condition per S-box, read once from the algebraic normal form of the S-box. The family can therefore be analysed by hand and enumerated by computer.

The simplest member uses a single quadratic component of the S-box and survives all thirty-two rounds on a class of 2^{108} keys. Allowing cubic components and a pattern that alternates between even and odd rounds yields an invariant that survives all thirty-two rounds for half of all keys, with suitably chosen tweaks. Both classes are determined exactly, and a computer classification shows that no member of the family is larger. Since the conditions are periodic with the tweakey, these invariants pass any number of rounds, not only the thirty-two of Craft, so increasing the round count would not remove them.

The invariants give key recovery. Almost every key of Craft satisfies a great many of them at once, and the classification yields a library of hundreds of thousands. Running the library against an unknown key and collecting the linear relations each invariant reveals recovers, for essentially every key, an average of 61.3 of the 64 low key bits with about $2^{14.5}$ chosen-tweak queries and negligible time. Over a sample of 5000 random keys every one was broken, the remaining search staying below the 124-bit security claim in every case. All programs and data that reproduce these results are available in a public repository.

Keywords: nonlinear invariant attacks, Craft, weak keys, tweakable block ciphers, key recovery

1 Introduction

Context. Linear cryptanalysis looks for a linear function of the state that survives, with some bias, from the plaintext to the ciphertext. Nonlinear invariant attacks, introduced by Todo, Leander and Sasaki [TLS16], look instead for a Boolean function of the state, not necessarily linear, that every round maps to itself up to a constant, for a class of weak keys. Such a function is preserved through the whole cipher with correlation one, independently of the number of rounds, and Todo, Leander and Sasaki used it to attack the full Scream, iScream and Midori64. Nonlinear approximations go back to Knudsen and Robshaw [KR96], who observed that a key addition does not preserve them in general,

Email: math@robertocivino.it (Roberto Civino)

This work is licensed under a "CC BY 4.0" license.
Date of this document: 2026-11-22.

Table 1: Distinguishers and attacks on Craft. For linear and invariant entries the measure is the expected squared correlation, for differential entries the probability, and for key recoveries the time complexity. ST, RT and CT stand for the single-tweak, related-tweak and chosen-tweak models, KP for known plaintexts. The library key recovery needs only $2^{14.5}$ chosen-tweak queries; the 2^{67} listed for it is the average time of the exhaustive search that then finishes the key.

type	rounds	measure	keys	model	source
linear hull	14	$2^{-62.12}$	all	ST	[BLMR19]
differential	14	$2^{-63.80}$	all	ST	[HSN+19]
truncated differential	20	2^{-63}	2^{118}	ST	[GSS+20]
circle-linear trail	14	2^{-44}	2^{108}	ST, $T = 0$	[Civ26]
circle-linear trail	18	2^{-64}	2^{110}	ST, $T = 0$	[Civ26]
key recovery	23	2^{94}	2^{112}	ST	[LR22]
key recovery	26	2^{105}	2^{108}	RT	[LR22]
key recovery	26	2^{118}	all	ST	[MNPN25]
key recovery, 2 KP	25	2^{125}	all	any tweak	[LN26]
nonlinear invariant	32	1	2^{127}	CT	Thm. 3
key recovery, library	32	2^{67}	almost all	CT	Sec. 6.3

and were revisited by Beierle, Canteaut and Leander [BCL18] through a conjugation that turns a nonlinear approximation into a linear one. Beyne gave the invariant attack its structural setting, describing invariants as eigenvectors of the correlation matrix of a round [Bey18] and then recasting linear cryptanalysis over an arbitrary finite abelian group in the geometric approach [Bey21]; in that language a correlation-one invariant is an eigenvector of eigenvalue ± 1, which is why it does not decay with the number of rounds. Beierle, Canteaut, Leander and Rotella [BCLR17] gave a criterion on the round constants that excludes the invariants preserved separately by the S-box layer and by the linear layer. In general the search for such a function is a problem about the whole state, tractable only by restricting to structured classes of functions.

Craft. Craft [BLMR19] is a 64-bit tweakable block cipher with a 128-bit key and a 64-bit tweak, designed by Beierle, Leander, Moradi and Rasoolzadeh for efficient protection against differential fault attacks. Its components are involutory, and its tweakey schedule reuses four round tweakeys with period four. The designers claim 124-bit security in the related-tweak model, and argue resistance to invariant attacks after eight rounds with the criterion of [BCLR17]. The cipher has since been studied in the differential, truncated-differential, related-key, zero-correlation, integral, rectangle and meet-in-the-middle models [EY19, HSN+19, GSS+20, MA20, HBS21, LR22, MNPN25, LN26]. For every key, the furthest attacks reach 26 of the 32 rounds in the single-tweak model [MNPN25], and 25 rounds from only two known plaintexts, in time 2^{125} [LN26]. Weak-key analyses go further on restricted classes of keys, with a 20-round distinguisher on 2^{118} keys [GSS+20] and, through an equivalent description of the cipher valid on a class of weak tweak-keys, attacks on 23 rounds with a single tweak and on 26 rounds with related tweaks [LR22]. The full cipher has been attacked only with related keys [EY19], a model for which no security is claimed. Table 1 collects these results.

Linear cryptanalysis with bibrace characters. In a previous paper [Civ26] we analysed Craft with linear approximations taken over a different group. Linear cryptanalysis measures the correlation of a cipher with the characters of the group used to define differences, and on the same set of nibbles there are other elementary abelian group

structures besides $\oplus$, those coming from binary bibraces [CF25, CF26], or equivalently from alternative operations [CBS19, CCI24]. Over such a group half of the masks are quadratic Boolean functions. For a suitable bibrace the S-box of Craft has four relations of correlation one, which form a subgroup of the dual preserved by the S-box; trails inside this subgroup pay only at MixColumn, the cost and the weak-key condition of each MixColumn are computed exactly from a quadratic form attached to the column, and the round constants impose no condition at all. The result was a family of weak-key linear distinguishers up to eighteen rounds. At fourteen rounds a single trail has squared correlation 2^{-44} on a class of 2^{108} keys, against $2^{-62.12}$ for the designers' linear hull, which holds for every key. That analysis is the starting point of the present work. The quadratic characters it builds are the functions that, once pushed from a cost at each round to correlation one at each round, become the invariants studied here; turning the one into the other is what this paper develops.

This work. We study nonlinear invariants of full Craft built one nibble at a time. Each function of the family reads the eight high nibbles of the state linearly and the eight low nibbles through components of the S-box, one nibble at a time, with a pattern of masks that may change from round to round with period one, two or four (Definition 1); we call such functions *cell separable*. For these functions the condition for a round to preserve the invariant splits into one condition per S-box application (Lemma 1). Each of them involves three nibble masks and one tweakey nibble, and the combinations that can be satisfied are computed once from the algebraic normal form of the S-box (Proposition 1). The proofs are therefore done by hand, and the members of the family with period dividing four can be enumerated by computer.

The simplest member of the family uses only the quadratic component `5` of the S-box and is the same at every round. It survives all thirty-two rounds on a class of 2^{108} keys, each with 2^{36} weak tweaks (Theorem 1). The number thirty-two plays no role here, because the same relation holds through any number of rounds; its condition repeats with the tweakey, so adding rounds, the usual way to restore a margin, would not remove it. We give it in full because it is the clearest instance of the construction, with a short proof and a weak-key class one can write down by hand, and because the whole family is the same mechanism with richer masks.

The richer masks are where the large classes come from. A member that alternates between two patterns at even and odd rounds, and uses cubic components of the S-box, survives all thirty-two rounds for every key satisfying a single linear condition, that is for half of all keys, provided the tweak is chosen in a set of 2^{30} values determined by the key. This is the second contribution, and it explains the two features that make the family reach so far. The even and the odd rounds carry separate conditions, on K_0 and on K_1 respectively, so a key is constrained on one half only; and the cubic components replace a two-bit condition on a nibble by a one-bit condition, which halves the codimension of the class. Both classes are proven exactly, and a computer classification of the family shows that no member of period dividing four has a larger class (Remark 3); this last statement rests on an exhaustive enumeration.

The third contribution, and the main one, is that these invariants are not isolated. Testing one invariant under chosen tweaks reveals the tweakey bits that its conditions involve and nothing else, appending key-guessing rounds adds nothing, and no member of the family touches the high half of the key (Propositions 4, 5 and 6). The lever is therefore the number of invariants a key satisfies, and almost every key of Craft satisfies many. The classification produces a library of hundreds of thousands of invariants (Section 6.3); running them one after another, and collecting the linear relations each one reveals, recovers the low half of the key for essentially every key. The 2^{127} invariant and its one-round shift are the worked example inside this scheme, determining 34 relations on their class with

about $2^{18.6}$ chosen-tweak queries (Theorem 4); the library goes much further.

Consequences for CRAFT: the high half of the key is out of reach of this family, so the best it can do is the 64 low key bits, after which 2^{64} keys remain to search (Proposition 6). The library approaches this limit. Over a sample of 5000 random keys it recovered on average 61.3 of the 64 low key bits with about $2^{14.5}$ chosen-tweak queries and negligible time, leaving an exhaustive search of about 2^{67} on average (Section 6.3). Every key of the sample was broken below the 124-bit security claim.

Organisation. Section 2 recalls CRAFT and the algebraic normal form of its S-box. Section 3 introduces the family and reduces the round condition to the S-box. Sections 4 and 5 treat the 2^{108} and the 2^{127} invariants, and Section 6 the key recovery.

2 Preliminaries

This section fixes the notation, recalls the specification of CRAFT, and records the algebraic normal form of its S-box, which is the only computation the rest of the paper relies on.

2.1 Notation

Let $V = \mathbb{F}_2^4$. A nibble $x \in V$ has bits $x_0, \dots, x_3$, with x_0 the least significant one, and is written in hexadecimal, so that for instance `5` is the nibble with $x_0 = x_2 = 1$ and $x_1 = x_3 = 0$. For $a, x \in V$ we write $\langle a, x\rangle = \bigoplus_i a_i x_i$ for the scalar product, and for a map $S : V \to V$ we write $a \cdot S$ for its component function $x \mapsto \langle a, S(x)\rangle$. For $k \in V$ and $f : V \to \mathbb{F}_2$ the derivative of f in direction k is $D_k f(x) = f(x) \oplus f(x \oplus k)$, and k is a *linear structure* of f if $D_k f$ is constant. The *radical* of a symmetric bilinear form ψ on V is $\operatorname{rad}\psi = \{x \in V : \psi(x, y) = 0 \text{ for every } y \in V\}$.

