Circulant Diffusion Layers: Cost Bounds and Rectangle Reuse

Enyan Li Gaoli Wang

Shanghai Key Laboratory of Trustworthy Computing,
Software Engineering Institute, East China Normal University,
Shanghai 200062, China

lienyan6@gmail.com glwang@sei.ecnu.edu.cn

Abstract

We study the addition complexity of sparse invertible circulant matrices over $\mathbb{F}_2$ and odd prime fields. Diffusion is measured by branch numbers, and prime field coefficients are restricted to $\{0, \pm 1\}$. Finding optimal general circuits is often difficult, particularly at large matrix orders. Rectangle reuse offers a simpler alternative for explicit circuit construction. We compare the general circuit cost C_g with the rectangle reuse cost C_{rect}, and show that rectangle reuse can approach, and sometimes attain, the general circuit optimum while providing shallow implementations in practical cryptographic parameter settings.

To the best of our knowledge, we establish the first lower bounds on both C_g and C_{rect} for general matrices in terms of matrix order and branch numbers. We then exploit cyclic structure to obtain cost classifications, sharper bounds, and explicit constructions at minimum weight. Over $\mathbb{F}_2$, we give a complete classification of rectangle costs over all invertible weight-three circulants and determine their exact minima at every order $n \geq 4$. The minimum costs in both the general circuit and rectangle reuse models equal $3n/2$ at every even order, with connected constructions attaining these minima. Over every odd prime field, connected weight-three constructions likewise attain both minima, namely $\lceil 3n/2 \rceil$, for every $n \geq 4$. For branch number five, we prove the rectangle lower bound $2n$ and give connected families attaining it at infinitely many orders. For branch number six, we exploit cyclic structure to sharpen the general rectangle bounds and determine certified exact rectangle minima at selected orders. Finally, we characterize the effects of connectivity and circuit depth. For weight-three circulants over $\mathbb{F}_2$ or $\mathbb{F}_p$ with $p \geq 5$, the minimum gate count at depth two equals the rectangle cost. These results provide theoretical cost benchmarks and explicit constructions for designing circulant diffusion layers with low gate counts and shallow circuits.

Keywords: Lightweight cryptography; diffusion layers; circulant matrices; branch number; addition complexity.

1 Introduction

1.1 Background and Motivation

Symmetric cryptographic primitives support applications ranging from embedded devices and the Internet of Things to blockchain protocols and privacy-preserving computation. These applications impose different requirements on security, implementation area, latency, and arithmetic complexity. Appropriately parameterized symmetric primitives are also expected to remain essential in systems designed for the post-quantum era [16]. Established designs such as AES, SHA-3, and Ascon rely on binary operations and efficient Boolean implementations [17, 18, 21]. Their linear layers are expected to provide strong diffusion

with low implementation area and latency, commonly assessed through XOR gate count and circuit depth. Sparse matrices are therefore attractive candidates, provided that they retain the required diffusion properties. Alongside these designs, arithmetization-oriented (AO) primitives have been developed for applications such as zero-knowledge (ZK) proofs and fully homomorphic encryption (FHE), with representative examples including Poseidon and HERA [1, 8]. For these primitives, desirable linear layers combine suitable diffusion with compact arithmetic circuits requiring few field additions and constant multiplications. Efficient linear layers reduce evaluation and witness generation costs in ZK proofs. In FHE evaluation, inexpensive coefficients help control computational cost and noise growth [1, 9]. Thus, field coefficients with inexpensive multiplication are desirable in FHE. Research on linear layer design for such symmetric primitives over large fields remains limited. Designing linear layers with low cost and latency while meeting prescribed diffusion requirements is therefore relevant to both binary and prime field settings.

The design of linear diffusion layers has two complementary directions: constructing matrices that combine suitable cryptographic properties with efficient implementations, and optimizing the circuits of matrices that have already been specified. Both directions are represented in the literature on lightweight diffusion and automatic circuit search [2, 6, 25]. Branch numbers provide a standard measure of diffusion and support bounds on the activity of differential and linear trails [3]. However, invertibility and a prescribed branch number do not directly determine the minimum number of operations needed to evaluate a matrix. A sparse representation gives a simple implementation bound, but shared intermediate values and cancellations can substantially change the cost. Existing circuit optimizations, lower bounds for particular implementation models, and certified results for small dimensions provide important foundations [2, 6, 22]. They motivate a complementary structural question: for a specified matrix class and diffusion requirement, what cost is unavoidable, and which matrices attain it? We focus on the construction direction through this question, seeking bounds and equality conditions that guide matrix selection across dimensions.

Circulant matrices form a natural class for this investigation. Their cyclic symmetry gives compact descriptions. The repeated patterns of input dependence can be exploited in implementations. They also occur at different granularities in established designs: AES uses a circulant MixColumns matrix over $\mathbb{F}_{2^8}$, whereas Ascon applies sparse binary circulant maps within its state words [18, 21]. Previous work has investigated the cryptographic properties and efficient implementations of cyclic binary mappings [4, 15]. These features motivate the study of sparse circulant matrices over $\mathbb{F}_2$ and odd prime fields $\mathbb{F}_p$. Their suitability as linear layers depends on both diffusion and implementation efficiency. Branch numbers measure diffusion, while gate count and circuit depth describe key aspects of implementation efficiency. However, determining the general circuit cost C_g, corresponding to the g-XOR cost in the binary setting [2], is often difficult, particularly at large matrix orders. Inspired by classical rectangular common subexpression extraction [11], we therefore focus on a special form of reuse called rectangle reuse, with cost C_{rect}. It extracts shared linear forms from rank one submatrices whose rows are proportional. This restriction provides a tractable setting for analyzing gate count and circuit depth together. Rectangle reuse facilitates explicit circuit construction across matrix orders but may require more gates than an optimal general circuit. This raises a theoretical question: how closely can this simpler reuse strategy approach the general circuit optimum across different matrix orders and branch number requirements? Since invertibility and connectivity are often basic requirements in cryptographic design, we keep invertibility explicit and emphasize connected constructions. Our goal is to relate diffusion requirements and cyclic structure to provable cost bounds, characterize optimal patterns, and obtain explicit constructions.

1.2 Related Work

The computational complexity of linear transformations has long been studied through circuit complexity and coding theory [7, 12, 19]. These studies connect matrix properties with computation under different restrictions on depth, allowed gates, and size measures. In symmetric cryptography, closely related questions arise in the construction and implementation of diffusion matrices. Duval and Leurent search directly for lightweight circuits yielding MDS matrices, while Xiang et al. optimize implementations through matrix decompositions [6, 25]. Recent automatic methods also jointly consider cryptographic properties and implementation constraints. Cui et al. combine branch number, invertibility, XOR count, and circuit depth in an STP-based model, and obtain tight gate count bounds for dimensions 4, 6, and 8 with branch number four; Weng and Wang develop a SAT-based approach to constructing efficient linear layers [2, 23]. Implementation lower bounds have also been studied for block MDS matrices under restricted circuit models [22]. Existing work on linear layer optimization primarily addresses prescribed matrices, selected matrix families, or fixed parameter ranges.

Cyclic binary diffusion matrices have been studied directly by Dehnavi et al. and Rishakani et al. [4, 15]. Rishakani et al. [15] study the branch numbers of cyclic binary matrices and construct efficient diffusion layers by composing sparse cyclic mappings. Of particular relevance, their Lemma 4.1 gives a shared pair implementation for weight-three mappings whose rotation offsets form an arithmetic progression. Guo et al. investigate cyclic structure in rotational-XOR diffusion primitives, with diffusion measured at the word level [10]. Our study builds on these results and examines cost bounds and optimality across dimensions, including signed constructions over odd prime fields and the effect of connectivity requirements.

Rectangular submatrix extraction also appears in arithmetic optimization. Hosangadi et al. use rectangles in a kernel intersection matrix to extract common subexpressions, with a local addition saving of $(r-1)(s-1)$ for an $r \times s$ rectangle [11]. Their procedure introduces intermediate variables that can be reused in subsequent extraction steps. We specify admissible rectangles and their disjointness directly in the original matrix to define C_{rect}, and compare this restricted reuse model with the general circuit cost C_g. Our analysis combines these cost models with branch number constraints and circulant structure to derive cost lower bounds and explicit constructions.

1.3 Contributions

Circulant matrices are widely used as linear layers in block ciphers and hash functions. We study these matrices and explore the tradeoffs between diffusion and implementation cost. We focus on sparse invertible matrices with prescribed branch numbers and measure cost by the number of two-input additions or subtractions. Our contributions are as follows.

1. **General cost lower bounds.** To the best of our knowledge, we establish the first theoretical lower bounds on general circuit addition complexity under prescribed branch number requirements for arbitrary matrix dimensions. These bounds do not require circulant structure. For rectangle reuse, we combine diffusion constraints, invertibility, and row and column capacities to derive explicit cost bounds at minimum weight. These bounds are in terms of the order and branch number requirement. This framework provides theoretical benchmarks for matrix selection and circuit optimization. We further examine how connectivity and circuit depth affect attainable costs in the circulant classes below.

2. **Binary circulants for lightweight diffusion.** Motivated by the demand for sparse diffusion layers in binary field block ciphers and hash functions, we study circulants with branch numbers four and six at minimum weight. At branch number four, we first give a complete classification of rectangle costs over all invertible binary circulants of weight three and determine their exact minima at every feasible order. We also give connected families that attain the general circuit lower bound. We find that rectangle reuse is optimal among circuits of depth at most two, and any implementation using fewer gates must have greater depth. At branch number six, we strengthen the rectangle lower bound using circulant support structure and determine exact rectangle minima at selected orders. The constructions and finite optimizations provide concrete choices for balancing XOR count, diffusion, depth, and connectivity.

3. **Signed circulants over odd prime fields.** Motivated by AO block ciphers and hash functions designed over large prime fields, we study signed circulants with branch numbers four, five, and six at minimum weight. At branch number four, we completely classify rectangle costs according to support structure and coefficient signs. Connected constructions attain both cost minima over every odd prime field and at every admissible order. At branch number five, we establish the rectangle lower bound and construct connected families attaining it at infinitely many orders. We establish a stronger rectangle bound at branch number six under the stated connectivity and order assumptions. Finite optimizations at representative orders complement these families and bounds. We also characterize the depth-two optimality of rectangle reuse in the relevant classes and examine the effects of connectivity and field characteristic.

Across many of the practical orders, fields, and branch number requirements, rectangle reuse achieves gate counts close to, and in some cases equal to, the general circuit optimum or its established lower bound. It also provides shallow implementations, while further gate savings through general reuse may require greater depth. Rectangle reuse therefore offers a useful balance between gate count and circuit depth, making the circuits studied here valuable candidates for linear layers in practical symmetric cipher design.

Table 1 summarizes these contributions and points to the theorems and propositions giving explicit lower bound formulas.

1.4 Organization

Section 2 specifies the scope of the study and introduces the basic concepts, cost models, and general lower bounds. The next two sections study cost lower bounds and constructions for individual parameters under common branch number requirements. Section 3 considers circulants over $\mathbb{F}_2$ with branch numbers four and six. Section 4 considers signed circulants over odd prime fields $\mathbb{F}_p$ with branch numbers four, five, and six. Section 5 describes the search and optimization methods and the associated code. Section 6 concludes the paper.

2 Preliminaries, Cost Models, and General Bounds

We study the addition cost of circulant linear layers. Our main goal is to derive cost lower bounds under prescribed invertibility and diffusion requirements. We require the matrices to have full rank so that the transformations are invertible. They must also satisfy specified lower bounds on their branch numbers to ensure a prescribed level of diffusion. Our cost measure counts additions in the linear transformation: XORs over $\mathbb{F}_2$ and field additions or subtractions over odd prime fields. Actual implementation performance also depends on circuit topology, instruction selection, and interaction with S-boxes and other cipher

Table 1: Summary of cost lower bounds, matrix constructions, finite optimizations, connectivity, and circuit depth. The first two rows do not require circulant structure. In the remaining rows, the circulant structure is required and the weight is $b-1$. Here b is the branch number. The parameter p is an odd prime and coefficients over $\mathbb{F}_p$ are signed.

Setting	Lower bounds	Attaining constructions / finite optimization	Connectivity	Depth: rectangle reuse	Depth: general circuits
General C_{rect} Section 2.3.2	Theorem 2.9.	—	Unrestricted.	—	—
General C_g Section 2.3.1	Theorem 2.4; Proposition 2.5.	—	Unrestricted.	—	—
$\mathbb{F}_2$, $b=4$ Section 3.1	Theorem 3.4; Proposition 3.5.	Theorem 3.4; Table 3.	Theorem: connected constructions at some orders; unrestricted at others. Table reports both classes.	Two.	For fixed matrices, improving on rectangle cost requires depth at least three.
$\mathbb{F}_2$, $b=6$ Section 3.2	Theorem 3.12.	Table 4.	Finite minima coincide with and without connectivity.	At least three.	Depth at least three.
$\mathbb{F}_p$, $b=4$ Section 4.1	Theorem 4.4.	Theorem 4.4; Table 5.	Connected constructions attain both cost minima.	Two.	For fixed matrices, improving on rectangle cost requires depth at least three when $p \geq 5$.
$\mathbb{F}_p$, $b=5$ Section 4.2	Theorem 4.8.	Theorem 4.9; Proposition 4.10; Table 6.	The theorem gives connected families; nontrivial lifts are disconnected. Finite minima are reported without a connectivity requirement.	Two.	Depth-two minimum equals rectangle cost. Any improvement requires depth at least three.
$\mathbb{F}_p$, $b=6$ Section 4.3	Theorem 4.14.	Table 7.	Finite minima coincide with and without connectivity.	Three for the listed finite optima.	Depth at least three.

components. Circuit search and implementation optimization have been studied in, for example, [2, 26]. Optimizing circuits for existing linear layers is not the main objective of this work.

This section introduces the necessary concepts and notation, including branch numbers, connectivity, and the rules for counting additions. We specify the circulant constructions under consideration and the scope of our results. We then establish basic lemmas used in later sections and general lower bounds that do not depend on a particular circulant structure.

2.1 Basic concepts, notation, and scope

Linear layers and weight. Let $\mathbb{F}$ be a field. A linear layer of dimension n is an invertible matrix $M \in \mathbb{F}^{n\times n}$ acting on column vectors by $x \mapsto Mx$. Coordinates are indexed by $\mathbb{Z}_n = \mathbb{Z}/n\mathbb{Z}$, represented by $0, \dots, n-1$. For $x \in \mathbb{F}^n$, its support and Hamming weight are

$$\operatorname{supp}(x) = \{i \in \mathbb{Z}_n : x_i \neq 0\}, \qquad \operatorname{wt}(x) = |\operatorname{supp}(x)|.$$

The support of a matrix is the set of its nonzero positions, and $\operatorname{wt}(M)$ denotes its total number of nonzero entries. A matrix has uniform weight w if every row and every column contains w nonzero entries; in this case $\operatorname{wt}(M) = nw$. The weight of a circulant matrix is the number of nonzero entries in each row.

Differential and linear branch numbers. The differential branch number [3] is

$$B_d(M) = \min_{0\neq x\in\mathbb{F}^n} \big(\operatorname{wt}(x) + \operatorname{wt}(Mx)\big). \tag{2.1}$$

We use the transpose convention for the linear branch number:

$$B_l(M) = B_d(M^{\mathsf{T}}). \tag{2.2}$$

Indeed, if v is an output mask, its induced input mask is $M^{\mathsf{T}}v$, so the total mask weight is $\operatorname{wt}(v) + \operatorname{wt}(M^{\mathsf{T}}v)$. Invertibility also gives $B_d(M^{-1}) = B_d(M)$. For an invertible matrix, considering vectors of weight one gives

$$2 \le B_d(M) \le n+1, \qquad 2 \le B_l(M) \le n+1. \tag{2.3}$$

If both branch numbers are at least b, every row and every column must have weight at least $b-1$. A layer is *maximum distance separable* (MDS) if $B_d(M) = B_l(M) = n+1$. Equivalently, every square submatrix of M is invertible.

Circulants and the shift convention. Let $P \in \mathbb{F}^{n\times n}$ be the cyclic permutation matrix specified by

$$P = \begin{pmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & 0 & 1 \\ 1 & 0 & \cdots & 0 & 0 \end{pmatrix}.$$

For the input column vector $x = (x_0, \dots, x_{n-1})^{\mathsf{T}}$,

$$Px = (x_1, x_2, \dots, x_{n-1}, x_0)^{\mathsf{T}}.$$

Thus P shifts the input coordinates cyclically by one position to the left.

For $a = (a_0, \ldots, a_{n-1}) \in \mathbb{F}^n$, the circulant with first row a is

$$M = \mathrm{Circ}(a_0, \ldots, a_{n-1}) = \sum_{d=0}^{n-1} a_d P^d, \qquad M_{i,j} = a_{j-i}. \tag{2.4}$$

All subscripts in this formula are taken modulo n. Each row of M is obtained by shifting the preceding row cyclically by one position to the right. Its generating polynomial and coefficient support are

$$f(X) = \sum_{d=0}^{n-1} a_d X^d \in \mathbb{F}[X], \qquad D = \{d \in \mathbb{Z}_n : a_d \neq 0\}.$$

Thus $M = f(P)$, its weight is $w = |D|$, and the support of row i is $i + D$. For $d \in \mathbb{Z}_n \setminus \{0\}$, define the translation overlap

$$c(d) = |D \cap (D - d)| = |\{j \in \mathbb{Z}_n : a_j \neq 0,\ a_{j+d} \neq 0\}|. \tag{2.5}$$

This is the support intersection size of rows whose indices differ by d, and $c(d) = c(-d)$. Over a field of odd characteristic, when all nonzero coefficients belong to $\{1, -1\}$, we also use

$$c_\pm(d) = |\{j \in \mathbb{Z}_n : a_j a_{j+d} = \pm 1\}|, \qquad c(d) = c_+(d) + c_-(d). \tag{2.6}$$

These statistics distinguish support overlap from the sign relations that determine whether overlapping contributions can cancel.