2.2 CRAFT

We recall the specification of [BLMR19]. The state is a 4×4 array of nibbles, $x = (x_0, \dots, x_{15}) \in V^{16}$; we index cells $0, \dots, 15$ in row-major order, so that cell i lies in row $\lfloor i/4 \rfloor$ and column $i \bmod 4$, and a column c consists of the cells $\{c, c+4, c+8, c+12\}$ read as rows $0, 1, 2, 3$. We call the cells of rows 0 and 1 the *high* cells, $\mathcal{H} = \{0, \dots, 7\}$, and those of rows 2 and 3 the *low* cells, $\mathcal{L} = \{8, \dots, 15\}$. The cipher has 32 rounds; for $0 \le t < 31$,

$$\mathcal{R}_t = \mathrm{SB} \circ \mathrm{PN} \circ \mathrm{ATK}_t \circ \mathrm{ARC}_t \circ \mathrm{MC},$$

written as a composition and therefore applied from right to left. The final round omits PN and SB, keeping the three operations that come first. When a single round is considered we write rc and tk for the vectors that ARC_t and ATK_t add to the state, so that the round reads $x' = \mathrm{SB}(\mathrm{PN}(\mathrm{MC}(x) \oplus rc \oplus tk))$. We describe now each of the five operations.

MixColumn (MC). Each column is multiplied by the involutory binary matrix

$$M = \begin{pmatrix} 1 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad \text{i.e.} \quad \begin{aligned} y_c &= x_c \oplus x_{c+8} \oplus x_{c+12}, & y_{c+8} &= x_{c+8}, \\ y_{c+4} &= x_{c+4} \oplus x_{c+12}, & y_{c+12} &= x_{c+12}. \end{aligned} \tag{1}$$

M is a *half-state* matrix, since it copies the low rows 2 and 3 and mixes only the high rows 0 and 1.

AddConstants (ARC). Two nibbles are added to cells 4 and 5 only, namely $RC_4[t]$ to cell 4 and $RC_3[t]$ to cell 5, produced by a 4-bit and a 3-bit LFSR. Both cells are high, in row 1, and every function considered in this paper reads the high cells linearly; Lemma 1 shows that for this reason the constants only shift the constant of an invariant and never enter the condition for it to survive a round.

AddTweakey (ATK). CRAFT is tweakable, meaning that besides the plaintext and the secret key it takes a third input, the tweak T, which is public and freely chosen. For us it is chosen by the adversary, since the attacks of Section 6 are in the chosen-tweak model. From $K = K_0 \| K_1$ and T, four round tweakeys are derived once and reused with period four,

$$TK_0 = K_0 \oplus T, \quad TK_1 = K_1 \oplus T, \quad TK_2 = K_0 \oplus \mathcal{Q}(T), \quad TK_3 = K_1 \oplus \mathcal{Q}(T),$$

where $\mathcal{Q}$ permutes the nibble positions of the tweak, $\mathcal{Q}(T)[i] = T[\mathcal{Q}(i)]$, with

$$\mathcal{Q} = [12, 10, 15, 5, 14, 8, 9, 2, 11, 3, 7, 4, 6, 0, 1, 13],$$

and $TK_{t \bmod 4}$ is added in round t. So K_0 is used at the even rounds and K_1 at the odd ones, each appearing twice per period but masked by two different images of the tweak.

PermuteNibbles (PN). An involutory permutation of the cells, cell i being moved to cell $\mathcal{P}(i)$, with

$$\mathcal{P} = [15, 12, 13, 14, 10, 9, 8, 11, 6, 5, 4, 7, 1, 2, 3, 0].$$

Reading it by rows, $\mathcal{P}$ exchanges rows 0 and 3 and rows 1 and 2, composing each exchange with a permutation of the columns. In particular it exchanges the high half of the state with the low half at every round.

SubBox (SB). The involutory 4-bit S-box of MIDORI [BBI+15],

$$\mathrm{Sb}_0 = [\mathtt{c}, \mathtt{a}, \mathtt{d}, 3, \mathtt{e}, \mathtt{b}, \mathtt{f}, 7, 8, 9, 1, 5, 0, 2, 4, 6],$$

applied to every cell. It has differential uniformity 4 and linearity 8.

2.3 The algebraic normal form of Sb_0

The four coordinates of Sb_0 are

$$\begin{aligned} S_0 &= x_1 \oplus x_0x_2 \oplus x_0x_3 \oplus x_0x_1x_2 \oplus x_0x_1x_3 \oplus x_1x_2x_3, \\ S_1 &= x_0 \oplus x_2 \oplus x_0x_2 \oplus x_0x_3 \oplus x_2x_3, \\ S_2 &= 1 \oplus x_0 \oplus x_3 \oplus x_0x_3 \oplus x_0x_1x_2 \oplus x_0x_1x_3 \oplus x_1x_2x_3, \\ S_3 &= 1 \oplus x_0x_1 \oplus x_1x_3 \oplus x_2x_3 \oplus x_0x_1x_3 \oplus x_1x_2x_3. \end{aligned} \tag{2}$$

Writing $C_0 = x_0x_1x_2 \oplus x_0x_1x_3 \oplus x_1x_2x_3$ and $C_3 = x_1x_3(x_0 \oplus x_2)$, the cubic part of $a \cdot \mathrm{Sb}_0$ is $(a_0 \oplus a_2)\, C_0 \oplus a_3\, C_3$. Hence $a \cdot \mathrm{Sb}_0$ is quadratic exactly for $a \in \{\mathtt{2}, \mathtt{5}, \mathtt{7}\}$ and cubic otherwise, and

$$\mathtt{5}\cdot\mathrm{Sb}_0 = 1 \oplus x_0 \oplus x_1 \oplus x_3 \oplus x_0x_2, \qquad \mathtt{2}\cdot\mathrm{Sb}_0 = S_1, \qquad \mathtt{7}\cdot\mathrm{Sb}_0 = 1 \oplus x_1 \oplus x_2 \oplus x_3 \oplus x_0x_3 \oplus x_2x_3. \tag{3}$$

3 Cell-separable invariants

This section sets up the family of invariants studied in the paper. We define the mask sequences that generate them, reduce the condition for a round to preserve one to a small condition per S-box application, and compute from the algebraic normal form of Sb_0 which conditions can be met. This is the machinery that the two invariants of Sections 4 and 5 and the library of Section 6.3 all rest on.

3.1 The family and the round condition

Informally, we want to carry one bit of information through the cipher, built the same way at every round. It is a linear function of the eight high cells, plus a correction that undoes, cell by cell, what the S-box did to the eight low cells one round earlier. We record that correction separately for each round index modulo four. The index set is $\mathbb{Z}/4$ because the tweakey schedule of Section 2.2 repeats with period four, so a mask that recurs with a period dividing four meets the same tweakey, and hence the same condition, every time it recurs.

Definition 1 (cell-separable sequence)**.** A *mask sequence* is a family $\alpha = (\alpha_t)_{t\in\mathbb{Z}/4}$ with $\alpha_t \in V^{\mathcal{H}}$, one nibble mask for each high cell and each round index modulo four. It defines, for every t, the state function

$$I_t(x) = \bigoplus_{j\in\mathcal{H}} \langle \alpha_t[j], x_j\rangle \;\oplus\; \bigoplus_{j\in\mathcal{H}} \langle \alpha_{t-1}[j], \mathrm{Sb}_0(x_{\mathcal{P}(j)})\rangle,$$

where x is the state entering MixColumn of round t. For $\beta \in V^{\mathcal{H}}$ let $m(\beta) \in V^{\mathcal{L}}$ be the masks that MixColumn induces on the low cells, $m_{8+c}(\beta) = \beta[c]$ and $m_{12+c}(\beta) = \beta[c] \oplus \beta[4+c]$. The *period* of α is the smallest $p \mid 4$ with $\alpha_{t+p} = \alpha_t$.

The two indices t and $t-1$ are staggered by one round on purpose. The low cells of the state entering round t are the S-box outputs produced in round $t-1$, so the correction that cancels them must use the mask α_{t-1} that was linear on the high cells one round earlier. This threading of consecutive rounds is what lets I_t and I_{t+1} relate through a single round, as the next lemma makes precise.

The simplest member of the family, treated in full in Section 4 below, is the period-one sequence $\alpha_t[j] = \mathsf{5}$ for all t, j. A sequence is *cell separable* in the sense that each I_t is a sum of functions of single cells.

The round condition of Lemma 1 below asks, one low cell at a time, that a three-term combination of the S-box, with an input mask, a linear middle mask and an output mask read through a key-nibble shift, be constant in the cell. The set of key nibbles that make it constant is the object of the next definition.

Definition 2 (free triples)**.** For $a, m, c \in V$ let

$$\mathcal{K}(a,m,c) = \{k \in V :\; w \mapsto \langle a, \mathrm{Sb}_0(w)\rangle \oplus \langle m, w\rangle \oplus \langle c, \mathrm{Sb}_0(w\oplus k)\rangle \text{ is constant}\}.$$

The triple (a, m, c) is *free* for k if $k \in \mathcal{K}(a, m, c)$.

Lemma 1 (round condition)**.** *Let $x' = \mathrm{SB}(\mathrm{PN}(\mathrm{MC}(x) \oplus rc \oplus tk))$ be a round with tweakey tk. Then $I_{t+1}(x') = I_t(x) \oplus \kappa$ for a constant κ if and only if, for every high cell j, with $i = \mathcal{P}(j)$,*

$$tk_i \in \mathcal{K}\big(\alpha_{t-1}[j],\; m_i(\alpha_t),\; \alpha_{t+1}[j]\big).$$

The constant is $\kappa = \bigoplus_{j\in\mathcal{H}} \langle \alpha_t[j], (rc\oplus tk)_j\rangle \oplus \bigoplus_j \kappa_j(tk_{\mathcal{P}(j)})$, where $\kappa_j(k)$ is the constant value of the triple at cell j. The high tweakey cells and the round constants enter κ only.