Proposition 2.1. *For the circulant $M = f(P)$, the following hold:*

1. *$B_d(M) = B_l(M)$.*
2. *M is invertible if and only if $\gcd(f(X), X^n - 1) = 1$ in $\mathbb{F}[X]$.*

Neither assertion requires the characteristic of $\mathbb{F}$ to be coprime to n.

The proof is given in Appendix A.1.

For circulants we therefore abbreviate the common branch number as $B(M)$. For other matrix classes, the two branch numbers are retained unless their equality has been established.

Connectivity. The support graph of M is bipartite, with one vertex for each input and one for each output. An edge joins input j to output i when $M_{i,j} \neq 0$. We call the layer connected if this graph is connected.

Fields and the scope of our work. This work studies scalar circulant constructions with two coefficient classes: binary circulants over $\mathbb{F}_2$, and signed circulants over $\mathbb{F}_p$ with p an odd prime and entries in $\{0, 1, -1\}$. The binary case is motivated by linear layers in conventional block ciphers and hash functions. Examples include the binary representation of AES MixColumns [18] and the θ step in Keccak [17]. More specifically, SPEEDY-192 applies a binary circulant matrix of order 32 and weight seven to each state column [13]. Ascon applies a binary circulant matrix of order 64 and weight three within each of its five state words [21].

The prime field case targets AO symmetric primitives for FHE and ZK proofs. Examples include Pasta [5] and YuX [14] for FHE, and Poseidon [8] and Rescue-Prime [20] for ZK applications. These designs often operate over large prime fields and use diffusion matrices

with general field coefficients. The choice of coefficients affects the arithmetic cost of the linear layer and, in FHE, the growth of ciphertext noise [24]. We focus on coefficients in $\{0, 1, -1\}$ so that the linear layer can be evaluated using only field additions and subtractions. This choice is used in the diffusion layers evaluated homomorphically in YuX [14] and its variant YuX† [24].

We measure implementation cost by the number of additions, including subtractions. Setting aside other factors such as circuit topology, fewer additions give a lower implementation cost. Our design objective is to determine the minimum number of additions required by circulant constructions that are invertible and meet a prescribed branch number bound. Section 2.2 defines the cost models and counting rules.

Low-weight constraints and branch number targets. Sections 3 and 4 study practical low-weight linear layers with branch number requirements $B(M) \geq b$ for $b \in \{4, 5, 6\}$. These values represent the practical branch number targets of linear layers in most symmetric cryptographic designs. If we additionally require the minimum weight $w = b - 1$, then $B(M) = b$.

For binary circulants, invertibility forces w to be odd: otherwise $f(1) = 0$, contradicting Proposition 2.1. For odd w, the sum of the output coordinates equals the sum of the input coordinates in $\mathbb{F}_2$. Hence $\mathrm{wt}(Mx)$ and $\mathrm{wt}(x)$ have the same parity, and $B(M)$ is even. In particular, the binary targets $B \geq 5$ and $B \geq 6$ coincide, and their smallest possible circulant weight is five.

Over an odd prime field, the lower target $B = 3$ illustrates why size restrictions matter even at weight two. After a nonzero scalar factor and a support shift, a signed binomial has the form $I \pm P^d$ with $d \neq 0$. The matrix $I - P^d$ is singular, while $I + P^d$ is invertible exactly when $n/\gcd(n, d)$ is odd. In that case its branch number is three. Thus such a layer exists exactly when n has an odd divisor greater than one. Its support graph has $\gcd(n, d)$ components, so it is connected exactly when $\gcd(n, d) = 1$. A connected weight-two layer exists exactly for odd $n \geq 3$.

2.2 Addition Cost Models and Basic Tools

2.2.1 Addition cost models

Counting additions. We count operations in a circuit computing Mx from the input coordinates $x_0, \ldots, x_{n-1}$. Each gate adds or subtracts two available signals and has cost one. Its operands may be input coordinates or outputs of earlier gates. The same signal may be used as both operands. Copying, fanout, coordinate permutations, and negation are free. The zero signal is also available for free. Multiplication by any other constant must be implemented using the counted gates. Over $\mathbb{F}_2$, a gate is a two-input XOR. Over an odd prime field, it is a field addition or subtraction.

The depth of an input is zero. The depth of a gate is one plus the maximum depth of its operands, and the circuit depth is the maximum depth of an output. Unless a depth constraint is stated, the cost is minimized over circuits of arbitrary depth.

For a matrix with entries in $\{0, 1, -1\}$ and no zero rows, computing each output independently gives the direct cost

$$C_0(M) = \mathrm{wt}(M) - n. \tag{2.7}$$

If every row has weight w, this is $n(w - 1)$. We write $C_g(M)$ for the minimum number of gates in a circuit computing Mx. Intermediate signals may be reused throughout the circuit, and coefficients may cancel. In the binary case, this is the cost of general XOR (g-XOR) circuits [2].

The s-XOR metric counts in-place updates of the form $x_i \leftarrow x_i \oplus x_j$ [2]. It describes implementations using a fixed number of storage locations. We study addition costs without this storage restriction and use C_g as the general circuit cost. Lower bounds on C_g also apply to s-XOR implementations.

Rectangle reuse. Inspired by the rectangle-based common subexpression extraction of Hosangadi, Fallah, and Kastner [11], we consider a rectangle reuse model. It shares sums of original input coordinates among several outputs. This connects sharing decisions directly to matrix structure and often makes sharp bounds and exact optima easier to establish than for general circuits.

A rectangle is a set of matrix positions $Q = I \times J$, where I is a set of rows and J is a set of columns. Write $r_Q = |I|$ and $s_Q = |J|$, with $r_Q, s_Q \geq 2$. Over $\mathbb{F}_2$, the rectangle is admissible if all its entries are one. Over an odd prime field, it is admissible if

$$M_{i,j} = \alpha_i \beta_j \quad (i \in I,\ j \in J), \qquad \alpha_i, \beta_j \in \{1, -1\}. \tag{2.8}$$

Thus the entries form a nonzero matrix of rank one. Compute the shared sum $\sum_{j \in J} \beta_j x_j$ once and use it with sign α_i in output i for each $i \in I$. Let $\mathcal{R}(M)$ denote the set of all admissible rectangles in M. Figure 1 illustrates rectangle reuse over $\mathbb{F}_2$ in (a) and over an odd prime field in (b).

(a) Binary rectangle

$$\begin{pmatrix} \boxed{1} & \boxed{1} & 0 & \boxed{1} \\ \boxed{1} & \boxed{1} & 1 & \boxed{1} \\ 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix}$$

Shared sum: $x_0 + x_1 + x_3$.
Used in outputs 0 and 1.

(b) Signed rectangle

$$\begin{pmatrix} \boxed{1} & \boxed{-1} & \boxed{1} & 0 \\ \boxed{-1} & \boxed{1} & \boxed{-1} & 1 \\ 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \end{pmatrix}$$

Shared sum: $x_0 - x_1 + x_2$.
Outputs 0 and 1 use opposite signs.

Figure 1: Admissible rectangles with row set $I = \{0, 1\}$. The column sets are $J = \{0, 1, 3\}$ in (a) and $J = \{0, 1, 2\}$ in (b). The two outlined parts in (a) form a single rectangle with nonconsecutive columns.

A rectangle scheme selects a collection $\mathcal{S} \subseteq \mathcal{R}(M)$ whose rectangles have disjoint sets of matrix positions. Their row sets and column sets may overlap. Each shared sum is computed independently from the original inputs. Each output then combines its shared sums and remaining individual terms. We denote the minimum cost of such a scheme by $C_{\mathrm{rect}}(M)$.

Proposition 2.2. *For a matrix $M \in \{0, 1, -1\}^{n \times n}$ with no zero rows,*

$$C_{\mathrm{rect}}(M) = \mathrm{wt}(M) - n - \max_{\mathcal{S} \subseteq \mathcal{R}(M)} \sum_{Q \in \mathcal{S}} (r_Q - 1)(s_Q - 1), \tag{2.9}$$

where the rectangles in $\mathcal{S}$ have disjoint sets of matrix positions. The empty collection is allowed.

Proof. Sharing an $r_Q \times s_Q$ rectangle replaces s_Q individual terms by one signal in each of r_Q outputs. The output computations save $r_Q(s_Q - 1)$ gates. Computing the shared sum costs $s_Q - 1$ gates, so the net saving is $(r_Q - 1)(s_Q - 1)$. Disjoint matrix positions ensure that the savings add over the selected rectangles. Every scheme in the model has this form, which gives the stated minimum. □

In particular,

$$C_g(M) \le C_{\mathrm{rect}}(M) \le C_0(M). \tag{2.10}$$

Nested reuse can improve on a rectangle scheme. For example, consider the three output forms

$$x_0 + x_1 + x_2, \qquad x_0 + x_1 + x_2 + x_3, \qquad x_0 + x_1.$$

Independent computation uses six gates. A rectangle scheme uses four, by sharing either $x_0 + x_1$ among all three outputs or $x_0 + x_1 + x_2$ among the first two. General reuse computes

$$t_1 = x_0 + x_1, \qquad t_2 = t_1 + x_2, \qquad t_3 = t_2 + x_3$$

in three gates and takes (t_2, t_3, t_1) as the outputs. General circuits can also use cancellation, for example to obtain $x_0 + x_1$ from $t_2 - x_2$.

Bounds in the rectangle model let us identify matrices that are optimal for rectangle reuse. For several low-weight parameter families studied below, rectangle constructions also attain the lower bound for C_g. Other examples have only a small gap between the rectangle cost and the general circuit lower bound. Optimizing C_{rect} therefore provides a practical way to construct linear layers with low addition cost.

2.2.2 Construction and verification tools

The component structure of the support graph also affects addition cost: disconnected layers can be implemented component by component. Coordinate permutations expose this decomposition without changing the cost models.

Connectivity and equivalent constructions. For the admissible circulant constructions studied below, we seek connected families attaining the addition-count lower bounds wherever possible. Connectivity allows diffusion across the full set of coordinates, whereas a disconnected layer mixes coordinates only within its independent components.

For a circulant, translate the coefficient support D so that $0 \in D$. The number of connected components is then $\gcd(n, D)$, the greatest common divisor of n and the elements of D. In particular, the layer is connected if and only if $\gcd(n, D) = 1$. Indeed, a path of two edges between output vertices changes the index by an element of $D - D$. The reachable output indices form the subgroup generated by these differences. When $0 \in D$, this subgroup is generated by D, giving the component count above.

Permuting input or output coordinates and negating individual coordinates preserve both branch numbers and both addition costs. For circulants, translating D, reflecting D to $-D$, or multiplying its indices by an integer coprime to n gives equivalent constructions under coordinate permutations. The coefficients move with their support positions in these transformations.

If coordinate permutations put M in the form $M_1 \oplus \cdots \oplus M_t$, then

$$C_{\mathrm{rect}}(M) = \sum_{i=1}^{t} C_{\mathrm{rect}}(M_i), \qquad C_g(M) \le \sum_{i=1}^{t} C_g(M_i). \tag{2.11}$$

Every admissible rectangle lies within a single component, proving the first statement. Implementing the components separately gives the second. This also describes constructions obtained by increasing the rotation step. If $d \mid n$ and $M = h(P^d)$ for a polynomial h with $\deg h < n/d$, reordering coordinates by their residues modulo d gives d identical blocks, each equal to $h(P_{n/d})$, where $P_{n/d}$ is the cyclic permutation matrix of order n/d.

Checking branch numbers with small supports. After choosing a construction, we must verify its branch number. The following observation reduces this verification to inputs of small weight in the forward and inverse directions.

Lemma 2.3. *Let M be invertible and let $b \geq 3$ be an integer. Then $B_d(M) \geq b$ if and only if*

$$\mathrm{wt}(x) + \mathrm{wt}(Mx) \geq b, \qquad \mathrm{wt}(x) + \mathrm{wt}(M^{-1}x) \geq b \tag{2.12}$$

for every x with $1 \leq \mathrm{wt}(x) \leq \lfloor (b-1)/2 \rfloor$.

Proof. If $B_d(M) \geq b$, both inequalities follow from the definition and $B_d(M^{-1}) = B_d(M)$. Conversely, a violation of the branch number bound gives a nonzero pair $y = Mx$ with $\mathrm{wt}(x) + \mathrm{wt}(y) \leq b-1$. At least one of x and y has weight at most $\lfloor (b-1)/2 \rfloor$. If it is x, the forward check detects the violation. If it is y, the inverse check detects it. □

For $b = 5$ and $b = 6$, the lemma requires only inputs of weights one and two in each direction. Applying it to M^{T} checks the corresponding linear branch number requirement. For circulants, Proposition 2.1 gives equality of the two branch numbers.

2.3 General lower bounds

We derive lower bounds on the general addition cost from the differential and linear branch numbers. For rectangle reuse, we combine these diffusion requirements with row and column weights and invertibility to obtain further cost lower bounds. The results in this subsection do not require circulant structure.

2.3.1 General addition circuits

Theorem 2.4. *Let $M \in \mathbb{F}^{n \times n}$ be invertible. Then*

$$B_l(M) \geq 4 \quad \Longrightarrow \quad C_g(M) \geq 2n - \left\lfloor \frac{2n}{B_d(M)} \right\rfloor, \tag{2.13}$$

$$B_l(M) \geq 5 \quad \Longrightarrow \quad C_g(M) \geq 2n - \left\lfloor \frac{n}{B_d(M)} \right\rfloor. \tag{2.14}$$

The proof is given in Appendix A.2.

To the best of our knowledge, Theorem 2.4 is the first to establish theoretical lower bounds on general circuit addition complexity under prescribed branch number requirements for arbitrary matrix dimensions. Weight information gives another bound that can be useful when rows are comparatively dense.

Proposition 2.5. *If M is invertible and every row has weight at least $w \geq 2$, then*

$$C_g(M) \geq n + w - 2. \tag{2.15}$$

Proof. The n outputs are pairwise nonproportional and none is an input up to sign. Consider the first output produced by the circuit. Its dependence on at least w distinct inputs requires at least $w-1$ gates among its ancestors, including its own gate. Indeed, a connected graph with g gates of fanin at most two can involve at most $g+1$ distinct input vertices. The other $n-1$ outputs require distinct later gates, giving the bound. □

For circulants, $B_d = B_l = B$. Theorem 2.4 therefore gives $C_g \geq \lceil 3n/2 \rceil$ when $B \geq 4$, and $C_g \geq 2n - \lfloor n/b \rfloor$ when $B \geq b \geq 5$. These circuit lower bounds also apply to C_{rect} by (2.10).

2.3.2 Rectangle size constraints and cost lower bounds

We now assume that M has entries in $\{0, 1, -1\}$ and that every row and every column has weight $w \geq 2$. The matrix is invertible, with $B_d(M) \geq b_d$ and $B_l(M) \geq b_l$. The two branch numbers control different dimensions of a reusable rectangle.

Proposition 2.6. *Every admissible $r \times s$ rectangle satisfies*

$$r \leq R := \min\{w - 1, w + 1 - \lceil b_d/2 \rceil\}, \qquad s \leq S := \min\{w - 1, w + 1 - \lceil b_l/2 \rceil\}. \tag{2.16}$$

Proof. Choose two columns of the rectangle. Their restrictions to its r rows are proportional, so a combination with two nonzero coefficients cancels those entries. Each column has at most $w-r$ remaining nonzero entries. This gives $b_d \leq 2+2(w-r)$, hence $r \leq w+1-\lceil b_d/2 \rceil$. If $r = w$, the two entire columns would be proportional, contradicting invertibility. Thus $r \leq w - 1$. Apply the same argument to M^{T} to bound s. □

Invertibility also limits how much of the matrix can be covered by independent rank one rectangles.

Lemma 2.7. *For any rectangle scheme $\mathcal{S}$,*

$$\sum_{Q \in \mathcal{S}} (r_Q s_Q - 1) \leq n(w - 1). \tag{2.17}$$

Proof. Write M as the sum of its selected rectangle matrices, extended by zeros outside their positions, and its uncovered nonzero entries. Every summand has rank one. There are $|\mathcal{S}|$ rectangle summands and $nw - \sum_{Q \in \mathcal{S}} r_Q s_Q$ uncovered entries. Subadditivity of rank gives

$$n \leq |\mathcal{S}| + nw - \sum_{Q \in \mathcal{S}} r_Q s_Q,$$

which is the stated inequality. □

The rank budget bounds total savings. The limited number of positions in each row and column gives additional capacity bounds.

Proposition 2.8. *Suppose every admissible rectangle has height at most R and width at most S. If $R, S \geq 2$, then*

$$C_{\mathrm{rect}}(M) \geq \left\lceil n(w-1)\frac{R+S-2}{RS-1} \right\rceil, \tag{2.18}$$

$$C_{\mathrm{rect}}(M) \geq n(w-1) - \left\lfloor n\frac{R-1}{R}\left(w - \left\lceil \frac{w}{S} \right\rceil\right) \right\rfloor, \tag{2.19}$$

$$C_{\mathrm{rect}}(M) \geq n(w-1) - \left\lfloor n\frac{S-1}{S}\left(w - \left\lceil \frac{w}{R} \right\rceil\right) \right\rfloor. \tag{2.20}$$

If $R < 2$ or $S < 2$, then $C_{\mathrm{rect}}(M) = n(w-1)$.