Proof. Write $y = \mathrm{MC}(x) \oplus rc$. Since $\mathcal{P}$ exchanges $\mathcal{H}$ and $\mathcal{L}$, a high output cell is $x'_j = \mathrm{Sb}_0(y_{\mathcal{P}(j)} \oplus tk_{\mathcal{P}(j)}) = \mathrm{Sb}_0(x_{\mathcal{P}(j)} \oplus tk_{\mathcal{P}(j)})$, because MixColumn copies the low rows and the constants live in row 1; and for a low output cell i one has $\mathrm{Sb}_0(x'_i) = y_{\mathcal{P}(i)} \oplus tk_{\mathcal{P}(i)}$ with $\mathcal{P}(i)$ high. Hence

$$I_{t+1}(x') = \bigoplus_{j\in\mathcal{H}} \langle \alpha_{t+1}[j], \mathrm{Sb}_0(x_{\mathcal{P}(j)} \oplus tk_{\mathcal{P}(j)})\rangle \oplus \bigoplus_{j\in\mathcal{H}} \langle \alpha_t[j], y_j \oplus tk_j\rangle.$$

In the second sum, $\bigoplus_j \langle \alpha_t[j], y_j\rangle = \bigoplus_{j\in\mathcal{H}} \langle \alpha_t[j], x_j\rangle \oplus \bigoplus_{i\in\mathcal{L}} \langle m_i(\alpha_t), x_i\rangle$ by the definition of m. Subtracting $I_t(x)$ leaves

$$I_{t+1}(x') \oplus I_t(x) = \bigoplus_{j\in\mathcal{H}} \Big[\langle \alpha_{t-1}[j], \mathrm{Sb}_0(x_i)\rangle \oplus \langle m_i(\alpha_t), x_i\rangle \oplus \langle \alpha_{t+1}[j], \mathrm{Sb}_0(x_i \oplus tk_i)\rangle \Big]_{i=\mathcal{P}(j)}$$

up to a constant, a sum of functions of the single cells x_i, $i \in \mathcal{L}$. It is constant if and only if each summand is, which is Definition 2; the constants are as stated. □

The lemma is about one round, but its force is that it is about *every* round at once. Once a key and tweak satisfy the per-round conditions, Lemma 1 holds at each round, so the relation $I_0(P) \oplus I_R(\cdot) = \text{const}$ survives any number of rounds R, with correlation one. The same invariant would pass a version of the cipher with more rounds, and this holds for every member of the family.

Lemma 2 (structure of the free triples)**.** *The free triples have the following structure.*

(i) For fixed k the free triples form a subgroup G_k of V^3, and $(a,m,c) \in G_k$, $(a,m,c') \in G_k$ imply $c = c'$. For every k, $(a,m,a) \in G_k$ iff $D_k(a \cdot \mathrm{Sb}_0) = \langle m, \cdot\rangle \oplus$ const.

(ii) Every nonempty $\mathcal{K}(a,m,c)$ is a coset of a subspace of V.

(iii) If $a \cdot \mathrm{Sb}_0$ is quadratic with associated bilinear form ψ_a, then for every k there is a unique $m_k(a)$ with $(a, m_k(a), a) \in G_k$, namely $\langle m_k(a), x\rangle = \psi_a(x,k)$, and $\mathcal{K}(a,m,a) = \{k : m_k(a) = m\}$ is either empty or a coset of $\operatorname{rad}\psi_a$, of size four.

(iv) If $a \cdot \mathrm{Sb}_0$ is cubic, $\mathcal{K}(a,m,a) \neq \emptyset$ only if k is a linear structure of $a \cdot \mathrm{Sb}_0$, i.e. $D_k(a \cdot \mathrm{Sb}_0)$ is affine.

Proof. (i) The sum of two constant functions is constant, and $(0,0,c)$ is free only for $c = 0$ since Sb_0 has no affine component. (ii) is the standard fact that $\{k : D_k f \text{ has prescribed affine part}\}$ is empty or a coset: if k, k' are in $\mathcal{K}(a,m,c)$ then, composing the two constancy identities, $\langle c, \mathrm{Sb}_0(w \oplus k)\rangle \oplus \langle c, \mathrm{Sb}_0(w \oplus k')\rangle$ is constant, and the set of $\delta = k \oplus k'$ with $D_\delta(c \cdot \mathrm{Sb}_0)$ constant is a subspace U_c; conversely $\mathcal{K}(a,m,c) \supseteq k \oplus U_c$. (iii) For quadratic $f = a \cdot \mathrm{Sb}_0$, $D_k f(x) = \psi_a(x,k) \oplus f(k)$, which is affine with linear part $\psi_a(\cdot,k)$; $k \mapsto \psi_a(\cdot,k)$ is linear with kernel $\operatorname{rad}\psi_a$ of dimension two, since the quadratic forms in (3) have rank two. (iv) is immediate from (i). □

Proposition 1 (the free triples of Sb_0)**.** *Write $\tau(a) = (a_0 \oplus a_2,\, a_3) \in \mathbb{F}_2^2$.*

(i) Quadratic components, $\tau(a) = (0,0)$, $a \in \{2,5,7\}$. The quadratic part of $a \cdot \mathrm{Sb}_0$, its radical $\operatorname{rad}\psi_a$ and the masks $m_k(a)$ are as follows:

a	*quadratic part*	$\operatorname{rad}\psi_a$	$m_k(a)$
5	x_0x_2	{0, 2, 8, a}	$k_2\,1 \oplus k_0\,4 \;\in \{0,5,4,1\}$
2	$x_0x_2 \oplus x_0x_3 \oplus x_2x_3$	{0, 2, d, f}	$(k_2 \oplus k_3)1 \oplus (k_0 \oplus k_3)4 \oplus (k_0 \oplus k_2)8 \in \{0,5,9,c\}$
7	$x_0x_3 \oplus x_2x_3$	{0, 2, 5, 7}	$k_3\,5 \oplus (k_0 \oplus k_2)8 \in \{0,5,8,d\}$

In particular $\mathcal{K}(5,0,5) = \mathrm{Ann} = \{0,2,8,a\}$, $\mathcal{K}(5,5,5) = \{5,7,d,f\}$, $\mathcal{K}(5,4,5) = \{1,3,9,b\}$, $\mathcal{K}(5,1,5) = \{4,6,c,e\}$ and $\mathcal{K}(2,5,2) = \{5,7,8,a\}$.

(ii) Cubic components. The linear structures of $a \cdot \mathrm{Sb}_0$ *are* $\{0, d\}$ *if* $\tau(a) = (1, 0)$, $\{0, 5\}$ *if* $\tau(a) = (0, 1)$ *and* $\{0, 8\}$ *if* $\tau(a) = (1, 1)$. *Consequently* $\mathcal{K}(a, 0, a) = \{0\}$ *and* $\mathcal{K}(a, m^*(a), a) = \{k^*(a)\}$ *for the unique nonzero linear structure* $k^*(a)$ *and the linear part* $m^*(a)$ *of* $D_{k^*(a)}(a \cdot \mathrm{Sb}_0)$*; all other* $\mathcal{K}(a, m, a)$ *are empty. In particular*

$$\mathcal{K}(4, b, 4) = \{d\},\ \ \mathcal{K}(1, e, 1) = \{d\},\ \ \mathcal{K}(b, 2, b) = \mathcal{K}(e, 2, e) = \{8\},\ \ \mathcal{K}(8, a, 8) = \{5\},$$

and $\mathcal{K}(b, 0, b) = \mathcal{K}(e, 0, e) = \{0\}$.

Proof in Appendix A.

Remark 1 (the role of the linear layers). An exact invariant is not a correlation that attenuates by composition. In the correlation-matrix picture of [Bey18, Bey21] it is an eigenvector with eigenvalue ± 1 on its weak-key class, and a sign raised to any power is still a sign. So once the round condition of Lemma 1 holds at every round there is no decay left to account for, and the question is only why that condition can hold for all thirty-two rounds at once. Three further features of Craft, none of them a statement about the S-box alone, are what make this possible.

First, PN is an involution that exchanges $\mathcal{H}$ and $\mathcal{L}$ as whole halves, cell by cell, which is why a cell-separable sequence maps to a cell-separable sequence of the same shape every round (proof of Lemma 1) instead of spreading into ever more cells; a denser, MDS-style diffusion would not have this property, and Definition 1's family would not be closed under one round. Second, MC does mix the nibbles of a column, a high output cell depending on the low cells of its own column, but it mixes them linearly and never across columns. A function that is a sum of single-cell terms is therefore carried to another such sum, through the explicit mask map m of Definition 1, with no cross-cell product appearing; this is what keeps Lemma 2 and Proposition 1 a per-cell, sixteen-value affair instead of an exponential one. Third, the tweakeys $TK_0, \ldots, TK_3$ repeat with period four and are injected by plain XOR with no further nonlinear key mixing, so the round condition, once met for one period, is automatically met for all thirty-two rounds without accumulating new, independent constraints; this is why the weak-key classes of Theorems 1 and 3 below are a sizeable fraction of the key space rather than vanishing with the number of rounds.

The S-box supplies only a fourth ingredient: that three of its fifteen nonzero component functions are quadratic with a two-dimensional radical (Proposition 1(i)), which is what makes the round condition satisfiable at all, by a nonempty coset of keys, rather than forcing correlation strictly below one everywhere. All four features have to hold together; removing any one of them removes the attack, not merely weakens it.

This is also why the designers' own resistance proof does not cover the family. They argue resistance to invariant attacks after eight rounds [BLMR19, Sec. 5.8], using the round-constant criterion of [BCLR17]. That criterion considers functions invariant under the S-box layer and under each linear and tweakey layer *separately*, and shows that the smallest $(\mathrm{PN} \circ \mathrm{MC})$-invariant subspace containing the relevant constant differences is the whole space, so within that framework an invariant must be affine while the S-box has no linear component. No I_t of this paper is of that kind. The S-box layer maps I_t to the function with its linear and its quadratic (or cubic) half exchanged, and only the full round, through PN, brings it back to itself up to a key-dependent constant. The round constants never interfere, for the reason given in the proof of Lemma 1: they touch only cells 4 and 5, both high, and every sequence of Definition 1 reads a high cell linearly, so a constant can only shift κ and never enter the survival condition.

4 The period-one invariant and its 2^{108} class

This section derives, from the general machinery above, the simplest member of the family, the all-5 sequence. The free-triple values it needs are already available from Proposition 1(i),

so no further S-box computation is needed. Like every member of the family, the resulting invariant survives any number of rounds, by the paragraph after Lemma 1; we state it at the full thirty-two.