Proof. A rectangle of size $r \times s$ saves $(r-1)(s-1)$ additions and contributes $rs - 1$ to the rank budget in Lemma 2.7. Thus the ratio $(r-1)(s-1)/(rs-1)$ measures the saving per unit of this budget. It increases with both r and s for $r, s \geq 2$, so every admissible rectangle satisfies

$$(r-1)(s-1) \leq \frac{(R-1)(S-1)}{RS-1}(rs-1).$$

Summing over the selected rectangles and applying the rank budget, the total saving is at most

$$n(w-1)\frac{(R-1)(S-1)}{RS-1}.$$

Subtracting this from the direct cost proves (2.18).

For (2.19), distribute the saving of each $r \times s$ rectangle equally among its r rows. Each such row receives $(r-1)(s-1)/r \le (R-1)(s-1)/R$. In a fixed row, the selected rectangles occupy disjoint groups of nonzero positions. Treat each uncovered nonzero entry as a singleton group. These groups partition the row's w nonzero positions, and each group has size at most S. Hence there are at least $\lceil w/S \rceil$ groups.

The sum of the group sizes minus one for each group equals w minus the number of groups. Singleton groups contribute zero, so the sum of $s-1$ over the rectangle groups is at most $w - \lceil w/S \rceil$. The saving assigned to this row is therefore at most

$$\frac{R-1}{R}\left(w - \left\lceil \frac{w}{S} \right\rceil\right).$$

Summing over all n rows bounds the total saving. This saving is an integer, so we may take the floor of its upper bound. Subtracting from $n(w-1)$ proves (2.19). The same argument applied column by column, with R and S interchanged, proves (2.20). If one dimension bound is less than two, no admissible rectangle exists. □

If only a differential branch number bound is available, we may use $R = \min\{w-1, w+1-\lceil b_d/2 \rceil\}$ and $S = w-1$. For example, $b_d \ge 5$ and $w = 4$ still give $C_{\text{rect}} \ge 2n$ from the row capacity bound.

2.3.3 Rectangle cost lower bounds at minimum weight

Suppose now that both branch numbers are at least $b \ge 4$ and that the row and column weight is the minimum possible value $w = b-1$. Proposition 2.6 bounds both rectangle dimensions by $\lfloor b/2 \rfloor$. Combining the rank budget with the row and column capacities gives the following closed bounds.

Theorem 2.9. *Let $M \in \{0,1,-1\}^{n\times n}$ be invertible, with row and column weight $w = b-1$ and $B_d(M), B_l(M) \ge b$. For $R \ge 2$,*

$$b = 2R+1: \qquad C_{\text{rect}}(M) \ge \left\lceil \frac{2(2R-1)}{R+1}\, n \right\rceil, \tag{2.21}$$

$$b = 2R: \qquad C_{\text{rect}}(M) \ge \left\lceil \frac{2(R-1)(2R-1)}{R^2}\, n \right\rceil. \tag{2.22}$$

The proof is given in Appendix A.3.

Table 2 records the three branch number targets used in Sections 3 and 4. The circuit column follows from Theorem 2.4, while the rectangle column uses the additional minimum weight assumption. The stronger rectangle bounds obtained from circulant structure will be proved in those sections.

Integer refinements for fixed dimensions. At a fixed n, the bounds can be strengthened by counting rectangles of each permitted size. Impose the rank budget and enumerate the possible partitions of a row's w positions into rectangle widths and uncovered single entries. The total number of row occurrences of each width must agree with the rectangle counts. Apply the same constraints to column heights, and maximize the total saving over these integer counts. Subtracting this maximum from $n(w-1)$ gives a lower bound on C_{rect}. These constraints record necessary counts. They do not ensure that the rectangles can be placed in a matrix meeting the branch number requirements.

Table 2: General lower bounds when both branch numbers are at least b and every row and column has weight $w = b - 1$.

b	w	Lower bound on C_g	Lower bound on C_{rect}
4	3	$\lceil 3n/2 \rceil$	$\lceil 3n/2 \rceil$
5	4	$2n - \lfloor n/5 \rfloor$	$2n$
6	5	$2n - \lfloor n/6 \rfloor$	$\lceil 20n/9 \rceil$

3 Circulant Layers over $\mathbb{F}_2$

Over $\mathbb{F}_2$, additions and subtractions coincide, and each counted gate is a two-input XOR. A weight-one circulant is a power of P, so it only permutes coordinates and has addition cost zero. We study layers with $w > 1$. Invertibility forces w to be odd, as shown in Section 2.1, and hence $w \geq 3$. Every row and every column of a circulant has the same weight w. Since every column has odd weight, the input and output Hamming weights have the same parity. Their sum is therefore even, so $B(M)$ is even. Consequently, the requirement $B(M) \geq 3$ is equivalent to $B(M) \geq 4$. We begin with weight three, the minimum weight for this target. Similarly, $B(M) \geq 5$ is equivalent to $B(M) \geq 6$. Section 3.2 studies the classification and construction of layers with branch number at least six.

3.1 Branch number four and weight three

Throughout this subsection, M is an invertible binary circulant of order n and weight three. We take $n \geq 4$, since the weight-three matrix of order three is the all-one matrix and is singular. All minima are over matrices with this order, field, and weight. The connectivity requirement is specified separately.

3.1.1 Exact rectangle costs at weight three

Proposition 3.1. *Every invertible binary circulant of weight three has $B_d(M) = B_l(M) = 4$.*

Proof. A weight-one input gives an output of weight three, so $B_d(M) \leq 4$. Since every column has odd weight, the sum of the output coordinates equals the sum of the input coordinates. Hence $\mathrm{wt}(x) + \mathrm{wt}(Mx)$ is even. For a nonzero input, invertibility ensures a nonzero output. The only possible total weight below four would therefore be two. This would require an input and output both of weight one, contrary to the column weight. Thus $B_d(M) = 4$, and Proposition 2.1 gives the same linear branch number. □

Proposition 2.6 now restricts every admissible rectangle to size 2×2. Define a graph $\Gamma(M)$ whose vertices are the output rows. Join two rows when they share two nonzero column positions. Their common positions form an admissible rectangle. Two distinct rows cannot share all three positions, since they would then be equal and M would be singular.

A matching in this graph is a set of edges with no common endpoint. Let $\mu(\Gamma(M))$ denote the maximum number of edges in a matching of $\Gamma(M)$.

Proposition 3.2. *The rectangle cost is*

$$C_{\text{rect}}(M) = 2n - \mu(\Gamma(M)). \tag{3.1}$$

Proof. Each 2×2 rectangle saves one addition. A row has only three nonzero positions, so it cannot accommodate two position-disjoint rectangles, each of which would occupy two positions in that row. Thus the edges corresponding to a rectangle scheme form a matching. Conversely, rectangles corresponding to a matching use disjoint row sets, so their matrix

positions are disjoint and they form a valid scheme. Maximizing the number of rectangles therefore gives the maximum matching size. Subtract it from the direct cost $2n$ to obtain (3.1). □

The support D determines this matching problem completely. We call two support positions antipodal if their difference is $n/2$, which requires even n.

Theorem 3.3. *The rectangle cost has the following classification.*

1. *If n is even and D contains an antipodal pair, then*

$$C_{\mathrm{rect}}(M) = \frac{3n}{2}.$$

2. *If D has no antipodal pair and $D = \{u, u+d, u+2d\}$, put $g = \gcd(n, d)$ and $q = n/g$. Such a support is invertible exactly when $3 \nmid q$. In the invertible case, $q \geq 5$ and*

$$C_{\mathrm{rect}}(M) = 2n - g\left\lfloor \frac{q}{2} \right\rfloor.$$

3. *For every other invertible support, $C_{\mathrm{rect}}(M) = 2n$.*

The proof is given in Appendix B.1. The second case was also considered by Rishakani et al. [15, Lemma 4.1], who gave a similar scheme for sharing XOR computations when the support forms an arithmetic progression. For supports in the third case, invertibility is checked by $\gcd(f, X^n - 1) = 1$. The classification shows that no rectangle reuse is possible in this case.

3.1.2 Optimal costs, attaining constructions, and the effect of connectivity

The classification yields an exact minimum rectangle cost at every order where weight three is feasible.

Theorem 3.4. *For even $n \geq 4$,*

$$\min C_g = \min C_{\mathrm{rect}} = \frac{3n}{2}. \tag{3.2}$$

For odd $n \geq 5$, write $n = 3^e n_0$, where $e \geq 0$ is an integer, $n_0 \geq 1$, and $3 \nmid n_0$. Then

$$\min C_{\mathrm{rect}} = \begin{cases} (3n + 3^e)/2, & n_0 > 1, \\ 2n, & n_0 = 1. \end{cases} \tag{3.3}$$

If $e = 0$, the general circuit minimum also satisfies

$$\min C_g = \min C_{\mathrm{rect}} = \frac{3n+1}{2}. \tag{3.4}$$

The minima in (3.2) *and* (3.4) *are attained by connected matrices. The lower bound on C_{rect} in* (3.3) *does not in general give the same lower bound on C_g.*

Proof. Theorem 2.4 gives $C_g \geq \lceil 3n/2 \rceil$ for every matrix in this subsection. For even n, take $M = I + P + P^{n/2}$. Its square is P^2, so it is invertible. Compute the $n/2$ sums $x_i + x_{i+n/2}$ for $0 \leq i < n/2$. Every output is one of these sums plus one further input, giving $n/2 + n = 3n/2$ gates. The support contains $0, 1$, so the matrix is connected. This proves (3.2).

For odd n, there are no antipodal pairs. By Theorem 3.3, a matrix with rectangle cost below $2n$ must have arithmetic progression support. For such a support, $q = n/g$ is odd, and its cost is

$$2n - g\frac{q-1}{2} = \frac{3n+g}{2}.$$

Invertibility requires $3 \nmid n/g$. Thus g contains the full factor 3^e of n, giving $g \geq 3^e$. If $n_0 > 1$, the choice

$$D = \{0, 3^e, 2 \cdot 3^e\}$$

has $g = 3^e$ and $q = n_0$, is invertible, and attains $(3n + 3^e)/2$. Every support without reuse costs $2n$, which is larger. This proves the first odd case.

If $n_0 = 1$, every arithmetic progression support is singular, since its cycle length is a positive power of three. All invertible supports therefore cost $2n$. This class is nonempty: take $M = I + P + P^3$. The polynomial $1 + X + X^3$ is irreducible over $\mathbb{F}_2$, as it has no root in $\mathbb{F}_2$. Its roots lie in $\mathbb{F}_8^*$ and have order seven, so it is coprime to $X^n - 1$ for every power of three. This proves the second odd case.

Finally, if $e = 0$, the support $\{0, 1, 2\}$ attains $(3n+1)/2$, which equals the general circuit lower bound. Its support graph is connected. This proves (3.4). □

When $e \geq 1$ and $n_0 > 1$, the attaining support above has 3^e connected components. Reordering coordinates separates it into 3^e identical copies of the layer $I + P_{n_0} + P_{n_0}^2$ of order n_0. Requiring connectivity removes this way of attaining the minimum.

Proposition 3.5. *Let $n \geq 9$ be odd with $3 \mid n$. Every connected invertible binary circulant of weight three has no admissible rectangle and satisfies*

$$C_{\mathrm{rect}}(M) = 2n. \tag{3.5}$$

Proof. If a rectangle existed, the odd-order classification would force $D = \{u, u+d, u+2d\}$. Its support graph has $\gcd(n, d)$ components, so connectivity requires $\gcd(n, d) = 1$. Its cycle length would then be n, divisible by three, contradicting the invertibility criterion in Theorem 3.3. Thus there is no rectangle saving, and the cost is $2n$. □

For example, at $n = 15$ the support $\{0, 3, 6\}$ attains rectangle cost 24 with three components. Every admissible connected matrix has rectangle cost 30. The connected class is nonempty, as witnessed by $\{0, 1, 3\}$.

3.1.3 General circuits beyond rectangle reuse: constructions and depth tradeoffs

The absence of a reusable rectangle does not prevent a general circuit from saving gates. The following construction uses a shared three-term intermediate expression and cancels a different input in each of two outputs.

Proposition 3.6. *For $n \geq 8$ with $7 \nmid n$, the connected matrix $M = I + P + P^3$ is invertible, has branch number four, and satisfies*

$$C_{\mathrm{rect}}(M) = 2n, \qquad C_g(M) \leq 2n - \left\lfloor \frac{n}{5} \right\rfloor. \tag{3.6}$$

The general circuit upper bound has an implementation of depth three.

Proof. The polynomial argument in Theorem 3.4 shows that $I + P + P^3$ is invertible exactly when $7 \nmid n$. Its support contains $0, 1$, so it is connected, and Proposition 3.1 gives branch number four. The six nonzero directed differences between support positions are $\pm 1, \pm 2, \pm 3$. They are distinct modulo n for $n \geq 8$, so two rows share at most one input. No rectangle is available, giving $C_{\mathrm{rect}}(M) = 2n$.

The outputs are $y_i = x_i + x_{i+1} + x_{i+3}$. For five consecutive outputs, first compute

$$t_1 = x_1 + x_3, \qquad t_2 = x_3 + x_4, \qquad t_3 = x_4 + x_5, \qquad t_4 = x_2 + t_2.$$

Then compute

$$\begin{aligned} y_0 &= x_0 + t_1, \quad y_1 = t_4 + t_1, \quad y_2 = t_4 + t_3, \\ y_3 &= x_6 + t_2, \quad y_4 = x_7 + t_3. \end{aligned}$$

The two central outputs follow from the cancellations

$$\begin{aligned} t_4 + t_1 &= (x_2 + x_3 + x_4) + (x_1 + x_3) = x_1 + x_2 + x_4, \\ t_4 + t_3 &= (x_2 + x_3 + x_4) + (x_4 + x_5) = x_2 + x_3 + x_5. \end{aligned}$$

There are four intermediate gates and five output gates, for a total of nine instead of ten. The expressions t_1, t_2, t_3 have depth one, t_4 has depth two, and the outputs have depth at most three.

Partition the outputs into $\lfloor n/5 \rfloor$ consecutive groups of five and a remainder. For each full group starting at output y_j, apply the same nine-gate construction with every input and output index increased by j. Compute each remaining output in two gates. Indices are taken modulo n, and the eight input positions within each template are distinct for $n \geq 8$. The total cost is

$$9 \left\lfloor \frac{n}{5} \right\rfloor + 2 \left(n - 5 \left\lfloor \frac{n}{5} \right\rfloor \right) = 2n - \left\lfloor \frac{n}{5} \right\rfloor.$$

Different groups may share inputs, which are available with free fanout. Each group computes its own intermediate expressions, so the full circuit has depth at most three. □

For $n = 3^e \geq 9$, the minimum rectangle cost is $2n$ and the construction is invertible. Thus the minima over all admissible weight-three matrices satisfy

$$\left\lceil \frac{3n}{2} \right\rceil \leq \min C_g \leq 2n - \left\lfloor \frac{n}{5} \right\rfloor < 2n = \min C_{\mathrm{rect}}. \tag{3.7}$$

For odd $n \geq 9$ with $3 \mid n$ and $7 \nmid n$, the same strict separation holds when both minima are restricted to connected matrices, by Proposition 3.5.

The depth required to improve on rectangle reuse can also be determined exactly.

Proposition 3.7. *For an invertible binary circulant of weight three, the minimum gate count among circuits of depth at most two equals $C_{\mathrm{rect}}(M)$. Consequently, every implementation with fewer than $C_{\mathrm{rect}}(M)$ gates has depth at least three.*

Proof. A signal of depth at most one is an input, a sum of two distinct inputs, or zero. Each output has weight three and therefore needs a final gate at depth two. Two inputs produce weight at most two, while the XOR of two two-input sums has even weight. Thus the final gate must combine a two-input sum with one further input. These three input positions must be distinct, since a repetition would cancel and leave weight one.

Every output therefore uses one original input pair and a third input. A pair computed at depth one can be shared among outputs, and these occurrences define a rectangle. Pairs used by only one output are private computations. Each row uses only one pair, so the shared

rectangles have disjoint matrix positions. The circuit is a rectangle scheme and cannot have cost below $C_{\mathrm{rect}}(M)$.

Conversely, every rectangle scheme in this class shares only pairs. Compute the shared and private pairs at depth one and complete all outputs at depth two. An optimal rectangle scheme therefore attains its cost at depth two, proving equality. □

The circuit in Proposition 3.6 attains depth three while improving on its matrix's rectangle cost. Thus three is the least depth permitting such an improvement in this class. Direct computation of each output still has depth two. Thus, designing binary circulant matrices of weight three and branch number four involves a tradeoff between circuit depth and the number of two-input XOR gates. For some orders, a reduction in gate count requires greater depth.

3.1.4 Constructions at representative orders

For $4 \le n \le 33$, we enumerated every three-point support D that includes position 0. The search and optimization framework used here and in later sections, including its integer programming formulation and associated code, is described in Section 5. Any three-point support can be put in this form by subtracting one of its positions from all three positions modulo n. This common shift preserves invertibility, branch number, and both costs. We checked invertibility by polynomial gcd and computed the maximum matching of $\Gamma(M)$ independently of the cost classification. All 4,617 invertible supports in this enumeration agreed with Theorem 3.3, and the minimum costs agreed with Theorem 3.4. These counts include supports equivalent under further coordinate permutations. We also verified the nine-gate template and its implementations with shifted indices as exact binary linear expressions for $8 \le n \le 128$. Detailed verification data are given in Appendix B.2.

Table 3 gives representative orders. Each rectangle entry is an exact minimum in the indicated matrix class. Intervals in the C_g columns combine the lower bounds in Theorem 2.4, the circuit upper bound in Proposition 3.6, and the inequality $C_g(M) \le C_{\mathrm{rect}}(M)$.

Table 3: Weight-three binary circulants with branch number four. The two column groups use different connectivity requirements. Intervals are certified bounds on the corresponding minimum C_g.

	No connectivity constraint		Connected	
n	$\min C_{\mathrm{rect}}$	$\min C_g$	$\min C_{\mathrm{rect}}$	$\min C_g$
4	6	6	6	6
5	8	8	8	8
6	9	9	9	9
7	11	11	11	11
8	12	12	12	12
9	18	$[14, 17]$	18	$[14, 17]$
15	24	$[23, 24]$	30	$[23, 27]$
21	33	$[32, 33]$	42	$[32, 42]$
27	54	$[41, 49]$	54	$[41, 49]$

For $n = 4, 6, 8$, the values $6, 9, 12$ also agree with the XOR optima reported by Cui et al. [2] for binary diffusion layers of branch number four. Here they follow from the general circuit lower bound and explicit connected circulant constructions.

The orders 15, 21, and 27 illustrate different features of the results. At order 15, the minimum rectangle cost depends on whether connectivity is required. Order 21 includes a factor of seven, so the construction in Proposition 3.6 is singular. Order 27 illustrates the

strict separation between the two minima at powers of three, even without a connectivity requirement.

The intervals at orders 9, 15, 21, and 27 remain open for general circuits. Their upper endpoints are supplied by the nine-gate construction in Proposition 3.6 or by rectangle implementations. Constructions with fewer gates may exist within the stated class of invertible binary circulants of weight three.

3.2 Branch number six and weight five

The requirement $B(M) \geq 6$ forces $w \geq 5$. At weight five, this requirement is exactly $B(M) = 6$. Throughout this subsection, the matrices are invertible binary circulants of weight five. Unlike the weight-three case, a reusable rectangle can have three rows or three columns. We first establish a rectangle cost lower bound by distinguishing supports with and without a 3×3 rectangle. We then report exact minimum rectangle costs at selected orders.

3.2.1 A rectangle cost lower bound at weight five

Substituting $w = 5$ and $B_d = B_l = 6$ into Theorem 2.9 and Theorem 2.4 gives the following general lower bounds on C_{rect} and C_g, independent of circulant structure:

$$C_{\text{rect}}(M) \geq \left\lceil \frac{20n}{9} \right\rceil, \qquad C_g(M) \geq 2n - \left\lfloor \frac{n}{6} \right\rfloor. \tag{3.8}$$

We can strengthen the rectangle bound to $\lceil 12n/5 \rceil$ by using the circulant support. We first give an equivalent condition for branch number six in this setting.

Proposition 3.8. *An invertible binary circulant of weight five has branch number six if and only if $c(d) \leq 3$ for every nonzero d and every row of M^{-1} has weight at least five. The inverse M^{-1} is also circulant, so all its row and column weights are equal. If all nonzero overlaps satisfy $c(d) \leq 2$, then $B(M) = 6$ without any additional test of the row weight of M^{-1}.*

The proof is given in Appendix B.3.

For later constructions, the possible obstruction among three columns can be described more narrowly. Under $c(d) \leq 3$, a sum of weight one would require pairwise overlaps $3, 2, 2$ with empty triple intersection, or pairwise overlaps $3, 3, 3$ with triple intersection of size one. Indeed, the weight formula in the proof forces these two possibilities. The only other numerical candidate is $3, 3, 1$ with empty triple intersection, but two subsets of size three in a column of weight five must intersect.

Proposition 2.6 bounds both sides of every rectangle by three. The allowed sizes are therefore 2×2, 2×3, 3×2, and 3×3. For a scheme $\mathcal{S}$, let $N_{22}, N_{23}, N_{32}, N_{33}$ count its rectangles of these four sizes. The exact packing objective is

$$C_{\text{rect}}(M) = 4n - \max_{\mathcal{S}} (N_{22} + 2N_{23} + 2N_{32} + 4N_{33}). \tag{3.9}$$

The positions occupied by distinct selected rectangles must remain disjoint.

With the rectangle sizes restricted as above, we combine Proposition 2.2, Lemma 2.7, and Proposition 2.8. We use $(r-1)(s-1)/(rs-1)$ as the ratio of XOR saving to rank budget contribution. This ratio is $1/3$ for 2×2 rectangles, $2/5$ for 2×3 and 3×2 rectangles, and $1/2$ for 3×3 rectangles. We therefore distinguish whether the support contains a 3×3 rectangle.

We first consider supports containing a 3×3 rectangle. The following classification identifies the structure needed to bound their rectangle savings.

Theorem 3.9. *Let $D \subseteq \mathbb{Z}_n$ have five elements and satisfy $c(d) \leq 3$ for every nonzero d. Its circulant support contains a 3×3 all-one submatrix if and only if $3 \mid n$ and D contains a coset of the subgroup of order three.*

The proof is given in Appendix B.4.

After a common shift, a support in this case has the form $D = \{0, n/3, 2n/3, \delta_1, \delta_2\}$ with $3 \mid n$. The additional positions δ_1, δ_2 have distinct nonzero residues modulo $n/3$. If they had the same residue, shifting by $n/3$ would produce overlap at least four, contradicting $B(M) = 6$. Write $H = I + P^{n/3} + P^{2n/3}$ and call its nonzero positions core positions. After reordering coordinates, they form $n/3$ disjoint 3×3 all-one blocks. We call each of these a core block: it is a reusable 3×3 rectangle whose nine positions all belong to the support of H. The two other shifts supply two additional positions in every row and column.

We next prove that every 2×3 rectangle lies within a core block. Such a rectangle consists of two rows and three columns in which both rows are nonzero. We therefore count the common nonzero columns of two rows, keeping track of whether each entry is core or additional in its own row.

It suffices to compare row 0 with row d, since subtracting the same index from all row and column indices preserves the support. Their nonzero columns are

	Core columns	Additional columns
Row 0	$0,\ n/3,\ 2n/3$	$\delta_1,\ \delta_2$
Row d	$d,\ d+n/3,\ d+2n/3$	$d+\delta_1,\ d+\delta_2.$

Column indices are taken modulo n. If $d = \pm n/3$, the three core columns coincide. There can be no further common column, since $c(d) \leq 3$. These two row differences therefore give only rectangles within a core block.

Now assume $d \not\equiv 0 \pmod{n/3}$. The two lists of core columns are disjoint. A common column must consequently use an additional entry in at least one of the two rows. The four possible ways to match a core entry with an additional entry are listed below. Each satisfied condition contributes exactly one common column.

Entry in row 0	Entry in row d	Condition modulo $n/3$	Common column
Additional	Core	$d \equiv \delta_1$	δ_1
Additional	Core	$d \equiv \delta_2$	δ_2
Core	Additional	$d \equiv -\delta_1$	$d+\delta_1$
Core	Additional	$d \equiv -\delta_2$	$d+\delta_2$

We also need to count common columns in which both entries are additional. The following lemma excludes two such columns.

Lemma 3.10. *Two distinct rows of M have at most one common column in which both entries are additional.*

Proof. Compare rows 0 and d, where $d \neq 0$. Two common columns with both entries additional would require $\delta_2 \equiv d + \delta_1$ and $\delta_1 \equiv d + \delta_2 \pmod{n}$. Since $\delta_1 \neq \delta_2$, these congruences give $\delta_2 - \delta_1 \equiv n/2 \pmod{n}$, so $P^{2\delta_1} = P^{2\delta_2}$. Recall that

$$M = H + P^{\delta_1} + P^{\delta_2}, \qquad H = I + P^{n/3} + P^{2n/3}.$$

All these summands are powers of P and commute. When squaring over $\mathbb{F}_2$, the cross terms cancel in pairs. Thus

$$H^2 = I + P^{2n/3} + P^{4n/3} = I + P^{2n/3} + P^{n/3} = H,$$

where $P^n = I$, and consequently

$$M^2 = H^2 + P^{2\delta_1} + P^{2\delta_2} = H.$$

The matrix H is rank deficient because rows whose indices differ by $n/3$ are equal. Hence M^2 is singular, which forces M to be singular, contradicting invertibility. □

Lemma 3.11. *If $d \not\equiv 0 \pmod{n/3}$, rows 0 and d of M have at most two common nonzero columns.*

Proof. The table above contributes at most two common columns. Indeed, the first two conditions cannot both hold because δ_1 and δ_2 have distinct residues, and neither can the last two. To obtain three common columns, we must therefore satisfy two conditions in the table and one additional/additional equality. There are just two possibilities for the pair of table conditions.

First, both conditions might involve δ_1. Then $d \equiv \delta_1 \equiv -\delta_1 \pmod{n/3}$, so $2\delta_1 \equiv 0$. The additional/additional equality requires $d \equiv \pm(\delta_2 - \delta_1)$. Combining these conditions gives either $\delta_2 \equiv 2\delta_1 \equiv 0$ or $\delta_2 \equiv 0$ modulo $n/3$. Both contradict the nonzero residue of δ_2. If both conditions involve δ_2, the same argument forces $\delta_1 \equiv 0 \pmod{n/3}$, again a contradiction.

Second, the two conditions might involve different extra offsets. They then require $\delta_2 \equiv -\delta_1$ and $d \equiv \pm\delta_1$ modulo $n/3$. The additional/additional equality also requires $d \equiv \pm(\delta_2 - \delta_1)$, which now equals $\pm 2\delta_1$ modulo $n/3$. Equating these two expressions for d gives either $\delta_1 \equiv 0$, already excluded, or $3\delta_1 \equiv 0$. In the latter case the three residues $0, \delta_1, -\delta_1$ are distinct and form a subgroup of order three modulo $n/3$.

To exclude this last possibility, we construct a nonzero vector x with $Mx = 0$. Define

$$x_i = \begin{cases} 1, & i \equiv 0 \text{ or } \delta_1 \pmod{n/3}, \\ 0, & \text{otherwise.} \end{cases}$$

In particular $x_0 = 1$, so $x \neq 0$. Since the definition depends only on the residue modulo $n/3$, we have $x_i = x_{i+n/3} = x_{i+2n/3}$. The three core terms therefore sum to x_i over $\mathbb{F}_2$, giving

$$\begin{aligned}(Mx)_i &= x_i + x_{i+n/3} + x_{i+2n/3} + x_{i+\delta_1} + x_{i+\delta_2} \\ &= x_i + x_{i+\delta_1} + x_{i+\delta_2}.\end{aligned}$$

Using $\delta_2 \equiv -\delta_1$ and $2\delta_1 \equiv -\delta_1$ modulo $n/3$, the remaining three terms can be checked separately for each possible residue of i:

$i \pmod{n/3}$	x_i	$x_{i+\delta_1}$	$x_{i+\delta_2}$
0	1	1	0
δ_1	1	0	1
$-\delta_1$	0	1	1
All other residues	0	0	0

For the last row, if adding or subtracting δ_1 brought i into the subgroup $\{0, \delta_1, -\delta_1\}$, then i would already belong to that subgroup. Hence neither shifted index can have residue 0 or δ_1. Every row of the table sums to zero over $\mathbb{F}_2$. Thus $Mx = 0$ with $x \neq 0$, contradicting invertibility. Both possibilities for three common columns have now been excluded. □

Thus rows from different core blocks have at most two common nonzero columns. Every 2×3 rectangle lies within a core block. Applying the same argument to M^{T}, whose extra offsets are $-\delta_1, -\delta_2$, gives the same conclusion for 3×2 rectangles. A 3×3 rectangle contains such a 2×3 rectangle and is also confined to one core block.

A rectangle crossing core blocks must therefore have size 2×2. It has either one core position or two diagonally opposite core positions, as shown in Figure 2.



(a) One core position. (b) Two opposite core positions.

Figure 2: The two possible forms of a 2×2 rectangle crossing core blocks, shown schematically. Blue diagonals represent core positions, the nonzero positions of H. Red diagonals represent additional positions. The four dots joined by black edges select the rectangle in each panel.

Indeed, zero core positions would give a four-cycle using just the two additional shifts, forcing $\delta_2 - \delta_1 = n/2$, which was excluded in the proof of Lemma 3.10. Three core positions force the fourth to be core. Two core positions in the same row or column would force the two additional residues to coincide.

Each core position occupied by a crossing rectangle also uses an additional position in its row and column. There are only two additional positions in either, so the occupied core positions have row and column degrees at most two. For $z = 0, 1, \ldots, 6$ occupied positions in a core block, the remaining internal saving is bounded by

z	0	1	2	3	4	5	6
Remaining saving	4	2	2	1	1	0	0.

(3.10)

Theorem 3.12. *Every invertible binary circulant of weight five and branch number six satisfies*

$$C_{\mathrm{rect}}(M) \geq \left\lceil \frac{12n}{5} \right\rceil. \tag{3.11}$$

If its support contains a 3×3 rectangle, equivalently a three-point subgroup coset, the stronger bound is

$$C_{\mathrm{rect}}(M) \geq \left\lceil \frac{22n}{9} \right\rceil.$$

Neither bound requires connectivity.

Proof. If a 3×3 rectangle exists, Theorem 3.9 and the rectangle structure just established apply. Crossing rectangles with one or two core positions may coexist. Let their respective counts be N_1, N_2. They occupy $N_1 + 2N_2$ core positions and use $3N_1 + 2N_2$ additional positions. Since there are $2n$ additional positions, $3N_1 + 2N_2 \leq 2n$. Every entry of the loss table (3.10) is at most $4 - 2z/3$. Also, since $N_2 \geq 0$, we have $N_1 - N_2 \leq N_1 + 2N_2/3 = (3N_1 + 2N_2)/3 \leq 2n/3$. Consequently the total saving, including crossing rectangles, is at most

$$\frac{4n}{3} - \frac{2}{3}(N_1 + 2N_2) + N_1 + N_2 = \frac{4n}{3} + \frac{N_1 - N_2}{3} \leq \frac{14n}{9}.$$

Subtracting this saving from $4n$ gives cost at least $22n/9$ for supports containing a 3×3 rectangle.

Now suppose no 3×3 rectangle exists. The remaining rectangle sizes have saving-to-budget ratios $1/3, 2/5, 2/5$. Lemma 2.7 bounds the total rank budget contribution by $4n$, so the total saving is at most $8n/5$. The cost is therefore at least $4n - 8n/5 = 12n/5$. Since $22n/9 > 12n/5$, the latter bound applies to both cases. Taking ceilings proves the stated bounds. □

Whether this lower bound is attainable at general orders remains open.

3.2.2 Finite optimal rectangle costs

Table 4 lists the exact minimum rectangle costs for $n \leq 16$ and for $n = 32, 64$. The small-order values were obtained by exhaustive search; the values at orders 32 and 64 were established by finite searches and optimality certificates. We include powers of two, such as 32 and 64, since these orders are often used in linear layers of block ciphers and hash functions. At every feasible order listed, the minimum is the same with and without a connectivity requirement. For comparison, the table also evaluates the rectangle cost lower bound $\lceil 12n/5 \rceil$ proved above. At all feasible orders listed, the exact minimum exceeds this bound, so the bound is generally not attained.

Table 4: Exact minimum rectangle costs for invertible binary circulants of weight five and branch number six. The rectangle minima apply both with and without a connectivity requirement.

n	$\lceil 12n/5 \rceil$	$\min C_{\mathrm{rect}}$
$1 \leq n \leq 9$	—	Infeasible
10	24	27
11	27	30
12	29	31
13	32	35
14	34	38
15	36	40
16	39	44
32	77	86
64	154	171

4 Circulant Layers over $\mathbb{F}_p$

Throughout this section, p is an odd prime and the matrix entries belong to $\{0, 1, -1\}$. Each counted gate performs a field addition or subtraction. As in the binary case, a weight-one layer only permutes coordinates and changes signs, so its addition cost is zero. The first nontrivial case is weight two and branch number three.

Section 2.1 gives the existence conditions for this case. After an overall sign change and a support shift, every invertible weight-two layer has the form $I + P^d$, where $d \neq 0$ and $n/\gcd(n, d)$ is odd. Its branch number is three. The addition cost is also exact.

Proposition 4.1. *Every invertible signed circulant M of weight two has branch number three and satisfies*

$$C_g(M) = C_{\mathrm{rect}}(M) = n. \tag{4.1}$$

These costs are attained by a circuit of depth one.

Proof. Each output contains two distinct input coordinates and therefore requires an addition or subtraction gate. Invertibility ensures that no two outputs are equal up to sign, so they require n distinct gates. Hence $C_g(M) \geq n$. At weight two, an admissible rectangle would contain both nonzero positions of two rows and make those rows proportional, contradicting invertibility. Thus no admissible rectangle exists, and $C_{\mathrm{rect}}(M) = C_0(M) = n$. Together with $C_g(M) \leq C_{\mathrm{rect}}(M)$, this proves the cost equality. Computing each output directly uses one addition or subtraction. All gates use only original inputs, so the depth is one. □

For every odd $n \geq 3$, the matrix $I + P$ attains these costs and is connected. More generally, if $q > 1$ is an odd divisor of n, then $I + P^{n/q}$ gives an order-n attaining matrix. If n is a power of two, no invertible signed circulant of weight two exists, and allowing greater weight gives $C_{\text{rect}}(M) \geq C_g(M) \geq n + 1$ by Proposition 2.5. We next consider branch numbers four, five, and six at their respective minimum weights.

4.1 Branch number four and weight three

Throughout this subsection, M is an invertible signed circulant over $\mathbb{F}_p$, with p odd, order $n \geq 4$, and weight three. We first determine which matrices have branch number four and classify their rectangle costs. We then give connected constructions attaining the general circuit lower bound at every such order.