Theorem 1 (the period-one invariant and the class $\mathcal{C}_{108}$)**.** *Let $\alpha_t[j] = \mathtt{5}$ for every $t \in \mathbb{Z}/4$ and every $j \in \mathcal{H}$. Since α is constant, I_t does not depend on t; write*

$$I(x) = \bigoplus_{j\in\mathcal{H}} \langle \mathtt{5}, x_j \rangle \;\oplus\; \bigoplus_{i\in\mathcal{L}} \langle \mathtt{5}, \mathrm{Sb}_0(x_i) \rangle,$$

linear on the high cells and quadratic, through $\mathtt{5}\cdot\mathrm{Sb}_0$, on the low ones. Then:

(i) a round carries I to itself up to a constant if and only if, cell by cell, $tk_i \in \{\mathtt{5},\mathtt{7},\mathtt{d},\mathtt{f}\}$ for $i = 8,\dots,11$ and $tk_i \in \mathrm{Ann}$ for $i = 12,\dots,15$; the high tweakey cells and the round constants are free;

(ii) propagating this condition through the four tweakeys $TK_0,\dots,TK_3$ (the same condition is imposed at every one of them, since the sequence has period one) gives: a weak tweak satisfies $T[i] \oplus T[\mathcal{Q}(i)] \in \mathrm{Ann}$ for $i = 8,\dots,15$, a class of 2^{48} tweaks; given a weak tweak, the coset of $K_0[i]$ and of $K_1[i]$ modulo Ann (for $i = 12,\dots,15$) or modulo $\{\mathtt{5},\mathtt{7},\mathtt{d},\mathtt{f}\}$ (for $i = 8,\dots,11$) is fixed, a further 2^{96} keys; and a key admits some weak tweak if and only if $K_0[i] \oplus K_1[i] \in \mathrm{Ann}$ for $i = 8,\dots,15$, together with the two further relations $K_1[8] \oplus K_1[11] \in \mathrm{Ann}$ and $K_1[13] \oplus K_1[15] \in \mathrm{Ann}$ induced by $\mathcal{Q}$: a class $\mathcal{C}_{108}$ of 2^{108} keys, each with 2^{36} weak tweaks, so 2^{144} weak pairs in all;

(iii) for every weak pair (K,T), every plaintext P and its ciphertext C under full 32-round Craft*,*

$$I(P) \oplus G(C) = \mu, \qquad G(C) = \bigoplus_{j\in\mathcal{H}} \langle \mathtt{5}, (\mathrm{MC}\,C)_j \rangle \;\oplus\; \bigoplus_{i\in\mathcal{L}} \langle \mathtt{5}, \mathrm{Sb}_0(C_i) \rangle \;\oplus\; \bigoplus_{i=8}^{11} \langle \mathtt{5}, C_i \rangle,$$

for a constant μ that depends neither on K nor on T.

Proof in Appendix A.

Proposition 2 (the three components of $\mathcal{C}_{108}$)**.** *The invariant of Theorem 1 is a sum of three cell-separable invariants, each itself a period-one invariant with a nonempty class: the mask $\mathtt{5}$ on high cells $\{1,5\}$ (column 1) alone, on $\{3,7\}$ (column 3) alone, or on $\{0,2,4,6\}$ (columns 0 and 2 together, which are exchanged by $\mathcal{P}$ and so cannot be separated). The first two have classes of 2^{124} keys each (rank 12, 4 pure key relations); the third has a class of 2^{118} keys (rank 24, 10 pure key relations). These account for $4+4+10 = 18$ of the twenty relations that cut out $\mathcal{C}_{108}$. The remaining two arise from requiring a single common tweak for the three invariants at once, which couples column 3 to columns $0,2$ through the schedule identity $T[11] = T[\mathcal{Q}(8)]$ and adds the relation $K_1[8] \oplus K_1[11] \in \mathrm{Ann}$ of Theorem 1(ii). Hence $\mathcal{C}_{108}$ is strictly smaller than the intersection of the three classes, of codimension $18+2 = 20$.*

Proof. In the notation of Proposition 3, each component is the period-one sequence with $A = B$ equal to the stated support, and all three satisfy its two inclusions $\mu(\mathrm{supp}\,A) \subseteq \mathrm{supp}\,B$, $\mu(\mathrm{supp}\,B) \subseteq \mathrm{supp}\,A$, so each is a cell-separable invariant. Its codimension is the one that proposition gives, $2|\,\mathrm{supp}\,A|$ plus one relation modulo Ann for each active coupling: four for column 1 (constrained cells $\mathcal{P}(1) = 12$, $\mathcal{P}(5) = 9$) and four for column 3 (cells 14, 11), with neither coupling active; and $8+2 = 10$ for columns $0,2$ (cells $15,13,10,8$), where the coupling $T[13] = T[\mathcal{Q}(15)]$ is active because both cells 13 and 15 are constrained. These give $4+4+10 = 18$ relations. The full all-$\mathtt{5}$ invariant constrains all eight low cells at once, so its other coupling $T[11] = T[\mathcal{Q}(8)]$ is active as well, cell 8 lying in columns $0,2$ and cell 11 in column 3; it adds the relation $K_1[8] \oplus K_1[11] \in \mathrm{Ann}$, for a total codimension 20, as computed directly in Theorem 1(ii). □

5 The period-two invariant and the class of 2^{127} keys

This section builds the invariant that produces the large classes. It alternates two mask patterns between even and odd rounds and, unlike the period-one invariant, uses the cubic components of the S-box; the reward is a weak-key class of half the key space. We first exhibit the sequence and its per-round conditions (Theorem 2), then determine the class exactly (Theorem 3), and finally, in Section 5.1, place it inside the purely quadratic skeleton it upgrades, which explains where it comes from.

Theorem 2 (period-two invariant)**.** *Let $A, B \in V^{\mathcal{H}}$ be*

$$A = (\mathtt{0},\mathtt{5},\mathtt{b},\mathtt{0},\ \mathtt{0},\mathtt{5},\mathtt{b},\mathtt{0}), \qquad B = (\mathtt{0},\mathtt{0},\mathtt{0},\mathtt{0},\ \mathtt{4},\mathtt{2},\mathtt{0},\mathtt{0}),$$

and let $\alpha_t = A$ for even t and $\alpha_t = B$ for odd t. The two state functions are

$$I_{\mathrm{ev}}(x) = \langle \mathtt{5}, x_1\rangle \oplus \langle \mathtt{b}, x_2\rangle \oplus \langle \mathtt{5}, x_5\rangle \oplus \langle \mathtt{b}, x_6\rangle \oplus \langle \mathtt{2}, \mathrm{Sb}_0(x_9)\rangle \oplus \langle \mathtt{4}, \mathrm{Sb}_0(x_{10})\rangle,$$
$$I_{\mathrm{od}}(x) = \langle \mathtt{4}, x_4\rangle \oplus \langle \mathtt{2}, x_5\rangle \oplus \langle \mathtt{5}, \mathrm{Sb}_0(x_{12})\rangle \oplus \langle \mathtt{b}, \mathrm{Sb}_0(x_{13})\rangle \oplus \langle \mathtt{5}, \mathrm{Sb}_0(x_9)\rangle \oplus \langle \mathtt{b}, \mathrm{Sb}_0(x_8)\rangle.$$

Round t carries I_t to I_{t+1} up to a constant if and only if

$$t \text{ even}: \quad tk_{10} = \mathtt{d}, \quad tk_9 \in \{\mathtt{5},\mathtt{7},\mathtt{8},\mathtt{a}\};$$
$$t \text{ odd}: \quad tk_8 = \mathtt{0}, \quad tk_9 \in \{\mathtt{0},\mathtt{2},\mathtt{8},\mathtt{a}\}, \quad tk_{12} \in \{\mathtt{1},\mathtt{3},\mathtt{9},\mathtt{b}\}, \quad tk_{13} = \mathtt{8}.$$

The other twelve tweakey cells and the round constants are free.

Proof in Appendix A.

The functions have algebraic degree three on cells 8, 10 and 13; the cubic components of the S-box are what distinguishes this sequence from the quadratic invariant of Section 4 (see Proposition 3).

Theorem 3 (the class $\mathcal{C}_{127}$)**.** *Let*

$$\mathcal{C}_{127} = \{\, K = K_0 \| K_1 :\ \langle \mathtt{5}, K_0[9]\rangle = \langle \mathtt{5}, K_1[9]\rangle \,\},$$

a set of 2^{127} keys. For $K \in \mathcal{C}_{127}$ let $\delta = K_0[9] \oplus K_1[9]$ and define a tweak $T(K)$ by

$$T[10] = T[7] = K_0[10] \oplus \mathtt{d}, \qquad T[8] = T[11] = K_1[8], \qquad T[13] = T[0] = K_1[13] \oplus \mathtt{8},$$
$$T[12] = T[6] = K_1[12] \oplus \mathtt{1}, \qquad T[9] = T[3] = K_1[9] \oplus \mathtt{8} \cdot (1 \oplus \langle \mathtt{9}, \delta\rangle),$$

the cells $1, 2, 4, 5, 14, 15$ being arbitrary. Then:

(i) for every $K \in \mathcal{C}_{127}$, every tweak in $T(K) \oplus W$, where W is the subspace of tweaks with arbitrary cells $1, 2, 4, 5, 14, 15$, cells $12, 6$ in Ann, and cells $9, 3$ in $\{\mathtt{0}, \mathtt{2}\}$, satisfies the conditions of Theorem 2 at all four tweakeys; there are 2^{30} such tweaks and no others;

(ii) a key admits a tweak satisfying the conditions if and only if $K \in \mathcal{C}_{127}$; the weak pairs (K, T) number 2^{157};

(iii) for every weak pair, every plaintext P and its ciphertext C under full 32-round CRAFT,

$$I_{\mathrm{ev}}(P) \oplus I_{\mathrm{od}}(C) = \gamma, \qquad \gamma = \bigoplus_{t \text{ odd}} \langle \mathtt{4}, RC_4[t]\rangle \oplus \bigoplus_{t \text{ even}} \langle \mathtt{5}, RC_3[t]\rangle \oplus \bigoplus_{t \text{ odd}} \langle \mathtt{2}, RC_3[t]\rangle,$$

a constant that depends neither on the key nor on the tweak.

Proof in Appendix A.

Remark 2 (the other sequences of the same class). The sequence with $A' = (0,5,e,0,0,5,e,0)$ and $B' = (0,0,0,0,1,2,0,0)$ satisfies Theorem 2 with $tk_{12} \in \{4,6,c,e\}$ in place of $\{1,3,9,b\}$ (the triples $(1,e,1)$, $(e,2,e)$, $(e,0,e)$, $(5,1,5)$ replace their counterparts) and has the same class. Shifting either sequence by one round, $\alpha'_t = \alpha_{t+1}$, exchanges the roles of even and odd rounds, hence of K_0 and K_1; the class is again $\mathcal{C}_{127}$, since the relation is symmetric. These four sequences are exactly the largest-class sequences the classification finds at period two (Remark 3).