4.1.1 Exact rectangle costs at weight three

Invertibility alone no longer guarantees branch number four. The following criterion identifies the additional support restriction.

Proposition 4.2. *The matrix M has $B_d(M) = B_l(M) = 4$ if and only if its support D is not a coset of a subgroup of order three. The excluded supports are precisely*

$$D = \{u, u + n/3, u + 2n/3\}, \qquad 3 \mid n.$$

The proof is given in Appendix C.1.

For the rest of this subsection assume $B(M) = 4$. Proposition 2.6 restricts admissible rectangles to size 2×2. Define $\Gamma(M)$ on the output rows, joining two rows when they share two columns and the corresponding 2×2 block has rank one. The rank-one condition ensures that the two outputs can use the same input sum or difference, possibly with opposite signs.

As in Proposition 3.2,

$$C_{\text{rect}}(M) = 2n - \mu(\Gamma(M)). \tag{4.2}$$

Indeed, each rectangle saves one gate. Since each row has only three nonzero positions, two position-disjoint rectangles cannot use the same row. A rectangle scheme therefore gives a matching in $\Gamma(M)$. Conversely, a matching gives rectangles on disjoint row sets and hence a valid scheme. This proves the formula.

Theorem 4.3. *For an invertible signed circulant of weight three and branch number four, the rectangle cost is classified as follows.*

1. *If D contains an antipodal pair, then*

$$C_{\text{rect}}(M) = \frac{3n}{2}.$$

2. *Suppose D has no antipodal pair and $D = \{u, u + d, u + 2d\}$. Put $g = \gcd(n, d)$ and $q = n/g$, and let $a, b, c \in \{1, -1\}$ be the coefficients at these three positions, respectively. Then $q \geq 5$, and*

$$C_{\text{rect}}(M) = \begin{cases} 2n - g\lfloor q/2 \rfloor, & a = c, \\ 2n, & a = -c. \end{cases}$$

   The middle coefficient b does not affect this cost.

3. *Every remaining support has $C_{\text{rect}}(M) = 2n$.*

Proof. If D contains $u, u+n/2$, let their coefficients be a, c. Rows $i, i+n/2$ share the columns $i+u, i+u+n/2$, on which their entries form $\left(\begin{smallmatrix} a & c \\ c & a \end{smallmatrix}\right)$. Its determinant is $a^2 - c^2 = 0$, so these positions form an admissible rectangle for either choice of signs. Pairing every row with its opposite gives a matching of size $n/2$, the largest possible. Equation (4.2) gives the first cost.

The support argument in the proof of Theorem 3.3 also applies here: two directed steps of the same length among three support points form either a two-point cycle or a three-point path. The cycle gives an antipodal pair, and the path gives a three-point arithmetic progression. Thus these are the only supports that can have two common columns. In the progression case, a second choice of middle point would force $3d = 0$, giving the subgroup coset excluded by Proposition 4.2. Consequently, without an antipodal pair, the only possible shared row pairs have index difference d or $-d$.

For rows $i, i+d$, the common columns $i+u+d, i+u+2d$ give the block $\left(\begin{smallmatrix} b & c \\ a & b \end{smallmatrix}\right)$. Its determinant is $b^2 - ac = 1 - ac$. Since p is odd, it vanishes exactly when $a = c$. If $a = -c$, none of the possible row pairs gives an admissible rectangle, so the cost is $2n$. If $a = c$, the row graph joins i to $i \pm d$ and consists of g cycles of length q. Each cycle contributes $\lfloor q/2 \rfloor$ matching edges, as proved in Theorem 3.3. Substituting into (4.2) gives the stated cost. Here $q = 3$ would be the excluded subgroup case and $q = 4$ would give an antipodal pair, so $q \geq 5$.

All other supports have no pair of rows sharing two columns, hence no rectangle saving. Their cost is $2n$. □

This classifies the costs of matrices already assumed invertible. For a given signed support, invertibility is still checked by $\gcd(f, X^n - 1) = 1$ as in Proposition 2.1. In particular, identical support positions can have different rectangle costs when their signs differ.

4.1.2 Optimal costs and attaining constructions

Unlike the binary case, the general circuit lower bound is attained at every order $n \geq 4$.

Theorem 4.4. *For every odd prime p and every $n \geq 4$, the minima over invertible signed circulants of weight three and branch number four are*

$$\min C_g = \min C_{\text{rect}} = \left\lceil \frac{3n}{2} \right\rceil. \tag{4.3}$$

Both minima are attained by connected matrices. Explicit attaining matrices are

$$\begin{cases} I - P + P^2, & n \geq 5 \text{ odd}, \\ I + P + \sigma P^{n/2}, & n \geq 4 \text{ even}, \end{cases}$$

where $\sigma \in \{1, -1\}$ is chosen so that

$$(-2)^{n/2} \neq \sigma \pmod{p}. \tag{4.4}$$

All these costs have depth-two implementations.

Proof. Theorem 2.4 gives $C_g(M) \geq \lceil 3n/2 \rceil$ for every matrix in the class. It remains to prove that the displayed constructions are admissible and attain this number of gates.

First let $n = 2k + 1 \geq 5$ and $M = I - P + P^2$. The identity

$$(I + P)M = I + P^3$$

proves invertibility as follows. If $Mx = 0$, then $x_{i+3} = -x_i$ for every i. Every step-three cycle has length $n/\gcd(n, 3)$, which is odd. Following it back to the starting coordinate gives

$x_i = -x_i$, hence $x_i = 0$ because p is odd. Thus M has zero kernel. Its support $\{0, 1, 2\}$ is not a subgroup coset of order three for $n \geq 5$, so Proposition 4.2 gives $B(M) = 4$.

For an explicit implementation of $y_i = x_i - x_{i+1} + x_{i+2}$, compute

$$t_j = x_{2j+1} - x_{2j+2}, \qquad 0 \leq j < k,$$

and then

$$y_{2j} = x_{2j} - t_j, \qquad y_{2j+1} = t_j + x_{2j+3}, \qquad 0 \leq j < k.$$

All indices are modulo n. These k shared differences and $2k$ output gates compute the first $2k$ outputs. The remaining output $y_{2k} = x_{2k} - x_0 + x_1$ uses two gates. The total is $k + 2k + 2 = (3n + 1)/2$. Each shared difference has depth one, and all outputs can be completed at depth two.

Now let $n = 2k \geq 4$ and $M = I + P + \sigma P^k$. Since p is odd, 1 and -1 are distinct, so at least one choice of σ satisfies (4.4). To prove invertibility, put $Q = P^k$, so $Q^2 = I$. Every vector decomposes as

$$x = x_+ + x_-, \qquad x_+ = \frac{x + \sigma Q x}{2}, \qquad x_- = \frac{x - \sigma Q x}{2},$$

with $Qx_+ = \sigma x_+$ and $Qx_- = -\sigma x_-$. The matrices M and Q commute, so if $Mx = 0$, then both $Mx_+ = 0$ and $Mx_- = 0$. On these two parts, respectively,

$$Mx_+ = (2I + P)x_+, \qquad Mx_- = Px_-.$$

The second equation gives $x_- = 0$. The first gives $Px_+ = -2x_+$, and applying P a total of k times yields

$$Qx_+ = (-2)^k x_+.$$

But $Qx_+ = \sigma x_+$, so the chosen inequality forces $x_+ = 0$. This proves invertibility, also when p divides n. The support $\{0, 1, k\}$ is not a subgroup coset of order three, so again $B(M) = 4$.

Compute the k signed opposite pairs

$$t_i = x_i + \sigma x_{i+k}, \qquad 0 \leq i < k.$$

The two corresponding outputs are

$$y_i = t_i + x_{i+1}, \qquad y_{i+k} = \sigma t_i + x_{i+k+1}.$$

Negation is free and $\sigma^2 = 1$, so the implementation uses $k + n = 3n/2$ gates and has depth two. In both constructions, each shared pair is used by two outputs and defines an admissible rectangle, so these are rectangle implementations. Both supports contain $0, 1$, which proves connectivity. Together with the lower bound, this establishes (4.3). □

4.1.3 General circuits beyond rectangle reuse: constructions and depth trade-offs

The equality of the two minima does not imply $C_g(M) = C_{\text{rect}}(M)$ for every fixed matrix. We first identify the restriction imposed by depth two when $p \geq 5$, then show why characteristic three gives an exception.

Proposition 4.5. *Let $p \geq 5$. For every invertible signed circulant M of weight three and branch number four, the minimum gate count among circuits of depth at most two equals $C_{\text{rect}}(M)$. Thus an implementation using fewer gates than rectangle reuse must have depth at least three.*

The proof is given in Appendix C.2.

Over $\mathbb{F}_3$, the connected, invertible matrix

$$M = I + P - P^2 = \begin{pmatrix} 1 & 1 & -1 & 0 & 0 \\ 0 & 1 & 1 & -1 & 0 \\ 0 & 0 & 1 & 1 & -1 \\ -1 & 0 & 0 & 1 & 1 \\ 1 & -1 & 0 & 0 & 1 \end{pmatrix}$$

has branch number four and satisfies $C_g(M) \leq 9 < 10 = C_{\text{rect}}(M)$. Compute $t_1 = x_1 - x_2$ and $t_2 = x_2 + x_3$, then use $y_0 = x_0 + t_1$, $y_1 = t_1 - t_2$, and $y_2 = t_2 - x_4$; the other two outputs use two gates each, giving nine gates in total at depth two. This works because $\pm 2 = \mp 1$ in $\mathbb{F}_3$: the repeated input in $t_1 - t_2 = x_1 - 2x_2 - x_3 = x_1 + x_2 - x_3$ has an allowed signed coefficient. Thus two overlapping input pairs can produce a valid weight-three output, which is excluded when $p \geq 5$.

Thus the constructions in Theorem 4.4 attain both the minimum addition/subtraction count and the minimum circuit depth of two. For the large odd primes p commonly used in AO algorithms, rectangle reuse gives a circuit with minimum depth for any fixed matrix in this class of weight-three circulants with branch number four.

4.1.4 Constructions at representative orders

Table 5 lists optimal constructions over $\mathbb{F}_p$ at different orders for every odd prime p. All displayed matrices are connected and invertible, and their implementations have depth two.

Table 5: Optimal weight-three circulants with branch number four over every odd prime field. For even n, choose $\sigma \in \{1, -1\}$ according to (4.4).

n	Attaining matrix	$\min C_g = \min C_{\text{rect}}$
4	$I + P + \sigma P^2$	6
5	$I - P + P^2$	8
6	$I + P + \sigma P^3$	9
7	$I - P + P^2$	11
8	$I + P + \sigma P^4$	12
9	$I - P + P^2$	14
10	$I + P + \sigma P^5$	15
11	$I - P + P^2$	17
12	$I + P + \sigma P^6$	18
13	$I - P + P^2$	20
14	$I + P + \sigma P^7$	21
15	$I - P + P^2$	23
16	$I + P + \sigma P^8$	24
32	$I + P + \sigma P^{16}$	48
64	$I + P + \sigma P^{32}$	96

4.2 Branch number five and weight four

At weight four, a weight-one input gives an output of weight four, so $B(M) \leq 5$. We determine when equality holds and prove the rectangle cost lower bound $2n$. Two connected infinite families attain this rectangle bound over every odd prime field. Throughout, M is an invertible signed circulant of order n and weight four; connectivity is imposed only where stated.

4.2.1 Branch number criterion and rectangle cost lower bound

Use the signed overlap counts $c_+(d), c_-(d)$ of Section 2.1, with $c(d) = c_+(d) + c_-(d)$. These count the common nonzero positions of two shifted columns, distinguishing equal and opposite coefficients.

Proposition 4.6. *The matrix M has branch number five if and only if*

$$c(d) + \max\{c_+(d), c_-(d)\} \le 5 \quad (d \neq 0), \qquad \mathrm{wt}\big((M^{-1})_{0,*}\big) \ge 4. \tag{4.5}$$

Proof. Two columns at displacement d have $8 - c(d)$ positions in their combined support. In a linear combination with two nonzero input coefficients, equal input coefficients cancel the $c_-(d)$ positions with opposite column coefficients. Opposite input coefficients cancel the $c_+(d)$ positions with equal column coefficients. Every other coefficient ratio cancels no common position. Thus the smallest possible output weight for these two columns is $8 - c(d) - \max\{c_+(d), c_-(d)\}$. Requiring it to be at least three gives the first condition.

For an input of weight three, the only remaining violation is an output of weight one. Such an input would be a nonzero scalar multiple of a column of M^{-1}. All inverse columns have the same weight, equal to its row weight because the inverse is circulant. The second condition excludes this possibility. Inputs of weight at least four have nonzero outputs by invertibility, giving total weight at least five. Weight-one inputs give total weight exactly five. Conversely, $B(M) = 5$ forces the two-column condition, and $B(M^{-1}) = B(M)$ forces every inverse column to have weight at least four. □

The next criterion avoids computing the inverse for the low-overlap supports used below. For two columns sharing exactly two positions, call their coefficient ratios equal if the same ratio occurs at both positions, and opposite otherwise. The ratios belong to $\{1, -1\}$.

Lemma 4.7. *Suppose any two columns of M share at most two positions. Then $B(M) = 5$ if and only if no three columns form the following configuration. Every pair shares exactly two positions, and the triple intersection is empty. Exactly two column pairs have equal coefficient ratios. The remaining pair has opposite coefficient ratios.*

The proof is given in Appendix C.3.

Theorem 4.8. *If $B(M) = 5$, then*

$$C_{\mathrm{rect}}(M) \ge 2n. \tag{4.6}$$

Equality holds if and only if all $4n$ nonzero positions can be partitioned into n admissible 2×2 rectangles. In such a partition, each row and each column belongs to exactly two rectangles.

Proof. Proposition 2.6 restricts admissible rectangles to size 2×2. Each saves one gate from the direct cost $3n$. Since a row has four nonzero positions, it participates in at most two rectangles whose matrix positions are disjoint. Counting the two rows of every selected rectangle therefore bounds their number by n. The total saving is at most n, giving $C_{\mathrm{rect}}(M) \ge 2n$.

Equality requires n rectangles. Their $4n$ positions are disjoint, so they cover all nonzero positions. Each row and column is then covered twice. Conversely, such a partition saves n gates and attains the lower bound. □

4.2.2 Attaining constructions and the effect of connectivity

Theorem 4.9. *For every odd prime p, the following matrices are connected, invertible, have branch number five, and satisfy $C_{\mathrm{rect}}(M) = 2n$:*

1. *For every odd* $n \geq 9$*,*

$$M = (I + P)(I + P^3) = I + P + P^3 + P^4.$$

2. *For* $n \equiv 4 \pmod 8$ *and* $n \geq 12$*,*

$$M = I + P + P^2 - P^{n/2+1}.$$

Both have implementations with $2n$ *gates and depth two.*

Proof. For the first family, all cycles of P and P^3 have odd length. On a cycle of length q, the kernel equation for either factor makes successive coordinates negatives of each other. After q steps, $x_i = -x_i$, so $x_i = 0$ in odd characteristic. Thus $I + P$ and $I + P^3$ are both invertible, and so is their product.

For $n \geq 9$, the support $\{0, 1, 3, 4\}$ has overlap two only at displacements $\pm 1, \pm 3$; all other nonzero overlaps are at most one. All coefficients are positive, so every column pair with two common positions has equal ratios. The configuration excluded by Lemma 4.7 cannot occur, and hence $B(M) = 5$. Compute

$$t_i = x_i + x_{i+1}, \qquad y_i = t_i + t_{i+3}, \qquad 0 \leq i < n.$$

The first layer has n gates and the second has n gates. Each shared pair t_i is used by outputs i and $i - 3$, and each output combines pairs on four distinct input positions. These pairs give a complete rectangle partition. Thus $C_{\text{rect}}(M) = 2n$, with depth two.

For the second family, write $n = 2k$, where k is even and $k/2$ is odd. Put $Q = P^k$, so $Q^2 = I$. Use $x_+ = (x + Qx)/2$ and $x_- = (x - Qx)/2$, with $Qx_+ = x_+$ and $Qx_- = -x_-$, as in the proof of Theorem 4.4. If $Mx = 0$, commutation gives $Mx_\pm = (Mx \pm QMx)/2 = 0$. On these two parts,

$$Mx_+ = (I + P^2)x_+, \qquad Mx_- = (I + P)^2 x_-.$$

The first equation implies $P^2 x_+ = -x_+$. Repeating $k/2$ times gives $Qx_+ = (-1)^{k/2} x_+ = -x_+$, so $x_+ = 0$. On the part where $Qx = -x$, the map $I + P$ has zero kernel: $Px = -x$ would imply $Qx = (-1)^k x = x$, forcing $x = 0$. Since $I + P$ preserves this part, its square also has zero kernel there. Hence $x_- = 0$, proving invertibility.

The support is $\{0, 1, 2, k+1\}$, with only the last coefficient negative. Overlap two occurs at displacements $\pm 1, k - 1, k, k + 1$. At ± 1 both coefficient ratios are 1; at the other three displacements both are -1. All remaining nonzero overlaps are at most one. Again there is no pair with opposite ratios, so the low-overlap lemma gives $B(M) = 5$.

For every even index i, compute the two depth-one signals

$$t_{i,1} = x_i + x_{i+1}, \qquad t_{i,2} = x_{i+2} - x_{i+k+1}.$$

There are n such signals in total. The outputs are

$$y_i = \begin{cases} t_{i,1} + t_{i,2}, & i \text{ even}, \\ t_{i+1,1} - t_{i-k-1,2}, & i \text{ odd}. \end{cases}$$

All indices are modulo n. For an odd i, expansion of the second expression gives $x_{i+1} + x_{i+2} - x_{i-k+1} + x_i$; since $n = 2k$, $i - k + 1 \equiv i + k + 1$, as required. Each signal is used by two outputs, and each output combines two disjoint signed pairs. This is a rectangle implementation with $n + n = 2n$ gates and depth two, so it attains the bound. Both families have support positions $0, 1$, proving connectivity. □

Without a connectivity requirement, the two attaining families above can be extended to attain the bound at more orders.