5.1 Period-two sequences with quadratic components

This subsection explains where the 2^{127} sequence comes from. We isolate the purely quadratic members of the period-two family, the $\{0,5\}$ skeleton, classify which mask patterns give an invariant, and then show that the 2^{127} sequence is one of them with two cells upgraded to cubic components, the upgrade being exactly what drops the class relation from two bits to one. The 2^{127} sequence is thus best understood as the cubic upgrade of a purely quadratic skeleton. Identify a subset $A \subseteq \mathcal{H}$ with its indicator in $\mathbb{F}_2^8$, written (r,s) with r the row-0 part and s the row-1 part, and let $\mu : \mathbb{F}_2^8 \to \mathbb{F}_2^8$ be the linear map

$$\mu(r,s) = (\sigma(r \oplus s),\ \pi(r)), \qquad \sigma(c) = c+1 \bmod 4, \qquad \pi = (0\ 2),$$

which sends the induced low masks back to high cells: $\mu(H)[\mathcal{P}(i)] = m_i(H)$.

Proposition 3 (the $\{0,5\}$ skeleton)**.** *Let α be a period-two sequence, $\alpha_{\mathrm{ev}} = A$, $\alpha_{\mathrm{od}} = B$, all of whose nonzero entries are quadratic components, $A[j], B[j] \in \{0,2,5,7\}$, and suppose every triple of Lemma 1 is free for some k. Then every nonzero entry equals 5, and with* $\operatorname{supp} A, \operatorname{supp} B \subseteq \mathcal{H}$*,*

$$\mu(\operatorname{supp} A) \subseteq \operatorname{supp} B, \qquad \mu(\operatorname{supp} B) \subseteq \operatorname{supp} A.$$

Conversely every pair of subsets with these two inclusions gives a period-two sequence whose conditions are $tk_{\mathcal{P}(j)} \in$ Ann or $\in 5 \oplus$ Ann according to the bit $\mu(\operatorname{supp} A)[j]$ (even rounds, $j \in \operatorname{supp} B$) or $\mu(\operatorname{supp} B)[j]$ (odd rounds, $j \in \operatorname{supp} A$). Its weak-key class has codimension $2|\operatorname{supp} A \cap \operatorname{supp} B|$, plus one further relation modulo Ann *for each of the couplings $T[11] = T[\mathcal{Q}(8)]$ and $T[13] = T[\mathcal{Q}(15)]$ whose two tweakey cells, $(8,11)$ resp. $(13,15)$, are both constrained.*

Proof in Appendix A.

Example 1. $\operatorname{supp} A = \{4,5\}$, $\operatorname{supp} B = \{1,2,5,6\}$ is a solution: $\mu(\{4,5\}) = \{1,2\}$ and $\mu(\{1,2,5,6\}) = \{4,5\}$. The shared cell is 5, i.e. tweakey cell $\mathcal{P}(5) = 9$, with the relation $K_0[9] \oplus K_1[9] \in 5 \oplus$ Ann: a class of 2^{126} keys. Replacing 5 by b on cells $2,6$ of B and by $4,2$ on cells $4,5$ of A gives exactly the sequence of Theorem 2 (with the roles of A and B exchanged), and the point of the upgrade is at cell 9: the two conditions there are now modulo $\operatorname{rad}\psi_2 = \{0,2,d,f\}$ and modulo Ann, two *different* two-dimensional subspaces of the three-dimensional space $\ker\langle 5,\cdot\rangle$, so their sum is that whole space and the relation drops to the single bit $\langle 5, K_0[9] \oplus K_1[9]\rangle = 0$. The period-one solutions $\operatorname{supp} A = \operatorname{supp} B = \{1,5\}$, $\{3,7\}$, $\{0,2,4,6\}$ are the three components of $\mathcal{C}_{108}$ (Proposition 2); $\operatorname{supp} A = \{0,4\}$, $\operatorname{supp} B = \{2,6\}$ has disjoint supports but is coupled through $T[13] = T[\mathcal{Q}(15)]$, with relation $K_0[13] \oplus K_1[15] \in$ Ann.

Remark 3 (classification). The program `classify.c` enumerates, by branch and bound, every cell-separable sequence of a given period dividing four with a nonempty weak-key class, and can be run with a class-size threshold θ that prunes a branch only once no completion of it could still reach 2^θ; a run at threshold θ is therefore a complete census of every class of size at least 2^θ. At period one and period two the threshold costs nothing

and was set to zero, so `p1.txt` and `p2.txt` are unconditionally exhaustive; their largest classes are 2^{124} and 2^{127}, the latter the four sequences above and no other, re-derived by hand (Theorem 3), together with the 2^{108}, 2^{124} and 2^{118} classes of the period-one case (Theorem 1, Proposition 2). At period four the unpruned search is far larger (Remark 4), so `classify.c` was run with threshold $\theta = 120$, producing `p4_120.txt`; it is complete for every class of size at least 2^{120}, and its largest class of period exactly four is 2^{126}. Since $2^{126} < 2^{127}$ and the period-one and period-two censuses are complete, 2^{127} is the largest weak-key class of any cell-separable sequence of period dividing four, which is what Theorem 3 asserts. The three files `p1.txt`, `p2.txt` and `p4_120.txt` are the library used in Section 6.3.

Remark 4 (the cost of a complete period-four census). We did not run an unthresholded period-four census, and nothing here needs one. The two classes used for key recovery, $\mathcal{C}_{108}$ and $\mathcal{C}_{127}$, are period one and two and already exhaustively classified, and any period-four class the threshold-120 run misses is below 2^{120}, so it changes neither Theorem 3 nor any attack of Section 6. A full census would also be out of reach.

6 Key recovery

We now turn the invariants into key recovery. We first fix what a single invariant can and cannot give a chosen-tweak adversary, which bounds the whole approach at the 64 low key bits and sets up Proposition 4, the tool the rest of the section runs on. We then give one key recovery in full, on the class $\mathcal{C}_{127}$, proven by hand, and finally the library procedure that runs the whole family at once and reaches almost every key. This last part is the main result, and the $\mathcal{C}_{127}$ recovery is the worked example it generalises, so a reader after the attack alone can go from the next subsection straight to Section 6.3.

6.1 What an invariant can and cannot give

Fix a sequence α with a nonempty class. Its conditions form an affine system $A_K K \oplus A_T T = b$ over $\mathbb{F}_2$ in the $128 + 64$ bits of (K, T). Row reducing with the tweak columns first splits it into *tweak rows*, each expressing one tweak bit as an affine function of the key, and *key relations*, free of T, which define the class. Let $\rho(\alpha)$ be the rank of the key parts of all rows and $c(\alpha)$ the number of key relations; for the sequence of Theorem 2, $\rho = 18$ and $c = 1$, and for the sequence of Theorem 1, $\rho = 32$ and $c = 20$.

Proposition 4 (information of one invariant)**.** *Let the adversary query encryptions of random plaintexts under tweaks of her choice. For a key in the class of α she can determine the values of the $\rho(\alpha)$ independent linear functionals of the key appearing in the system, and no other function of the key, by testing the $2^{\rho(\alpha)-c(\alpha)}$ candidate tweak patterns. Each wrong pattern is rejected after two queries on average; each pattern is accepted only after N agreeing queries, so the expected number of queries is about $2^{\rho-c+1} + N$ and the false-positive probability at most $2^{\rho-c-N+1}$. If the key is outside the class every pattern is rejected, which reveals that the c relations do not all hold; when $c = 1$ this is one further linear bit.*

Proof. The tweak rows determine T modulo the free tweak bits from the values of the key functionals they contain; two keys with the same values are indistinguishable by the invariant, since the invariant holds for both with the same tweaks and fails for both with all others (Lemma 1 is an equivalence). The key relations are implied by membership. The cost and error estimates are those of a constant-versus-balanced test; a wrong pattern gives a function that is not constant, so its restriction to random inputs disagrees with the first value with probability close to $1/2$ per query. □

The accepted pattern is more than a list of key bits. It hands the adversary a concrete tweak under which the invariant has correlation one for the secret key, a *weak tweak*. Each weak key has many of them for a given sequence, $2^{p-(128-c)}$ where p is the number of weak pairs of the sequence, for instance 2^{30} for the 2^{127} class and 2^{36} for the 2^{108} class, and the pattern search returns one of them.

Proposition 5 (appended rounds)**.** *Let $R' < 32$ and consider using the invariant on rounds $0, \dots, R'-1$ and guessing tweakey material to compute $I_{R'}$ from the ciphertext through the last $32 - R'$ rounds. Two tweakey guesses that agree on the condition bits of Lemma 1 for rounds $R', \dots, 31$ produce the same function of the ciphertext up to a constant. Hence the key-guessing rounds determine exactly the condition bits of those rounds, which for $32 - R' \geq 4$ are the $\rho(\alpha)$ functionals of Proposition 4, and nothing else.*

Proof. If a guess satisfies the conditions of round t, Lemma 1 gives $I_t(x^{(t)}) = I_{t+1}(x^{(t+1)}) \oplus \kappa_t$, where $x^{(t)}$ is the state computed backwards with the guess; iterating from 31 down to R', $I_{R'}$ of the partially decrypted state equals the key-free function $I_{\mathrm{od}}(C)$ (or its even analogue) up to a constant. The partial decryption therefore depends on the guess only through the condition bits, and since the conditions are periodic with the schedule, a window of four or more rounds imposes the full set of conditions. □

Proposition 6 (blindness to the high key)**.** *For every cell-separable sequence, the high tweakey cells enter only the constants κ_t of Lemma 1, and over the 32 rounds their contributions cancel. Consequently no adversary using only invariants of this family, on any number of tweaks, obtains any information on $K_0[0..7], K_1[0..7]$, and the recoverable information is contained in the 64 low key bits.*

Proof. The first statement is Lemma 1: I_{t+1} is linear on the high cells and the high tweakey is added after MixColumn, so it appears as $\bigoplus_j \langle \alpha_t[j], tk_j \rangle$. Over the thirty-two rounds each TK_r is used eight times with the same mask α_r (the period divides four), so these terms occur an even number of times. The conditions of Lemma 1 involve low tweakey cells only. □

Propositions 5 and 6 are negative results, but they are what makes the design of the attacks straightforward: the only lever is the set of condition bits, and the only way to enlarge it is to use several invariants.

6.2 A proven key recovery on $\mathcal{C}_{127}$

This is the example the library generalises, and the only key-recovery figure in this paper that is a theorem rather than a measurement.

A single invariant already gives something. The pattern test of Proposition 4 applied to the sequence of Theorem 2 determines $K_0[10]$, $K_1[8]$, $K_1[13]$ entirely, $K_1[12]$ modulo Ann, and the three bits $\langle \mathtt{5}, K_0[9] \rangle$, $\langle \mathtt{9}, K_0[9] \rangle$, $\langle \mathtt{1}, K_1[9] \rangle$, that is 18 linear relations of which one is the class relation, with about 2^{18} queries. Running it together with the one-round shift of the sequence nearly doubles the number of key relations recovered.

Theorem 4 (twin pair)**.** *Running the pattern test of Proposition 4 on the sequence of Theorem 2 and then on its one-round shift determines, for every $K \in \mathcal{C}_{127}$, the 18 relations above together with $K_1[10]$, $K_0[8]$, $K_0[13]$, $K_0[12]$ modulo* Ann, *$\langle \mathtt{9}, K_1[9] \rangle$ and $\langle \mathtt{1}, K_0[9] \rangle$, that is 34 independent linear relations on K in all, with $2^{17} + 2^{16}$ patterns and about $2^{18.6}$ queries. The remaining key space has 2^{94} elements. For $K \notin \mathcal{C}_{127}$ the first test fails and reveals $\langle \mathtt{5}, K_0[9] \oplus K_1[9] \rangle = 1$, so one linear bit is obtained for every key, at cost 2^{18} queries.*

Proof. The 18 functionals of the first test are those read in the proof of Theorem 3, $4+4+4+2+2+2 = 18$ bits of key, among which $\langle 5, K_0[9]\rangle \oplus \langle 5, K_1[9]\rangle$ is the class relation; so $\rho = 18$, $c = 1$ and the number of patterns is 2^{17}. The shifted sequence exchanges the roles of K_0 and K_1 in that proof, so its 18 functionals are the listed ones, of which $\langle 5, K_0[9]\rangle$ and $\langle 5, K_1[9]\rangle$ are already known after the first test; hence 16 bits remain and the union has rank $18 + 16 = 34$. The costs and error estimates are Proposition 4, and the failure statement is its last sentence. □

6.3 Several invariants: the library procedure

This is the central part of the paper. A single invariant reveals only the tweakey bits of its own conditions (Proposition 4), and that is a handful of bits. What makes the family effective is that almost every key satisfies a great many invariants at once, and the bits they reveal accumulate up to the barrier of the high key half (Proposition 6). To exploit this we need the whole family in hand, so we first describe how the library is built and why it covers essentially every key, and then the procedure that reads a key out of it.

Building the library. The members of the family with period dividing four are enumerated by the program `classify.c`. A sequence is fixed by its masks $\alpha_t[j]$, with $t \in \mathbb{Z}/4$ and $j \in \mathcal{H}$, so the search space is the set of mask assignments, and a branch-and-bound walk over it keeps only the assignments for which every triple of Lemma 1 is free, that is for which the sequence is an invariant with a nonempty class. For each surviving sequence the program forms the affine system of its conditions, row reduces it as in Section 6 into the tweak rows and the pure key relations, and records the class size 2^{128-c} together with the number of weak pairs. The walk can be run with a class-size threshold θ that prunes a branch as soon as no completion of it could reach 2^θ, so that a run at threshold θ is a complete census of every class of size at least 2^θ (Remark 3). At period one and period two the threshold costs nothing and the census is unconditional, giving 568 and 340 716 sequences in the files `p1.txt` and `p2.txt`; at period four the unpruned search is far larger, so the file `p4_120.txt` is the threshold-120 census, 38 271 sequences of period dividing four. The attack is run with this file alone; its sequences are the ones with a large class, which are the ones a random key is likely to satisfy. The full lists `p1.txt` and `p2.txt` are kept for Remark 3, whose maximality statement, that no cell-separable sequence has a class larger than 2^{127}, needs the complete period-one and period-two censuses. Each sequence is stored as its masks with the pre-computed (c, pairs) pair, so the attack reloads them without recomputing the systems.

Why the library covers almost every key. A uniform key lies in the class of a given sequence of codimension c with probability 2^{-c}, and the library contains sequences of small codimension in large number, so a random key lies in many of their classes at once. We checked this directly on the cipher. For a sample of random keys, each key was matched against the library, a concrete weak tweak was built for every sequence whose class contains it, and the invariant $I_0(P) \oplus G(C)$ was confirmed constant on the full thirty-two-round cipher; every sampled key was found to lie in the class of a large number of sequences, with enough independent conditions among them to pin down, on a sample of 5000 keys, on average 61.3 of the 64 low key bits (Section 6.3). Every one of the 5000 keys was covered. This is why the procedure below is not an attack on a distinguished class but on essentially the whole key space.

The procedure. Proposition 4 says how to read the condition bits of any one invariant. Any invariant whose class contains the secret key contributes its functionals, and the union of these functionals, over all the invariants the key satisfies, is the information available.

The following adaptive procedure, implemented in `craft_invattack attack`, collects it, cheapest invariants first. For each sequence it tries whose class could contain the key, the procedure searches a small set of tweak patterns and, when the key is in the class, finds one that makes the invariant hold, a weak tweak in the sense above; the key bits that pattern fixes are then added to E.

Algorithm 1 (library key recovery).
Input: a library $\mathcal{S}$ of sequences with nonempty class, a budget $d_{\max}$ on the number of guessed bits per sequence, a query budget.
State: a system E of linear equations on the key, initially empty.

1. Sort $\mathcal{S}$ by increasing $c(\alpha)$, then by increasing number of tweak rows.
2. For $d = 0, 1, \ldots, d_{\max}$, repeat until no progress: for each untried $\alpha \in \mathcal{S}$,
   (a) reduce the key relations of α modulo E; if one is contradicted, discard α; let u be the number of independent relations left unknown;
   (b) reduce the key parts of the tweak rows modulo E; let d_α be the rank of the residuals; if $d_\alpha = u = 0$ discard α (nothing to learn); if $d_\alpha > d$ postpone α;
   (c) enumerate the 2^{d_α} values of the residual functionals; for each, set the tweak from the tweak rows and test constancy of $I_0(P) \oplus G(C)$ on up to $N = 40$ random plaintexts, stopping at the first disagreement;
   (d) on acceptance add the d_α guessed values and the u relations to E; on rejection of all values, if $u = 1$ add the negated relation to E.
3. Output E; the key is then found by exhaustive search over the complement of E, of size $2^{128-\operatorname{rank} E} \geq 2^{64}$.

Results

The rank of E at the end depends on which classes contain the secret key, and is measured. The delivered program runs Algorithm 1 on random keys and prints, for each one, the rank of E, the number of chosen-tweak queries, and whether E is consistent with the true key; it then aggregates the histogram of $\operatorname{rank} E$ over the sample. One lower bound holds without running the library. On $\mathcal{C}_{127}$, Theorem 4 already gives $\operatorname{rank} E \geq 34$ at cost $2^{18.6}$. And $\operatorname{rank} E \leq 64$ always, by Proposition 6, so the search that finishes the key costs at least 2^{64}.

Over a sample of 5000 random keys (seed 1), the library `p4_120.txt` recovered on average 61.3 of the 64 low key bits, with the count ranging from 48 to the full 64 and reaching 64 for 1216 of the 5000 keys, using on average about $2^{14.5}$ chosen-tweak queries per key and negligible time; E was consistent with the true key in all 5000 cases. The remaining search is then about 2^{128-r} for r recovered bits, so about 2^{67} on average and 2^{64} for the keys whose low half is recovered in full. The least favourable key of the whole sample was still recovered to 48 low bits, a remaining search of 2^{80}, so every one of the 5000 keys is broken below the designers' 124-bit claim, the worst by a margin of 2^{44}. The commands that reproduce the run are in the repository (Section 7).

6.4 Summary

Table 2 collects the variants. In the chosen-tweak model the 2^{127} class is half of all keys and the 2^{108} class is the one of Section 4. On $\mathcal{C}_{127}$ the twin pair followed by an exhaustive search over the remaining 2^{94} candidates recovers the whole key in time about 2^{94}, below the 124-bit security claimed by the designers [BLMR19]. The library procedure lowers this search key by key, to about 2^{67} on average, and by Proposition 6 it cannot go below 2^{64}.

Table 2: Attacks on full 32-round CRAFT from cell-separable invariants, with random known plaintexts. ST and CT are the single-tweak and chosen-tweak models. Time is in evaluations of the state functions, each below one encryption; memory is negligible, about 10^2 bytes per sequence for the library. Success probability is one for keys in the class, with false positives at most 2^{-22} per sequence ($N = 40$). "Relations" counts the independent linear equations on K known after the attack, class membership included, and "remaining" is the size of the search for the rest of the key.

variant	model	weak keys	data, time	relations	remaining
distinguisher, 2^{127}	CT	2^{127}	2^{18}	18	2^{110}
twin pair (Thm. 4)	CT	2^{127}	$2^{18.6}$	34	2^{94}
twin pair, $K \notin \mathcal{C}_{127}$	CT	2^{127}	2^{18}	1	2^{127}
library (Section 6.3)	CT	almost all	$2^{14.5}$	61.3 avg (≤ 64)	2^{67} avg ($\geq 2^{64}$)

7 Reproduction

All programs and data are available at https://github.com/robcivino/CraftNLinv-full. The program implementing the attack `craft_invattack.cpp` contains the cipher with the designers' test vector, the free-triple table with the self-test of Proposition 1, the affine system of a sequence for both the $\mathcal{C}_{108}$ and the $\mathcal{C}_{127}$ invariants and for every sequence of the library, the explicit weak-tweak construction, the pattern test and the library procedure of Algorithm 1, which reports the recovered low-key relations and the size of the exhaustive search that then finishes the key. The auxiliary program `classify.c` is the classification program that produced the three library files `p1.txt`, `p2.txt` and `p4_120.txt`, included here. Every result reported for the library in Section 6.3 is the output of `craft_invattack` on these files.