Proposition 4.10. *Let $M = f(P_m)$ be an order-m matrix with weight four, branch number five, and $C_{\mathrm{rect}}(M) = 2m$. For any positive integer ℓ, put $n = \ell m$. Then $f(P_n^\ell)$ has the same weight and branch number and satisfies $C_{\mathrm{rect}} = 2n$. A depth-two implementation of M lifts to depth two. If M is connected, the lifted matrix has exactly ℓ connected components.*

Proof. Reorder the coordinates by their residues modulo ℓ. The lifted matrix becomes ℓ identical copies of M. Invertibility and branch number are preserved, and (2.11) gives rectangle cost $\ell \cdot 2m = 2n$. Implementing the copies in parallel preserves depth. The component count follows from the same decomposition. □

Writing $n = 2^a r$ with r odd, these constructions attain $2n$ whenever $r \geq 9$, or whenever $a \geq 2$ and $r \geq 3$. For the first range, lift the odd-order construction at order r by 2^a. For the second, lift the construction at order $4r$ by 2^{a-2}. The extended matrices are generally disconnected.

4.2.3 General circuits beyond rectangle reuse: constructions and depth trade-offs

For every matrix in this class, Theorem 2.4 gives the general circuit lower bound

$$C_g(M) \geq 2n - \left\lfloor \frac{n}{5} \right\rfloor . \tag{4.7}$$

For the attaining constructions above, this yields $2n - \lfloor n/5 \rfloor \leq C_g(M) \leq 2n$. Their exact rectangle cost does not determine their unrestricted general circuit cost. With a depth-two restriction, however, the two costs coincide over every odd prime field.

Proposition 4.11. *For an invertible signed circulant of weight four and branch number five over any odd prime field, the minimum gate count among circuits of depth at most two equals $C_{\mathrm{rect}}(M)$. Consequently, an implementation with fewer gates has depth at least three. Every rectangle cost is attained at the minimum possible depth of two.*

Proof. Each output contains four distinct inputs and needs a final gate at depth two. Each operand has support size at most two, so the two operands must each contain two inputs and have disjoint supports. Thus every output combines two signed input pairs. In particular, overlapping pairs cannot produce a weight-four output, even in characteristic three.

Remove unused gates and duplicate pair computations. Any input pair shared by two or more outputs gives an admissible rectangle. The two pairs within each output have disjoint positions, so the shared rectangles form a valid rectangle scheme. This bounds the cost below by $C_{\mathrm{rect}}(M)$. Conversely, all rectangles have size 2×2. In each output, group any remaining original inputs into private pairs. Compute all shared and private pairs at depth one and combine the two pairs of each output at depth two, with the same gate count as the rectangle scheme. This proves equality. A depth-one signal cannot contain four distinct inputs, so depth two is minimal. □

Therefore, when the circuit depth is at most two, rectangle reuse is an optimal circuit implementation method in our setting.

4.2.4 Finite optimal costs and attainability

If the cost $2n$ in Theorem 4.8 is attained, the n selected rectangles partition all nonzero positions of the matrix. This partition gives an integer factorization as follows. Associate one shared signal with each rectangle. Put the two signed input coefficients of each signal in a separate row of V. In U, use the corresponding column to record the two signed uses of that signal in the outputs. The shared signals are Vx, and the outputs are $U(Vx)$, so

$$M = UV \quad \text{over the integers.} \tag{4.8}$$

Each signal uses two inputs and contributes to two outputs. Each output combines two signals, and each input belongs to two rectangles. Thus U, V are $n \times n$ signed matrices with exactly two nonzero entries in every row and column. Because the rectangles have disjoint matrix positions, each matrix entry receives a contribution from at most one shared signal. There is no cancellation or addition of contributions from different signals. The factorization therefore holds over the integers. Conversely, a factorization with these properties gives a complete admissible rectangle partition.

Proposition 4.12. *An integer signed matrix with a complete rectangle partition as above has determinant zero or a signed power of two. Consequently, its invertibility is the same over all odd prime fields. In particular, if it is singular over $\mathbb{F}_3$, it is singular over every odd prime field.*

Proof. In each factor of (4.8), the support graph has degree two at every row and column vertex, so it splits into cycles. Each cycle contains equally many row and column vertices and gives a square block after permutations. These cycle blocks form a block diagonal matrix for each of U and V. The determinant expansion of each cycle block has exactly two nonzero terms, each 1 or -1. It is therefore zero or ± 2. Multiplying the block determinants, and then the determinants of U and V, proves the stated form. No odd prime divides a nonzero power of two. □

Proposition 4.12 gives a quick test for excluding attainment of $2n$ at a given order. Enumerate signed generators satisfying the condition on two columns in (4.5), and retain those admitting a complete rectangle partition. For each retained candidate, it suffices to test invertibility over $\mathbb{F}_3$. If all are singular there, none can attain $2n$ over any odd prime field. For a candidate that passes the rank test, the remaining condition for branch number five is checked by Proposition 4.6.

A support shift and an overall sign change place a coefficient $+1$ at position zero, leaving $8\binom{n-1}{3}$ normalized signed generators. Complete enumeration excludes attainment for $n \in \{4, 5, 6, 7, 8, 10, 14, 16, 32, 64\}$: every candidate satisfying the condition on two columns and admitting a complete rectangle partition is singular over $\mathbb{F}_3$. At orders 32 and 64, the numbers of such normalized candidates passing the condition on two columns and admitting a complete rectangle partition are 1,680 and 7,440, respectively. Hence $C_{\text{rect}} \geq 2n + 1$ at all these orders whenever a matrix in the prescribed class exists. For $n = 9, 11, 12, 13, 15$, the constructions in Theorem 4.9 attain $2n$ over every odd prime field.

Table 6 summarizes these decisions and the exact minimum rectangle costs over $\mathbb{F}_3$ and $\mathbb{F}_5$, obtained from the completed finite searches and their optimality bounds. The minima in the table impose no connectivity requirement. Orders below four cannot accommodate weight four.

Connectivity changes some of the finite optima. Over $\mathbb{F}_3$, requiring connectivity raises the minima at orders 32 and 64 to 77 and 154, respectively. Over each of $\mathbb{F}_5, \mathbb{F}_7, \mathbb{F}_{11}$, the minima at these two orders remain 72 and 144 even when connectivity is required. These exact values are asserted only for the stated fields.

Table 6: Attainability of $2n$ over odd prime fields and exact minimum rectangle costs over $\mathbb{F}_3, \mathbb{F}_5$ for weight four and branch number five. *Yes* means attainment over every odd prime field; *No* means that no odd prime field admits attainment. A dash indicates that no invertible matrix of weight four and branch number five exists in the specified field.

n	$2n$	$2n$ attainable	$\min C_{\text{rect}}$ over $\mathbb{F}_3$	$\min C_{\text{rect}}$ over $\mathbb{F}_5$
4	8	No	—	—
5	10	No	—	12
6	12	No	14	14
7	14	No	16	16
8	16	No	18	18
9	18	Yes	18	18
10	20	No	22	22
11	22	Yes	22	22
12	24	Yes	24	24
13	26	Yes	26	26
14	28	No	30	30
15	30	Yes	30	30
16	32	No	36	36
32	64	No	72	72
64	128	No	144	144

4.3 Branch number six and weight five

We now consider invertible signed circulants of weight five with branch number six. We derive rectangle cost lower bounds and report the exact finite minima already determined. Connectivity is required only where stated.

4.3.1 Rectangle cost lower bounds

Proposition 2.6 allows only rectangles of sizes 2×2, 2×3, 3×2, and 3×3. Substituting weight five and branch number six into Theorem 2.9 gives $C_{\text{rect}}(M) \geq \lceil 20n/9 \rceil$. The stronger bound below uses the circulant structure. We first handle matrices containing an admissible 3×3 rectangle.

Lemma 4.13. *Suppose M is connected, $n \geq 15$, and its support contains an admissible 3×3 rectangle. Then*

$$C_{\text{rect}}(M) \geq \left\lceil \frac{22n}{9} \right\rceil.$$

The proof follows the core block counting argument in Section 3.2. The additional steps establish equal core coefficients and exclude the exceptional overlaps over odd prime fields, under the stated connectivity and order assumptions. The full proof is given in Appendix C.

Theorem 4.14. *Every connected invertible signed circulant of weight five and branch number six, of order $n \geq 15$ over an odd prime field, satisfies*

$$C_{\text{rect}}(M) \geq \left\lceil \frac{12n}{5} \right\rceil. \tag{4.9}$$

If no admissible 3×3 rectangle exists, this bound holds without a connectivity requirement or the restriction $n \geq 15$.

Proof. First suppose no admissible 3×3 rectangle exists. The remaining sizes 2×2, 2×3, and 3×2 save one, two, and two gates, respectively. Their contributions to the rank budget are three, five, and five. Hence each rectangle saves at most $2/5$ of its budget contribution.

Lemma 2.7 bounds the total budget by $4n$, so the total saving is at most $8n/5$. Therefore $C_{\mathrm{rect}}(M) \geq 4n - 8n/5 = 12n/5$.

If an admissible 3×3 rectangle exists, the preceding lemma gives $C_{\mathrm{rect}}(M) \geq 22n/9 > 12n/5$. Taking ceilings proves the result. □

For general circuits, Theorem 2.4 gives $C_g(M) \geq 2n - \lfloor n/6 \rfloor$.

4.3.2 Finite optimal rectangle costs

For $p \in \{3, 5, 7, 11\}$, the finite optimizations are complete for all orders $n \leq 16$ and for $n = 32, 64$. No matrix of weight five and branch number six exists at orders $n \leq 6$ for these primes. Table 7 gives the exact minimum rectangle costs at the feasible orders considered. At each listed order and prime, the minimum with a connectivity requirement equals the minimum without that requirement. All feasible minima have connected implementations of depth three, which is the minimum possible depth for an output containing five inputs.

At orders 32 and 64, complete enumeration can use a smaller candidate set. A cost below $3n$ requires a 2×3 or 3×2 rectangle, since 2×2 rectangles alone save at most n gates. By transposition equivalence, consider two rows with three common entries of the same coefficient ratio. The branch condition forces their total overlap to be exactly three. On the five support positions, the three steps of their row displacement cannot form a cycle at a power-of-two order: a step of order two gives an even number of edges, and a longer cycle would need at least four edges. They therefore form two paths with four and one vertices, or three and two vertices. After a support shift and an overall sign change, the corresponding generators are

$$1 + \sigma_1 X^d + X^{2d} + \sigma_1 X^{3d} + \sigma_2 X^e,$$
$$1 + \sigma_1 X^d + X^{2d} + \sigma_2 X^e + \sigma_1 \sigma_2 X^{e+d},$$

where $\sigma_1, \sigma_2 \in \{1, -1\}$ and all five exponents are distinct modulo n. Enumerating these forms is sufficient: every omitted matrix costs at least $3n$, above the minima in the table. Bounds for disconnected candidates also follow by adding the already determined costs of their smaller components.

Table 7: Exact minimum C_{rect} for signed circulants of weight five and branch number six. The values are the same with and without a connectivity requirement. A dash indicates that no matrix in this class exists.

n	$\mathbb{F}_3$	$\mathbb{F}_5$	$\mathbb{F}_7$	$\mathbb{F}_{11}$
7	—	20	20	20
8	24	24	22	22
9	22	24	22	22
10	27	28	27	27
11	30	30	30	30
12	33	32	31	31
13	35	35	35	35
14	39	39	38	38
15	40	40	40	40
16	44	44	44	44
32	88	88	86	86
64	176	176	171	171

5 Search and Verification Methods

For each feasible candidate, we enumerate every admissible rectangle within the dimension bounds of Proposition 2.6, including smaller rectangles contained in larger ones. Each rectangle has a binary selection variable, and each nonzero matrix position has a constraint allowing it to occur in at most one selected rectangle. Maximizing the sum of the savings $(r_Q - 1)(s_Q - 1)$ gives the rectangle packing problem in Proposition 2.2. We solve these integer programs using Gurobi and HiGHS. The related code is available at `https://github.com/williamarendt/circulant-diffusion-layers`.

In several of the experiments above, this code can also obtain search or optimization results at larger orders n within a short time. We report only a limited selection of representative orders to support the theoretical analysis. For practical linear layer design, the code can also guide matrix selection and construction at larger orders.

6 Conclusions

This paper aims to provide insights into linear layer design for block ciphers and hash functions. We focus on circulant matrices, a widely used matrix class, and study how to balance diffusion and implementation cost. We measure diffusion by branch numbers and cost by the number of two-input addition or subtraction gates. Our main objects are invertible circulants with entries in $\{0, 1\}$ over $\mathbb{F}_2$ or $\{0, 1, -1\}$ over odd prime fields, focusing on branch numbers four, five, and six at minimum weight. We consider both the general circuit cost C_g and the rectangle reuse cost C_{rect}.

We establish general lower bounds on C_g and C_{rect}. For binary circulants, we establish cost bounds under branch number requirements four, five, and six, study how connectivity requirements and circuit depth affect attainable costs, and give construction families and finite search results at practical orders. We also study the corresponding problems over odd prime fields, giving cost bounds, attaining families, and finite search results, together with the effects of connectivity, circuit depth, and field characteristic.

For many of the parameter settings studied, C_{rect} and C_g are close, while rectangle implementations offer lower circuit depth. Rectangle reuse therefore provides a useful reference for practical linear layer design. The gate count lower bounds established in this paper also provide theoretical bounds for matrix selection and circuit optimization, and a reference for estimating implementation costs.

References

[1] J. Cho, J. Ha, S. Kim, B. Lee, J. Lee, J. Lee, D. Moon, and H. Yoon. Transciphering framework for approximate homomorphic encryption. In *Advances in Cryptology – ASIACRYPT 2021, Part III*, pages 640–669, 2021. `https://eprint.iacr.org/2020/1335`.

[2] Tingting Cui, Xi Han, Yan He, Congkai Zhang, Siqin Yu, and Wei Zhao. An STP-based model toward designing binary diffusion layers with the lowest latency. *Cybersecurity*, 9:184, 2026. Article 184. `https://doi.org/10.1186/s42400-026-00617-z`.

[3] Joan Daemen and Vincent Rijmen. The wide trail design strategy. In Bahram Honary, editor, *Cryptography and Coding*, volume 2260 of *Lecture Notes in Computer Science*, pages 222–238. Springer, 2001. `https://doi.org/10.1007/3-540-45325-3_20`.

[4] S. M. Dehnavi, A. Mahmoodi Rishakani, and M. R. Mirzaee Shamsabad. Bitwise linear mappings with good cryptographic properties and efficient implementation. Cryptology ePrint Archive, Report 2015/225, 2015. https://eprint.iacr.org/2015/225.

[5] Christoph Dobraunig, Lorenzo Grassi, Lukas Helminger, Christian Rechberger, Markus Schofnegger, and Roman Walch. Pasta: A case for hybrid homomorphic encryption. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2023(3):30–73, 2023. https://doi.org/10.46586/tches.v2023.i3.30-73.

[6] S. Duval and G. Leurent. MDS matrices with lightweight circuits. *IACR Transactions on Symmetric Cryptology*, 2018(2):48–78, 2018. https://doi.org/10.13154/tosc.v2018.i2.48-78.

[7] A. Gál, K. A. Hansen, M. Koucký, P. Pudlák, and E. Viola. Tight bounds on computing error-correcting codes by bounded-depth circuits with arbitrary gates. *IEEE Transactions on Information Theory*, 59(10):6611–6627, 2013. https://doi.org/10.1109/TIT.2013.2270275.

[8] Lorenzo Grassi, Dmitry Khovratovich, Christian Rechberger, Arnab Roy, and Markus Schofnegger. Poseidon: A new hash function for zero-knowledge proof systems. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 519–535. USENIX Association, 2021. https://www.usenix.org/system/files/sec21-grassi.pdf.

[9] Lorenzo Grassi, Dmitry Khovratovich, and Markus Schofnegger. Poseidon2: A faster version of the Poseidon hash function. In *Progress in Cryptology – AFRICACRYPT 2023*, 2023. https://eprint.iacr.org/2023/323.

[10] Z. Guo, R. Liu, S. Gao, W. Wu, and D. Lin. Direct construction of optimal Rotational-XOR diffusion primitives. *IACR Transactions on Symmetric Cryptology*, 2017(4):169–187, 2017. https://doi.org/10.13154/tosc.v2017.i4.169-187.

[11] Anup Hosangadi, Farzan Fallah, and Ryan Kastner. Algebraic methods for optimizing constant multiplications in linear systems. *Journal of VLSI Signal Processing*, 49:31–50, 2007. https://doi.org/10.1007/s11265-007-0137-7.

[12] S. Jukna and I. Sergeev. Complexity of linear Boolean operators. *Foundations and Trends in Theoretical Computer Science*, 9(1):1–123, 2013. https://doi.org/10.1561/0400000063.

[13] Gregor Leander, Thorben Moos, Amir Moradi, and Shahram Rasoolzadeh. The SPEEDY family of block ciphers: Engineering an ultra low-latency cipher from gate level for secure processor architectures. Cryptology ePrint Archive, Paper 2021/960, 2021. https://eprint.iacr.org/2021/960.

[14] Fen Liu, Yongqiang Li, Huiqin Chen, Lin Jiao, Ming Luo, and Mingsheng Wang. YuX: Finite field multiplication based block ciphers for efficient FHE evaluation. *IEEE Transactions on Information Theory*, 70(5):3729–3749, 2024. https://doi.org/10.1109/TIT.2024.3349414.

[15] A. Mahmoodi Rishakani, S. M. Dehnavi, M. R. Mirzaee Shamsabad, and N. Bagheri. Cryptographic properties of cyclic binary matrices. *Advances in Mathematics of Communications*, 15(2):311–327, 2021. https://doi.org/10.3934/amc.2020068.

[16] National Institute of Standards and Technology. Post-quantum cryptography: Frequently asked questions. NIST Computer Security Resource Center. Accessed November 10, 2026. https://csrc.nist.gov/Projects/Post-Quantum-Cryptography/faqs.

[17] National Institute of Standards and Technology. SHA-3 standard: Permutation-based hash and extendable-output functions. Technical Report FIPS 202, National Institute of Standards and Technology, 2015. https://doi.org/10.6028/NIST.FIPS.202.

[18] National Institute of Standards and Technology. Advanced encryption standard (AES). Technical Report FIPS 197-upd1, National Institute of Standards and Technology, 2023. https://doi.org/10.6028/NIST.FIPS.197-upd1.

[19] R. Paturi and P. Pudlák. Circuit lower bounds and linear codes. *Journal of Mathematical Sciences*, 134(5):2425–2434, 2006. Preprint: ECCC TR04-004, 2004. https://eccc.weizmann.ac.il/report/2004/004/.

[20] Alan Szepieniec, Tomer Ashur, and Siemen Dhooghe. Rescue-Prime: a standard specification (SoK). Cryptology ePrint Archive, Paper 2020/1143, 2020. https://eprint.iacr.org/2020/1143.

[21] Meltem Sönmez Turan, Kerry A. McKay, Donghoon Chang, Jinkeon Kang, and John Kelsey. Ascon-based lightweight cryptography standards for constrained devices: Authenticated encryption, hash, and extendable output functions. Technical Report SP 800-232, National Institute of Standards and Technology, 2025. https://doi.org/10.6028/NIST.SP.800-232.

[22] A. Venkateswarlu, A. Kesarwani, and S. Sarkar. On the lower bound of cost of MDS matrices. *IACR Transactions on Symmetric Cryptology*, 2022(4):266–290, 2022. https://doi.org/10.46586/tosc.v2022.i4.266-290.

[23] T. Weng and G. Wang. New SAT-based model for constructing linear layers with good cryptographic properties and implementation. *Frontiers of Computer Science*, 20(10):2010812, 2026. https://doi.org/10.1007/s11704-025-40011-5.

[24] Baofeng Wu, Wen Kong, Dewei Kong, and Hailun Yan. MDS diffusion layers for arithmetization-oriented symmetric ciphers: The Rotational-Add construction. *IACR Transactions on Symmetric Cryptology*, 2025(3):827–867, 2025. https://doi.org/10.46586/tosc.v2025.i3.827-867.

[25] Z. Xiang, X. Zeng, D. Lin, Z. Bao, and S. Zhang. Optimizing implementations of linear layers. *IACR Transactions on Symmetric Cryptology*, 2020(2):120–145, 2020. https://eprint.iacr.org/2020/903.

[26] Buji Xu and Xiaoming Sun. Utilizing circulant structure to optimize the implementations of linear layers. arXiv:2511.18226, 2025. Version 2, revised May 24, 2026. https://arxiv.org/abs/2511.18226.

A Proofs for Section 2

A.1 Proposition 2.1 in Section 2.1

For the circulant $M = f(P)$, the following hold:

1. $B_d(M) = B_l(M)$.
2. *M is invertible if and only if* $\gcd(f(X), X^n - 1) = 1$ *in* $\mathbb{F}[X]$.

Neither assertion requires the characteristic of $\mathbb{F}$ to be coprime to n.

Proof. Let J be the coordinate reversal $(Jx)_i = x_{-i}$. Then $JPJ = P^{-1} = P^{\mathsf{T}}$, hence $JMJ = M^{\mathsf{T}}$. Independent permutations of input and output coordinates preserve Hamming weight and the differential branch number, proving the first assertion.

The algebra generated by P is isomorphic to $\mathbb{F}[X]/(X^n - 1)$ under $X \mapsto P$: the matrices $I, P, \ldots, P^{n-1}$ are linearly independent and $P^n = I$. If the gcd is one, a Bézout identity evaluated at P gives an inverse of $f(P)$. Otherwise, put $h = \gcd(f, X^n - 1)$ and $q = (X^n - 1)/h$. Then $q(P) \neq 0$, since q is nonzero of degree less than n, whereas $f(P)q(P) = 0$. Thus $f(P)$ is singular. This argument also covers repeated factors of $X^n - 1$. □

A.2 Theorem 2.4 in Section 2.3.1

Let $M \in \mathbb{F}^{n \times n}$ be invertible. Then

$$B_l(M) \geq 4 \quad \Longrightarrow \quad C_g(M) \geq 2n - \left\lfloor \frac{2n}{B_d(M)} \right\rfloor,$$
$$B_l(M) \geq 5 \quad \Longrightarrow \quad C_g(M) \geq 2n - \left\lfloor \frac{n}{B_d(M)} \right\rfloor.$$

Proof. For this proof, allow arbitrary scalar multiplications for free. A lower bound in this larger circuit class also applies to the addition model of Section 2.2. Choose a minimum circuit in the larger class, with m addition gates, and remove unused signals. Each gate equation defines a new signal using only inputs and earlier gate signals. In computation order, the coefficients of the newly defined signals form a triangular matrix with nonzero diagonal. Thus the gate equations are linearly independent.

Call the n inputs and n outputs the terminals. Their possible values form the linear space

$$\{(x, Mx) : x \in \mathbb{F}^n\}.$$

Its minimum nonzero Hamming weight is $B_d(M)$. A linear relation among the terminals must hold for every input x. It has the form $u^{\mathsf{T}}x + v^{\mathsf{T}}Mx = 0$, with $u = -M^{\mathsf{T}}v$. The number of nonzero coefficients in this relation is $\mathrm{wt}(M^{\mathsf{T}}v) + \mathrm{wt}(v)$, whose minimum over $v \neq 0$ is $B_l(M)$. In particular, $B_l(M) \geq 4$ excludes every nonzero linear relation involving at most three terminal coordinates.

Since $B_l(M) \geq 4$, no output is proportional to an input or to another output. Thus the outputs occupy n distinct gate signals. Write $h = m - n$ for the number of remaining gate signals, which we call internal signals. Each internal signal is a linear function of the input. Requiring all h of them to be zero imposes at most h independent linear constraints on x. The inputs satisfying these constraints therefore form a subspace of dimension at least $n - h$.

Each gate equation involves at most three signals. An equation involving only terminals would give a relation supported on at most three terminals, which is impossible. Substitute zero for the internal signals in all gate equations. Each remaining nontrivial equation either sets one terminal to zero or makes two terminals proportional. Their full solution space

consists exactly of the pairs (x, Mx) whose internal signals are zero. Since a pair (x, Mx) uniquely determines x, this terminal solution space has the same dimension as the restricted input space above.

Form a graph on the terminals using the two-terminal equations as edges. Within a connected component, the proportionality relations determine all coordinates from any one of them. The component is either forced to zero or has one free parameter. Let t be the number of free components. Thus t is the dimension of the terminal solution space, giving $t \geq n - h$.

To obtain a basis vector, set the free parameter of one component to one and all other free parameters to zero. The resulting vector still satisfies every gate equation and therefore belongs to $\{(x, Mx) : x \in \mathbb{F}^n\}$. Its nonzero coordinates are exactly those of the chosen component, so that component contains at least $B_d(M)$ terminals. The t free components have disjoint supports within the $2n$ terminal coordinates. Consequently,

$$n - h \leq t, \qquad B_d(M)t \leq 2n.$$

It follows that $h \geq n - \lfloor 2n/B_d(M) \rfloor$, proving (2.13).

Now suppose $B_l(M) \geq 5$. Each graph edge comes from a gate equation containing two terminals and one internal signal. Before substitution, such an equation has the form

$$a\xi + b\eta + cz = 0,$$

where ξ, η are terminal coordinates, z is an internal signal, and a, b, c are nonzero. Two such equations cannot contain the same internal signal: eliminating it would give a relation on at most four terminals. This relation cannot vanish identically, since that would make the two gate equations proportional, contrary to their independence. A nonzero relation of this size contradicts $B_l(M) \geq 5$. Hence there are at most h edge equations.

Each free component contains at least $B_d(M)$ terminals and needs at least $B_d(M) - 1$ edges to be connected. Distinct components use disjoint sets of edge equations, so at least $(B_d(M) - 1)t$ such equations are needed in total. Thus

$$h \geq (B_d(M) - 1)t \geq (B_d(M) - 1)(n - h).$$

Rearranging and using integrality gives $h \geq n - \lfloor n/B_d(M) \rfloor$, proving (2.14). □

A.3 Theorem 2.9 in Section 2.3.3

Let $M \in \{0, 1, -1\}^{n \times n}$ be invertible, with row and column weight $w = b - 1$ and $B_d(M), B_l(M) \geq b$. For $R \geq 2$,

$$b = 2R + 1 : \qquad C_{\text{rect}}(M) \geq \left\lceil \frac{2(2R-1)}{R+1}\, n \right\rceil,$$
$$b = 2R : \qquad C_{\text{rect}}(M) \geq \left\lceil \frac{2(R-1)(2R-1)}{R^2}\, n \right\rceil.$$

Proof. For both cases, Proposition 2.6 gives $r_Q, s_Q \leq R$ for every selected rectangle. If $b = 2R + 1$, then $w = 2R$, and the rank bound (2.18) gives

$$C_{\text{rect}}(M) \geq \left\lceil n(2R-1)\frac{2R-2}{R^2-1} \right\rceil = \left\lceil \frac{2(2R-1)}{R+1} n \right\rceil.$$

This proves (2.21).

Now let $b = 2R$, so $w = 2R - 1$. This weight gives an additional capacity restriction: two disjoint groups of R positions cannot fit in one row or column. In particular, a row belongs to at most one width R rectangle, and a column belongs to at most one height R rectangle.

First consider $R = 2$, for which $w = 3$. Every rectangle has size 2×2 and saves one addition. Since each row occurs in at most one such rectangle, there are at most $\lfloor n/2 \rfloor$ selected rectangles. Hence

$$C_{\mathrm{rect}}(M) \geq 2n - \lfloor n/2 \rfloor = \lceil 3n/2 \rceil,$$

as required by (2.22).

For the remaining case, assume $R \geq 3$ and fix a rectangle scheme $\mathcal{S}$. Counting the row occurrences of width R rectangles and the column occurrences of height R rectangles separately gives

$$\sum_{Q \in \mathcal{S}} r_Q \mathbf{1}_{s_Q = R} \leq n, \qquad \sum_{Q \in \mathcal{S}} s_Q \mathbf{1}_{r_Q = R} \leq n.$$

Here $\mathbf{1}$ is one when its subscript condition holds and zero otherwise. The combined count is at most $2n$. We use its average

$$T := \frac{1}{2} \sum_{Q \in \mathcal{S}} \left(r_Q \mathbf{1}_{s_Q = R} + s_Q \mathbf{1}_{r_Q = R} \right) \leq n$$

to express this capacity restriction symmetrically.

We bound the saving of each rectangle by combining its rank budget contribution with its contribution to T. If both dimensions are less than R, the saving per unit of rank budget satisfies

$$\frac{(r-1)(s-1)}{rs-1} \leq \frac{(R-2)^2}{(R-1)^2 - 1} = \frac{R-2}{R}.$$

Thus the rank budget alone bounds the savings of these smaller rectangles with coefficient $(R-2)/R$. For an $R \times R$ rectangle, the saving exceeds this bound by

$$(R-1)^2 - \frac{R-2}{R}(R^2 - 1) = \frac{2(R-1)}{R}.$$

Such a rectangle contributes R to T. Multiplying this contribution by $2(R-1)/R^2$ supplies exactly the difference. These two coefficients suggest the following bound for every allowed size:

$$(r-1)(s-1) \leq \frac{R-2}{R}(rs-1) + \frac{2(R-1)}{R^2} \frac{r\mathbf{1}_{s=R} + s\mathbf{1}_{r=R}}{2}. \tag{A.1}$$

The cases $r, s < R$ and $r = s = R$ were verified above. It remains to check a rectangle with exactly one dimension equal to R. By symmetry, take $r = R$ and $2 \leq s \leq R-1$. The right side of (A.1) minus the left side equals

$$\frac{1 + (R-1-s)(R^2 - R + 1)}{R^2} \geq \frac{1}{R^2} > 0.$$

Hence the bound holds for every admissible rectangle.

We can now sum (A.1) over the scheme. The first term is controlled by the rank budget $\sum_Q (r_Q s_Q - 1) \leq n(w-1) = (2R-2)n$ from Lemma 2.7. The second is controlled by $T \leq n$. Consequently, the total saving satisfies

$$\begin{aligned} \sum_{Q \in \mathcal{S}} (r_Q - 1)(s_Q - 1) &\leq \frac{R-2}{R} \sum_{Q \in \mathcal{S}} (r_Q s_Q - 1) + \frac{2(R-1)}{R^2} T \\ &\leq \frac{R-2}{R}(2R-2)n + \frac{2(R-1)}{R^2} n = \frac{2(R-1)^3}{R^2} n. \end{aligned}$$

Subtracting this saving from the direct cost gives

$$C_{\mathrm{rect}}(M) \geq (2R-2)n - \frac{2(R-1)^3}{R^2} n = \frac{2(R-1)(2R-1)}{R^2} n.$$

Taking the ceiling proves (2.22). □

Necessary structure for equality. Equality in the real bound before taking the ceiling forces equality throughout the corresponding proof. For odd $b = 2R + 1$, every selected rectangle must have size $R \times R$, and the rank budget must be saturated. The number of rectangles is then

$$\frac{(2R-1)n}{R^2-1},$$

and the number of uncovered nonzero entries is n minus this number.

For even $b = 2R$ with $R \geq 3$, equality in (A.1) permits only $R \times R$ and $(R-1) \times (R-1)$ rectangles. Saturating $T \leq n$ requires every row and every column to occur in exactly one $R \times R$ rectangle, so there are n/R such rectangles. Saturating the rank budget then requires

$$\frac{(R-1)^2 n}{R^2(R-2)}$$

rectangles of size $(R-1) \times (R-1)$. For $R = 2$, equality requires $n/2$ rectangles of size 2×2. These counts must be integers. They are necessary for equality in the real bound. Attaining its ceiling need not force the same structure.

B Supplementary Data and Proofs

B.1 Theorem 3.3 in Section 3.1.1

Let M be an invertible binary circulant of order $n \geq 4$, weight three, and branch number four, with first row support D. The rectangle cost has the following classification.

1. *If n is even and D contains an antipodal pair, then*

$$C_{\mathrm{rect}}(M) = \frac{3n}{2}.$$

2. *If D has no antipodal pair and $D = \{u, u+d, u+2d\}$, put $g = \gcd(n, d)$ and $q = n/g$. Such a support is invertible exactly when $3 \nmid q$. In the invertible case, $q \geq 5$ and*

$$C_{\mathrm{rect}}(M) = 2n - g\left\lfloor \frac{q}{2} \right\rfloor.$$

3. *For every other invertible support, $C_{\mathrm{rect}}(M) = 2n$.*

Proof. First suppose D contains $u, u + n/2$. For every row i, rows i and $i + n/2$ share the columns $i + u$ and $i + u + n/2$. Pairing these opposite rows gives a matching of size $n/2$, which is the largest possible. Equation (3.1) gives the first cost. Such binary supports are always invertible: after a support translation, write $M = I + P^{n/2} + P^d$. Then $M^2 = P^{2d}$, an invertible permutation matrix.

To classify all remaining reuse, suppose two rows share two columns. Their index difference gives a nonzero translation d with $c(d) = 2$. Equivalently, there are two directed steps of length d between points of D. With only three support points, these two steps must either form a path on three points or form a cycle on two points. The cycle requires $2d = 0$ and gives an antipodal pair. The path gives $D = \{u, u + d, u + 2d\}$. Thus every support admitting reuse belongs to one of the first two cases.

For the second case, remove the support translation u and consider $M = I + P^d + P^{2d}$. The permutation P^d consists of g cycles of length q. On each cycle the matrix is $I + P_q + P_q^2$. Its polynomial $1 + X + X^2$ has as roots the nontrivial cube roots of unity, so it is coprime to $X^q - 1$ exactly when $3 \nmid q$. This proves the invertibility criterion. Three distinct support

points require $q \geq 3$, while $q = 3$ is singular and $q = 4$ contains an antipodal pair. Hence the present invertible case has $q \geq 5$.

In this case the only translations with overlap two are d and $-d$. To see this, the two directed steps associated with any such translation must again form a path through all three support points, since an antipodal pair is excluded. If the middle point is $u + d$, the path traverses the endpoints $u, u + 2d$ in one of the two orders, so its step is d or $-d$. Any other step would therefore require one of these endpoints to be the middle point. If u were the middle, the sum of the other two points would equal $2u$, giving

$$2u = (u + d) + (u + 2d), \qquad 3d = 0 \pmod{n}.$$

Choosing $u + 2d$ as the middle gives the same condition. Since the support points are distinct, d would then have order three, and D would be a coset of the subgroup $\{0, d, 2d\}$. In particular, $D + d = D$, so rows i and $i + d$ would have identical supports and hence be equal over $\mathbb{F}_2$. This contradicts invertibility.