References

[BBI+15] Subhadeep Banik, Andrey Bogdanov, Takanori Isobe, Kyoji Shibutani, Harunaga Hiwatari, Toru Akishita, and Francesco Regazzoni. Midori: A block cipher for low energy. In *Advances in Cryptology — ASIACRYPT 2015*, volume 9453 of *Lecture Notes in Computer Science*, pages 411–436. Springer, 2015. doi:10.1007/978-3-662-48800-3_17.

[BCL18] Christof Beierle, Anne Canteaut, and Gregor Leander. Nonlinear approximations in cryptanalysis revisited. *IACR Transactions on Symmetric Cryptology*, 2018(4):80–101, 2018. doi:10.13154/tosc.v2018.i4.80-101.

[BCLR17] Christof Beierle, Anne Canteaut, Gregor Leander, and Yann Rotella. Proving resistance against invariant attacks: How to choose the round constants. In *Advances in Cryptology — CRYPTO 2017*, volume 10402 of *Lecture Notes in Computer Science*, pages 647–678. Springer, 2017. doi:10.1007/978-3-319-63715-0_22.

[Bey18] Tim Beyne. Block cipher invariants as eigenvectors of correlation matrices. In *Advances in Cryptology — ASIACRYPT 2018*, volume 11272 of *Lecture Notes in Computer Science*, pages 3–31. Springer, 2018. doi:10.1007/978-3-030-03326-2_1.

[Bey21] Tim Beyne. A geometric approach to linear cryptanalysis. In *Advances in Cryptology — ASIACRYPT 2021*, volume 13090 of *Lecture Notes in Computer Science*, pages 36–66. Springer, 2021. doi:10.1007/978-3-030-92062-3_2.

[BLMR19] Christof Beierle, Gregor Leander, Amir Moradi, and Shahram Rasoolzadeh. CRAFT: Lightweight tweakable block cipher with efficient protection against DFA attacks. *IACR Transactions on Symmetric Cryptology*, 2019(1):5–45, 2019. doi:10.46586/tosc.v2019.i1.5-45.

[CBS19] Roberto Civino, Céline Blondeau, and Massimiliano Sala. Differential attacks: using alternative operations. *Designs, Codes and Cryptography*, 87(2–3):225–247, 2019. doi:10.1007/s10623-018-0516-z.

[CCI24] Marco Calderini, Roberto Civino, and Riccardo Invernizzi. Differential experiments using parallel alternative operations. *Journal of Mathematical Cryptology*, 18(1):20230030, 2024. doi:10.1515/jmc-2023-0030.

[CF25] Roberto Civino and Valerio Fedele. Binary bibraces and applications to cryptography. *Mediterranean Journal of Mathematics*, 22(1):29, 2025. doi:10.1007/s00009-024-02793-z.

[CF26] Roberto Civino and Valerio Fedele. Classification of small binary bibraces via bilinear maps. *Results in Mathematics*, 81(2):53, 2026. doi:10.1007/s00025-026-02612-9.

[Civ26] Roberto Civino. Linear cryptanalysis with bibrace characters: Weak-key distinguishers for CRAFT. Cryptology ePrint Archive, Paper 2026/1767, 2026.

[EY19] Muhammad ElSheikh and Amr M. Youssef. Related-key differential cryptanalysis of full round CRAFT. In *Security, Privacy, and Applied Cryptography Engineering (SPACE 2019)*, volume 11947 of *LNCS*, pages 50–66. Springer, 2019. doi:10.1007/978-3-030-35869-3_6.

[GSS+20] Hao Guo, Siwei Sun, Danping Shi, Ling Sun, Yao Sun, Lei Hu, and Meiqin Wang. Differential attacks on CRAFT exploiting the involutory s-boxes and tweak additions. *IACR Transactions on Symmetric Cryptology*, 2020(3):119–151, 2020. doi:10.13154/tosc.v2020.i3.119-151.

[HBS21] Hosein Hadipour, Nasour Bagheri, and Ling Song. Improved rectangle attacks on SKINNY and CRAFT. *IACR Transactions on Symmetric Cryptology*, 2021(2):140–198, 2021. doi:10.46586/tosc.v2021.i2.140-198.

[HSN+19] Hosein Hadipour, Sadegh Sadeghi, Majid M. Niknam, Ling Song, and Nasour Bagheri. Comprehensive security analysis of CRAFT. *IACR Transactions on Symmetric Cryptology*, 2019(4):290–317, 2019. doi:10.13154/tosc.v2019.i4.290-317.

[KR96] Lars R. Knudsen and Matthew J. B. Robshaw. Non-linear approximations in linear cryptanalysis. In *Advances in Cryptology — EUROCRYPT 1996*, volume 1070 of *Lecture Notes in Computer Science*, pages 224–236. Springer, 1996. doi:10.1007/3-540-68339-9_20.

[LN26] Eran Lambooij and Patrick Neumann. A known-plaintext attack with minimal data complexity on 25-round CRAFT. *IACR Transactions on Symmetric Cryptology*, 2026(1):212–225, 2026. doi:10.46586/tosc.v2026.i1.212-225.

[LR22] Gregor Leander and Shahram Rasoolzadeh. Weak tweak-keys for the CRAFT block cipher. *IACR Transactions on Symmetric Cryptology*, 2022(1):38–63, 2022. doi:10.46586/tosc.v2022.i1.38-63.

[MA20] AmirHossein Ebrahimi Moghaddam and Zahra Ahmadian. New automatic search method for truncated-differential characteristics: application to Midori, SKINNY and CRAFT. *The Computer Journal*, 63(12):1813–1825, 2020. doi:10.1093/comjnl/bxaa004.

[MNPN25] Dounia M'Foukh, María Naya-Plasencia, and Patrick Neumann. The state-test technique on differential attacks: a 26-round attack on CRAFT and other applications. In *Advances in Cryptology — ASIACRYPT 2025*, Lecture Notes in Computer Science. Springer, 2025. Cryptology ePrint Archive, Report 2025/1597. doi:10.1007/978-981-95-5018-0_9.

[TLS16] Yosuke Todo, Gregor Leander, and Yu Sasaki. Nonlinear invariant attack: Practical attack on full SCREAM, iSCREAM and Midori64. In *Advances in Cryptology — ASIACRYPT 2016*, volume 10032 of *Lecture Notes in Computer Science*, pages 3–33. Springer, 2016. doi:10.1007/978-3-662-53890-6_1.

A Deferred proofs

The statements proved here are finite. Each can be checked exhaustively by computer, over the sixteen inputs of the S-box or the finite affine systems of a sequence, and the delivered program does so; we give the hand proofs as well. The statements themselves are in the body.

Proof of Proposition 1. (i) From (3): $D_k(\mathtt{5}\cdot \mathrm{Sb}_0) = k_0 \oplus k_1 \oplus k_3 \oplus k_0k_2 \oplus k_2x_0 \oplus k_0x_2$, so $m_k(\mathtt{5})$ has bit 0 equal to k_2 and bit 2 equal to k_0. For $\mathtt{2}\cdot \mathrm{Sb}_0 = S_1$ the derivative has linear part $x_0(k_2 \oplus k_3) \oplus x_2(k_0 \oplus k_3) \oplus x_3(k_0 \oplus k_2)$; it equals $\langle \mathtt{5}, x\rangle$ iff $k_0 = k_2$ and $k_3 = 1 \oplus k_0$, i.e. $k \in \{\mathtt{5},\mathtt{7},\mathtt{8},\mathtt{a}\}$. For $\mathtt{7}\cdot \mathrm{Sb}_0$ the linear part is $k_3(x_0 \oplus x_2) \oplus (k_0 \oplus k_2)x_3$. The radicals are read off the quadratic parts. (ii) The quadratic part of D_kC_0 is $x_0x_1(k_2 \oplus k_3) \oplus x_1x_2(k_0 \oplus k_3) \oplus x_1x_3(k_0 \oplus k_2) \oplus k_1(x_0x_2 \oplus x_0x_3 \oplus x_2x_3)$, which vanishes iff $k_1 = 0$ and $k_0 = k_2 = k_3$, i.e. $k \in \{\mathtt{0},\mathtt{d}\}$. That of D_kC_3 is $(k_0 \oplus k_2)x_1x_3 \oplus (x_0 \oplus x_2)(k_1x_3 \oplus k_3x_1)$, vanishing iff $k_0 = k_2$, $k_1 = k_3 = 0$, i.e. $k \in \{\mathtt{0},\mathtt{5}\}$. That of $D_k(C_0 \oplus C_3) = D_k(x_0x_1x_2)$ is $k_0x_1x_2 \oplus k_1x_0x_2 \oplus k_2x_0x_1$, vanishing iff $k \in \{\mathtt{0},\mathtt{8}\}$. Since the quadratic part of $D_k(a\cdot \mathrm{Sb}_0)$ is that of D_k of its cubic part, this proves the linear structures, and Lemma 2(i),(iv) give the shape of $\mathcal{K}(a,m,a)$. The explicit values: $D_{\mathtt{d}}S_2$ has linear part $x_0 \oplus x_1 \oplus x_3$ (mask $\mathtt{b}$), $D_{\mathtt{d}}S_0$ has $x_1 \oplus x_2 \oplus x_3$ (mask $\mathtt{e}$), $\mathtt{b}\cdot \mathrm{Sb}_0 = 1 \oplus x_0 \oplus x_1 \oplus x_2 \oplus x_0x_1 \oplus x_1x_3 \oplus x_0x_1x_2$ and $\mathtt{e}\cdot \mathrm{Sb}_0 = x_2 \oplus x_3 \oplus x_0x_1 \oplus x_0x_2 \oplus x_1x_3 \oplus x_0x_1x_2$ both have $D_{\mathtt{8}} = x_1 \oplus \mathrm{const}$ (mask $\mathtt{2}$), and $D_{\mathtt{5}}S_3 = x_1 \oplus x_3 \oplus \mathrm{const}$ (mask $\mathtt{a}$), all by direct expansion of (2). □

Proof of Theorem 1. (i) Since $\alpha_{t-1} = \alpha_t = \alpha_{t+1} = \mathtt{5}$ on every high cell, the triple of Lemma 1 at cell j, $i = \mathcal{P}(j)$, is $(\mathtt{5}, m_i(\mathtt{5}\text{-vector}), \mathtt{5})$ for every t. By the definition of m in Definition 1, $m_i = \mathtt{5}$ for $i = 8,\dots,11$ (row 2, where $m_{8+c} = \alpha[c] = \mathtt{5}$) and $m_i = \mathtt{0}$ for $i = 12,\dots,15$ (row 3, where $m_{12+c} = \alpha[c] \oplus \alpha[4+c] = \mathtt{5} \oplus \mathtt{5} = \mathtt{0}$). The two triples $(\mathtt{5},\mathtt{5},\mathtt{5})$ and $(\mathtt{5},\mathtt{0},\mathtt{5})$ are free exactly for $\mathcal{K}(\mathtt{5},\mathtt{5},\mathtt{5}) = \{\mathtt{5},\mathtt{7},\mathtt{d},\mathtt{f}\}$ and $\mathcal{K}(\mathtt{5},\mathtt{0},\mathtt{5}) = \mathrm{Ann}$, by Proposition 1(i). (ii) is a bookkeeping exercise on these two cosets, which we have reproduced exactly with `craft_invattack conditions 55555555 55555555 55555555 55555555`: this reports rank 48 (28 tweak-row pivots, 20 pure key relations), 2^{108} weak keys and 2^{144} weak pairs, and lists the twenty relations, which are precisely the eight cross relations $K_0[i] \oplus K_1[i] \in \mathrm{Ann}$, $i \in \{8,\dots,15\} \setminus \{9\} \cup \{9\}$ (read: for every i from 8 to 15, with $i = 11$ paired against $K_1[11]$ via $K_0[8] \oplus K_1[11]$ and $i = 13$ against $K_1[15]$), together with the two within-K_1 relations above. For (iii), iterate Lemma 1 over the thirty-one full rounds, read I off the state entering the first MixColumn, and treat the last round, which is $\mathrm{MC} \oplus rc \oplus tk$ only, exactly as in the proof of Theorem 3(iii) below, using $\mathcal{K}(\mathtt{5},\mathtt{5},\mathtt{5})$ and $\mathcal{K}(\mathtt{5},\mathtt{0},\mathtt{5})$ on the ciphertext side; because the sequence has period one, every tweakey- or

key-dependent contribution to the per-round constant occurs in all four tweakeys and the row constants are the only genuinely new term, so μ depends on neither K nor T. □

Proof of Theorem 2. By Lemma 1 the condition at round t and high cell j, with $i = \mathcal{P}(j)$, is

$$tk_i \in \mathcal{K}(\alpha_{t-1}[j],\, m_i(\alpha_t),\, \alpha_{t+1}[j]),$$

and $\alpha_{t-1} = \alpha_{t+1}$ for a period-two sequence. The induced low masks are $m(A) = (0,5,b,0,0,0,0,0)$ on cells $8,\dots,15$, since $A[c] = A[4+c]$ for every column, and $m(B) = (0,0,0,0,4,2,0,0)$. For even t the triples are $(B[j], m_{\mathcal{P}(j)}(A), B[j])$: they are $(0,0,0)$ except at $j = 4$ ($\mathcal{P}(4) = 10$), where the triple is $(4,b,4)$, and $j = 5$ ($\mathcal{P}(5) = 9$), where it is $(2,5,2)$. For odd t the triples are $(A[j], m_{\mathcal{P}(j)}(B), A[j])$: trivial except at $j = 1$ ($\mathcal{P}(1) = 12$): $(5,4,5)$; $j = 2$ (13): $(b,2,b)$; $j = 5$ (9): $(5,0,5)$; $j = 6$ (8): $(b,0,b)$. The sets $\mathcal{K}$ of these six triples are those of Proposition 1. The displayed forms of I_{ev}, I_{od} follow from Definition 1 with $\mathcal{P}(4) = 10$, $\mathcal{P}(5) = 9$, $\mathcal{P}(1) = 12$, $\mathcal{P}(2) = 13$, $\mathcal{P}(6) = 8$. □

Proof of Theorem 3. The conditions of Theorem 2 at the four tweakeys are twelve statements. Each fixes a nibble of $TK_t[i] = K_{t \bmod 2}[i] \oplus T[\sigma_t(i)]$, $\sigma_t = \mathrm{id}$ for $t < 2$ and $\sigma_t = \mathcal{Q}$ for $t \geq 2$, modulo the subspace of its coset; the relevant values are $\mathcal{Q}(10) = 7$, $\mathcal{Q}(9) = 3$, $\mathcal{Q}(8) = 11$, $\mathcal{Q}(12) = 6$, $\mathcal{Q}(13) = 0$. The six singleton conditions ($tk_{10} = d$ at TK_0, TK_2; $tk_8 = 0$ and $tk_{13} = 8$ at TK_1, TK_3) determine $T[10], T[7], T[8], T[11], T[13], T[0]$ completely, 24 affine equations in all. The two conditions on cell 12 fix $T[12]$ and $T[6]$ modulo Ann, 4 equations. The tweak cells 9 and 3 each carry two conditions: $K_0[9] \oplus T[9] \in \{5,7,8,a\}$, which reads $\langle 5, K_0[9] \oplus T[9]\rangle = 0$ and $\langle 9, K_0[9] \oplus T[9]\rangle = 1$, and $K_1[9] \oplus T[9] \in \mathrm{Ann}$, which reads $\langle 1, K_1[9] \oplus T[9]\rangle = \langle 5, K_1[9] \oplus T[9]\rangle = 0$. The three functionals $\langle 5,\cdot\rangle, \langle 9,\cdot\rangle, \langle 1,\cdot\rangle$ are independent, so these four equations fix three bits of $T[9]$ and leave exactly one relation free of T, namely $\langle 5, K_0[9]\rangle \oplus \langle 5, K_1[9]\rangle = 0$; the same at $T[3]$, giving the same relation. Altogether 36 equations of rank 35 in the 192 bits of (K_0, K_1, T), with a single pure-key relation: this proves (ii), the count $2^{192-35} = 2^{157}$, and that the key class is $\mathcal{C}_{127}$, of codimension one. The explicit $T(K)$ satisfies every condition: for cell 9, $K_1[9] \oplus T[9] \in \{0,8\} \subseteq \mathrm{Ann}$, and $K_0[9] \oplus T[9] = \delta \oplus 8(1 \oplus \langle 9, \delta\rangle)$; since $\langle 5, \delta\rangle = 0$, $\delta \in \{0,2,d,f\} \cup \{5,7,8,a\}$, the second set being exactly where $\langle 9, \delta\rangle = 1$, and adding 8 exchanges the two sets. The free part W has dimension $6 \cdot 4 + 2 \cdot 2 + 2 \cdot 1 = 30$, which proves (i). For (iii), iterate Lemma 1 over rounds $0,\dots,30$ and treat the last round as follows: with x the state entering the last MixColumn, $C = \mathrm{MC}(x) \oplus rc \oplus TK_3$, so $x_i = C_i \oplus TK_3[i]$ on low cells and $x_4 = C_4 \oplus C_{12} \oplus \mathrm{const}$, $x_5 = C_5 \oplus C_{13} \oplus \mathrm{const}$. Using the four free triples of the odd rounds at TK_3, $\langle 5, \mathrm{Sb}_0(C_{12} \oplus TK_3[12])\rangle = \langle 5, \mathrm{Sb}_0(C_{12})\rangle \oplus \langle 4, C_{12}\rangle \oplus \mathrm{const}$, $\langle b, \mathrm{Sb}_0(C_{13} \oplus 8)\rangle = \langle b, \mathrm{Sb}_0(C_{13})\rangle \oplus \langle 2, C_{13}\rangle \oplus \mathrm{const}$, $\langle 5, \mathrm{Sb}_0(C_9 \oplus TK_3[9])\rangle = \langle 5, \mathrm{Sb}_0(C_9)\rangle \oplus \mathrm{const}$ and $\langle b, \mathrm{Sb}_0(C_8 \oplus 0)\rangle = \langle b, \mathrm{Sb}_0(C_8)\rangle$, so that $I_{31}(x) = I_{\mathrm{od}}(C) \oplus \mathrm{const}$, the terms $\langle 4, C_{12}\rangle$ and $\langle 2, C_{13}\rangle$ cancelling against those coming from x_4, x_5. The total constant is $\bigoplus_{t=0}^{31} \kappa_t$ with κ_t as in Lemma 1. Each of $TK_0,\dots,TK_3$ occurs in exactly eight of the thirty-two rounds, and the sequence has period two, so every key- or tweak-dependent term of κ_t occurs an even number of times and cancels; what remains is the contribution of the round constants on cells 4 and 5, with $\alpha_t[4] = 0,4$ and $\alpha_t[5] = 5,2$ for t even, odd. □

Proof of Proposition 3. At even rounds the triple at j is $(B[j], m_{\mathcal{P}(j)}(A), B[j])$. If $B[j] = 0$ it is free only if $m_{\mathcal{P}(j)}(A) = 0$. If $B[j] \neq 0$ is quadratic, Lemma 2(iii) requires $m_{\mathcal{P}(j)}(A) \in \{m_k(B[j])\}_k$, a set that by Proposition 1(i) meets $\{0,2,5,7\}$ in $\{0,5\}$ for each of $2,5,7$. Since $m_{\mathcal{P}(j)}(A)$ is a sum of entries of A, hence in $\{0,2,5,7\}$, we get $m(A) \in \{0,5\}^{\mathcal{L}}$, i.e. $A[c] \in \{0,5\}$ (row 2) and $A[c] \oplus A[4+c] \in \{0,5\}$ (row 3), whence $A \in \{0,5\}^{\mathcal{H}}$; symmetrically B. With all entries in $\{0,5\}$ the two conditions become $\operatorname{supp} m(A) \subseteq \mathcal{P}(\operatorname{supp} B)$ and $\operatorname{supp} m(B) \subseteq \mathcal{P}(\operatorname{supp} A)$, which is the statement with μ. The explicit form of μ is read from $m_{8+c} = H[c]$, $m_{12+c} = H[c] \oplus H[4+c]$ and the pairs of $\mathcal{P}$: row-2 cell $8+c$ goes to the row-1 cell of column $\pi(c)$, row-3 cell $12+c$ to the row-0 cell of column $\sigma(c)$. The

conditions and the class follow from Proposition 1(i) and the tweakey schedule as in the proof of Theorem 3: a cell constrained at both parities shares the tweak cell $T[i]$ (and $T[\mathcal{Q}(i)]$) between $K_0[i]$ and $K_1[i]$, giving a two-bit relation modulo Ann; the tweak cell $T[11]$ is added to cell 11 of TK_0, TK_1 and to cell 8 of TK_2, TK_3, and $T[13]$ to cell 13 of TK_0, TK_1 and to cell 15 of TK_2, TK_3, which couples the two cells of each pair when both are constrained. □