Consequently, $\Gamma(M)$ joins row i precisely to rows $i + d$ and $i - d$. Starting from row i, repeatedly following the edge of step d visits

$$i,\ i + d,\ i + 2d,\ \ldots,\ i + (q - 1)d$$

before returning to i, because $q = n/\gcd(n, d)$ is the smallest positive integer for which $qd = 0 \pmod{n}$. These vertices form one cycle, and the n vertices partition into $g = n/q$ such cycles. On a cycle of length q, a matching has at most $\lfloor q/2 \rfloor$ edges because its edges use disjoint pairs of vertices. This number is attained by pairing consecutive vertices along the cycle, leaving one vertex unpaired when q is odd. Matchings can be selected independently in the g cycles, so

$$\mu(\Gamma(M)) = g \left\lfloor \frac{q}{2} \right\rfloor.$$

Substitution into (3.1) proves the second cost formula. The remaining supports have no reusable rectangle, so their cost is $2n$. □

B.2 Binary circulants of weight three

Table 8 reports the complete enumeration described in Section 3.1.4. The tested supports all include position 0. The connected counts refer to invertible supports with connected support graphs. The two minimum-cost columns respectively allow all such invertible supports and require connectivity. Every minimum also agrees with Theorem 3.4 and Proposition 3.5 where applicable.

Table 9 lists every odd order $9 \leq n \leq 128$ with $3 \mid n$ and $7 \nmid n$. The construction in Proposition 3.6 uses fewer gates than the minimum rectangle cost for connected matrices at every listed order. At orders 9, 27, and 81, it also improves on the minimum rectangle cost without a connectivity requirement. For each order, we checked all output coefficient vectors over $\mathbb{F}_2$ against those of $I + P + P^3$. All 8,228 output checks passed. The gate counts describe these explicit circuits and do not assert optimality. The two $\min C_{\text{rect}}$ columns give the minimum rectangle costs over invertible binary circulants of the same order and weight three, without and with a connectivity requirement. The former follows from Theorem 3.4. The latter equals $2n$ by Proposition 3.5.

Table 8: Complete support enumeration for $4 \le n \le 33$. Counts retain all supports containing position 0.

n	Tested	Invertible	Connected	$\min C_{\mathrm{rect}}$	$\min_{\mathrm{conn}} C_{\mathrm{rect}}$
4	3	3	3	6	6
5	6	6	6	8	8
6	10	6	6	9	9
7	15	9	9	11	11
8	21	21	18	12	12
9	28	18	18	18	18
10	36	36	30	15	15
11	45	45	45	17	17
12	55	39	30	18	18
13	66	66	66	20	20
14	78	54	45	21	21
15	91	54	48	24	30
16	105	105	84	24	24
17	120	120	120	26	26
18	136	96	72	27	27
19	153	153	153	29	29
20	171	171	132	30	30
21	190	99	90	33	42
22	210	210	165	33	33
23	231	231	231	35	35
24	253	189	132	36	36
25	276	276	270	38	38
26	300	300	234	39	39
27	325	234	216	54	54
28	351	255	198	42	42
29	378	378	378	44	44
30	406	258	168	45	45
31	435	345	345	47	47
32	465	465	360	48	48
33	496	375	330	51	66

Table 9: Nine-gate constructions and rectangle minima for odd $9 \le n \le 128$ with $3 \mid n$ and $7 \nmid n$.

n	Groups of five	Remaining outputs	Gates	$\min C_{\mathrm{rect}}$		Depth
				No connectivity constraint	Connected	
9	1	4	17	18	18	3
15	3	0	27	24	30	3
27	5	2	49	54	54	3
33	6	3	60	51	66	3
39	7	4	71	60	78	3
45	9	0	81	72	90	3
51	10	1	92	78	102	3
57	11	2	103	87	114	3
69	13	4	125	105	138	3
75	15	0	135	114	150	3
81	16	1	146	162	162	3
87	17	2	157	132	174	3
93	18	3	168	141	186	3
99	19	4	179	153	198	3

n	Groups of five	Remaining outputs	Gates	$\min C_{\text{rect}}$ No connectivity constraint	$\min C_{\text{rect}}$ Connected	Depth
111	22	1	200	168	222	3
117	23	2	211	180	234	3
123	24	3	222	186	246	3

B.3 Proposition 3.8 in Section 3.2.1

An invertible binary circulant of weight five has branch number six if and only if $c(d) \leq 3$ for every nonzero d and every row of M^{-1} has weight at least five. The inverse M^{-1} is also circulant, so all its row and column weights are equal. If all nonzero overlaps satisfy $c(d) \leq 2$, then $B(M) = 6$ without any additional test of the row weight of M^{-1}.

Proof. An input of weight one gives an output of weight five. For an input of weight two whose nonzero positions differ by d, the output weight is $10-2c(d)$. Thus the branch number requirement for these inputs is precisely $c(d) \leq 3$.

As noted at the beginning of Section 3, the input and output weights have the same parity. An input of weight three therefore has an output of odd weight. The output is nonzero, and the requirement on its weight is at least three. Thus the only violation occurs when the output has weight one. A bad input of weight three exists exactly when a column of M^{-1} has weight three. Since this inverse is an invertible binary circulant, its common row and column weight is odd. It cannot be one, since then M^{-1} and M would be permutation matrices. Its remaining possibilities are therefore three or at least five. Requiring its row weight to be at least five excludes precisely the bad inputs under consideration.

An input of weight four has a nonzero output of even weight, and an input of weight five has a nonzero output. All larger input weights already give a total weight of at least six. This proves the criterion.

For the last assertion, $c(d) \leq 2$ suffices to exclude a weight-three input with a weight-one output. The output is the XOR of three distinct columns, each of weight five. There are three pairwise support intersections and one intersection common to all three columns. Each pairwise intersection reduces the initial weight count by twice its size. A position common to all three columns is counted three times initially and subtracted twice for each of the three pairs. Thus the inclusion–exclusion expression adds back four times the size of the triple intersection. The added term is nonnegative, and each pairwise intersection has size at most two. The output weight is therefore at least $15 - 2(2 + 2 + 2) = 3$. This excludes an output of weight one and proves the second assertion. □

B.4 Theorem 3.9 in Section 3.2.1

Let $D \subseteq \mathbb{Z}_n$ have five elements and satisfy $c(d) \leq 3$ for every nonzero d. Its circulant support contains a 3×3 all-one submatrix if and only if $3 \mid n$ and D contains a coset of the subgroup of order three.

Proof. Let I and J be the row and column sets of the submatrix. They each have three elements, and $J - I \subseteq D$. Thus the three sets $J - i$, for $i \in I$, have a union of size at most five. If every pair intersected in at most one point, their union would have size at least six. Hence J and $J + d$ intersect in at least two points for some nonzero d.

If the intersection has size three, then $J + d = J$. Since J has three points, it is a coset of a subgroup of order three. Distinct translates of this coset are disjoint, so the union bound forces all three translates to be identical. Their common coset is contained in D.

Suppose the intersection has size two. The directed steps of length d on J either form a path through its three points or a cycle through two of them, as in the proof of Theorem 3.3. In the path case, translate J to $\{0, d, 2d\}$ and let q be the order of d. The case $q = 3$ has just been treated. Otherwise J has no nonzero period, and distinct translates intersect in at most two points. We are considering the three sets $J - i$, for $i \in I$. Each has three points. Since their union has at most five points, inclusion–exclusion shows that the sum of the three pairwise intersection sizes is at least four. Represent these three sets by three vertices, joining two vertices when the corresponding sets intersect. This graph is connected. Otherwise it would have at most one edge, and the sum of the pairwise intersection sizes would be at most two, a contradiction.

Since $J = \{0, d, 2d\}$, whenever two shifted copies of J intersect, the difference between their shifts is an integer multiple of d. Following the edges of the connected graph therefore shows that the three shifts belong to the same coset of $\langle d \rangle$. These three shifts are exactly the three positions in $-I$.

For $q = 4$ or 5, this coset contains only q positions. Fix any target position in it. The three positions in $-I$ and the three positions obtained by subtracting J from that target both lie in this coset. Since $3 + 3 > q$, the two sets intersect. Hence the target is a sum of a position in $-I$ and a position in J. Every position in the coset therefore belongs to $J - I \subseteq D$. Adding d permutes the whole coset, so all its q positions contribute to the overlap of D with its shift by d. This gives $c(d) \geq q \geq 4$, a contradiction. For $q \geq 6$, view the three points of $-I$ on the cycle generated by d. Adding $\{0, d, 2d\}$ fills up to two positions in each gap after an occupied interval. The total number of empty positions is $q - 3 \geq 3$. A union of at most five points permits at most two new positions, so only one gap can be nonempty. The points of $-I$ are therefore consecutive, and adding $J = \{0, d, 2d\}$ gives five consecutive points in $J - I$. This again gives $c(d) \geq 4$.

It remains to consider a set J containing an antipodal pair but not forming a three-term progression. Its only overlap of size two is at the antipodal shift. Since the three pairwise overlaps of the translates have total size at least four, two translating indices must also be antipodal. Adding these two sets produces two distinct complete antipodal pairs inside D. The antipodal overlap of D is then at least four, another contradiction.

Conversely, a three-point subgroup coset in D supplies three rows with three common columns, giving the required submatrix. □

C Proofs for Section 4

C.1 Proposition 4.2 in Section 4.1.1

Throughout these two proofs, M is an invertible signed circulant over an odd prime field $\mathbb{F}_p$, of order $n \geq 4$ and weight three.

The matrix M has $B_d(M) = B_l(M) = 4$ if and only if its support D is not a coset of a subgroup of order three. The excluded supports are precisely

$$D = \{u, u + n/3, u + 2n/3\}, \qquad 3 \mid n.$$

Proof. A weight-one input gives an output of weight three, so $B_d(M) \leq 4$. Suppose first that two distinct columns have the same three support positions. Their index difference is a nonzero translation preserving D. The cycles of this translation partition D into sets of equal size, so their length must be three. Thus D is a coset of a subgroup of order three. Conversely, such a support has distinct columns with identical support.

For two columns with identical support, the three ratios of their nonzero coefficients lie in $\{1, -1\}$. At least two ratios are equal. Choosing a two-input linear combination to cancel these two positions leaves output weight at most one. Hence $B_d(M) \leq 3$.

If D is not such a coset, distinct columns share at most two positions. In a linear combination with two nonzero input coefficients, each column then has at least one position outside the other column's support. These two positions cannot cancel, so the output has weight at least two. Inputs of weight at least three have nonzero outputs by invertibility, and therefore also have total input-output weight at least four. This proves $B_d(M) = 4$. The equality $B_l(M) = B_d(M)$ follows from Proposition 2.1. □

C.2 Proposition 4.5 in Section 4.1.3

Let $p \geq 5$. For every invertible signed circulant M of weight three and branch number four, the minimum gate count among circuits of depth at most two equals $C_{\text{rect}}(M)$. Thus an implementation using fewer gates than rectangle reuse must have depth at least three.

Proof. A signal at depth at most one has support size at most two. If it has two distinct inputs, their coefficients are 1 or -1. If it has just one input, its coefficient can also be 2 or -2 because a gate may use that input twice. Every output has weight three and therefore requires a final gate at depth two.

Suppose both operands of this final gate have support size two. Their supports must intersect in exactly one position to produce an output of weight three. At that position the resulting coefficient is 0, 2, or -2. Zero would reduce the output weight to two, while $\pm 2 \notin \{1, -1\}$ for $p \geq 5$. Thus this case is impossible. The final gate must instead combine a signed pair with a third, distinct input. The third input has coefficient 1 or -1; a doubled input cannot give a signed output.

After unused gates and duplicate pair computations are removed, each output uses one depth-one pair and one further input. Any pair used by two or more outputs defines an admissible rectangle. Each row uses only one pair, so these rectangles have disjoint matrix positions. The circuit is therefore a rectangle scheme and cannot cost less than $C_{\text{rect}}(M)$. Conversely, every rectangle scheme in this class computes its shared and private pairs at depth one and completes its outputs at depth two. This proves equality. □

C.3 Lemma 4.7 in Section 4.2.1

Let M be an invertible signed circulant of order n and weight four over an odd prime field.

Suppose any two columns of M share at most two positions. Then $B(M) = 5$ if and only if no three columns form the following configuration. Every pair shares exactly two positions, and the triple intersection is empty. Exactly two column pairs have equal coefficient ratios. The remaining pair has opposite coefficient ratios.

Proof. Two nonzero inputs leave at least two private positions in each column, so their output weight is at least four. Inputs of weight one or at least four already satisfy the target. We need only exclude a combination of three columns with output weight one.

If their triple intersection has size $r \geq 1$, each column has at least $4 - 2 - 2 + r = r$ positions outside the other two columns. At least three private positions therefore survive. If the triple intersection is empty and some pair intersects in at most one position, each of those two columns has a private position, so at least two positions survive.

The remaining case consists of three disjoint two-position groups, one for each column pair, with no private positions. In a group with equal ratios, its two entries cancel together or both remain nonzero. A group with opposite ratios always leaves at least one nonzero entry. With no such group, the output has even weight and cannot be zero by invertibility, so its weight is at least two. With at least two such groups, at least two entries survive. With exactly one, choose the three input coefficients to cancel the other two groups. In the

last group, one entry cancels and the other is a nonzero multiple of 2, giving output weight one. This is exactly the excluded configuration. □

C.4 Lemma 4.13 in Section 4.3.1

Let M be a connected invertible signed circulant of order $n \geq 15$, weight five, and branch number six over an odd prime field. If M contains an admissible 3×3 rectangle, then

$$C_{\mathrm{rect}}(M) \geq \left\lceil \frac{22n}{9} \right\rceil .$$

Proof. Any two rows of the rectangle have three common positions with the same coefficient ratio. Let d be their row displacement; their total overlap is $c(d)$, as for two columns. Combining these rows to cancel those positions leaves at most $10-c(d)-3$ nonzero positions. Branch number six requires this number to be at least four, so $c(d) = 3$.

The support argument in the proof of Theorem 3.9 now applies to the three selected rows. That argument considers three translates of a three-point set whose union has at most five points. Unless the set is a coset of a subgroup of order three, its progression and antipodal cases both produce two selected rows with at least four common positions. The preceding cancellation excludes both cases here. Thus, after a support shift, the three core positions are $0, m, 2m$, where $n = 3m$. The three rows of the core block are cyclic shifts of one another. Rank one makes consecutive rows proportional with a ratio in $\{1, -1\}$. Three cyclic shifts return to the original nonzero row, so the cube of this ratio is one. The ratio must therefore be one, and the three core coefficients are equal. After an overall sign change, we may write

$$M = H + \sigma_1 P^{\delta_1} + \sigma_2 P^{\delta_2}, \qquad H = I + P^{n/3} + P^{2n/3}, \qquad \sigma_1, \sigma_2 \in \{1, -1\}.$$

The residues of δ_1, δ_2 modulo $m = n/3$ are nonzero and distinct. Otherwise, two rows separated by m would have an additional common position besides their three equally signed core entries, contradicting the same two-row bound. Connectivity gives $\gcd(m, \delta_1, \delta_2) = 1$, and $m \geq 5$.

We next check that a rectangle extending outside a core block can only have size 2×2. Two distinct rows cannot share both additional columns: that would force $\delta_2 - \delta_1 = n/2$ modulo n. To see why this is impossible, let $E = P^{n/2}$. On its eigenspace with eigenvalue $-\sigma_1/\sigma_2$, the additional part $P^{\delta_1}(\sigma_1 I + \sigma_2 E)$ vanishes. This eigenspace has dimension $n/2$, whereas $\operatorname{rank} H \leq m < n/2$. It therefore contains a nonzero vector killed by H, giving a nonzero vector killed by M and contradicting invertibility.

For rows $0, d$ in different core blocks, their core columns are disjoint. The possible core/additional matches require $d \equiv \delta_1, \delta_2, -\delta_1, -\delta_2 \pmod m$, respectively. There are at most two such matches, because the additional residues are distinct. Three common columns would require two of these matches and one additional/additional match, the latter requiring $d \equiv \pm(\delta_2 - \delta_1) \pmod m$. If the two core/additional matches use the same extra residue, these congruences force the other residue to be zero. If they use different residues, they force $\delta_2 \equiv -\delta_1$ and $3\delta_1 \equiv 0 \pmod m$. The first possibility contradicts the nonzero residues; the second, together with $\gcd(m, \delta_1, \delta_2) = 1$, forces $m = 3$. Thus rows in different core blocks have at most two common columns. Applying the same argument to the transpose also confines every 3×2 rectangle to a core block.

As in the proof of Theorem 3.12, a crossing 2×2 rectangle has either one core position or two diagonally opposite core positions. Zero core positions would require the excluded antipodal additional shifts. Three core positions force the fourth to be core, and two core positions in the same row or column would force $\delta_1 \equiv \delta_2 \pmod m$. Let N_1, N_2 count the crossing rectangles with one and two core positions. They use $3N_1 + 2N_2$ additional positions, so $3N_1 + 2N_2 \leq 2n$.

The core positions form m disjoint 3×3 blocks. Every core position used by a crossing rectangle requires an additional position in its row and column. Thus at most two positions are removed from any row or column of a core block. The combinatorial loss table (3.10) therefore applies: after z removals the remaining internal saving is at most $4 - 2z/3$. This is an upper bound regardless of the additional signs, which can only exclude rectangles. Since $N_1 - N_2 \leq (3N_1 + 2N_2)/3 \leq 2n/3$, the total saving is at most

$$4m - \frac{2}{3}(N_1 + 2N_2) + N_1 + N_2 = \frac{4n}{3} + \frac{N_1 - N_2}{3} \leq \frac{14n}{9}.$$

Subtracting from the direct cost $4n$ proves the bound. □