Quantum Information & Computation
ISSN 1533-7146 Volume 25 (2025) pp. 586–597
©Cerebration Science Publishing

Article

Optimal T-depth Quantum Circuits for Implementing Arbitrary Boolean Functions

Suman Dutta [1,*], Anik Basu Bhaumik [2], Anupam Chattopadhyay [2] and Subhamoy Maitra [1]

[1] Applied Statistics Unit, Indian Statistical Institute, Kolkata 700108, West Bengal, India; sumand.iiserb@gmail.com (S.D.); subho@isical.ac.in or maitra.subhamoy@gmail.com (S.M.)
[2] College of Computing and Data Science, Nanyang Technological University, Singapore 639798, Singapore; anikbasu001@e.ntu.edu.sg (A.B.B.); anupam@ntu.edu.sg (A.C.)
* Corresponding author: sumand.iiserb@gmail.com

Received date: 1 June 2025; Accepted date: 1 September 2025; Published online: 11 November 2025

Abstract: In this paper, we present a generic construction for synthesizing an optimal T-depth quantum circuit for any arbitrary n-input, m-output Boolean function $f : \{0,1\}^n \to \{0,1\}^m$ with algebraic degree $k \leq n$, achieving an exact Toffoli (consequently T) depth of $\lceil \log_2 k \rceil$. This broadly generalizes the recent result establishing the optimal Toffoli (and T) depth for multi-controlled Toffoli decompositions (Dutta et al., Phys. Rev. A, 2025). The optimality of T-depth in this initiative is considered in the context of implementing an n-MCT, assuming the decomposition via Clifford plus Toffoli gates. The key technique involves inspecting the Algebraic Normal Form (ANF) of the Boolean function. Obtaining a benchmark for the minimum T-depth of such circuits is crucial for the efficient implementation of quantum algorithms by enabling greater parallelism, reducing time complexity, and minimizing circuit latency, making them suitable for near-term quantum devices with limited coherence times. The broader implications of our results include a provable lower bounds on T-depth for S-box and block cipher implementations, such as AES. Finally, we also explain the impact of our result in identifying the T-depth for the generic cryptanalysis of block ciphers using Grover's algorithm.

Keywords: AES; block cipher; Boolean function; cryptanalysis; efficient circuit implementation; S-box synthesis; T-depth; universal gate sets

1. Introduction

Quantum computing is one of the most fundamental aspects of quantum physics, with its origins traceable to Feynman's early insights into quantum simulation [1]. It is essential to note that quantum algorithms often rely on the efficient encoding of classical Boolean functions into quantum circuits through a process known as oracle construction. An oracle is a reversible quantum circuit that implements an unknown Boolean function and is integral to several landmark quantum algorithms, including Deutsch-Jozsa [2], Grover's search [3], Simon's hidden shift finding [4], and Shor's factoring algorithm [5].

Quantum gates, represented by unitary matrices, serve as the basic building blocks of quantum circuits. Unlike classical gates, they are inherently reversible. In fault-tolerant quantum computing [6], oracle construction becomes especially resource-intensive due to the high cost of implementing multi-controlled Toffoli (MCT) gates. These gates must be decomposed into a universal gate set, such as Clifford+T [7], where the key resource metrics like gate count, circuit depth, and qubit overhead become critical for the feasibility of large-scale quantum computations. Therefore, minimizing these resources is essential for improving algorithmic efficiency and mitigating computational errors.

Circuit depth is calculated by the number of layers of gates, assuming parallel execution on disjoint qubits. The T-depth explicitly measures the number of such layers containing T gates, a crucial measure in fault-tolerant quantum computation. Since T gates are resource-intensive to implement, often requiring costly magic state distillation, minimizing T-depth directly reduces circuit latency and execution time. This is particularly important for near-term quantum devices with limited coherence times. As a result, T-depth optimization has become a key focus in the design of efficient quantum circuits [8,9].

While recent advances in surface code implementations suggest that T-depth may not be the sole bottleneck in large-scale quantum circuit design [10,11], it nonetheless remains a critical metric. The overall cost of fault-tolerant

DOI: 10.2478/qic-2025-0032

quantum computation involves several factors, including T-count, qubit overhead, and scheduling strategies such as resource state distillation and teleportation. However, in the context of near-term quantum hardware with coherence times as the primary constraint, T-depth still plays a significant role in determining circuit latency and execution feasibility. Throughout the paper, we present our results under the standard assumption that T-depth is among the most resource-intensive parameters, and optimizing it is crucial for practical quantum circuit implementation.

In [12, Page 11], the authors emphasized the importance of precise resource estimation in implementing arbitrary n-input, m-output Boolean functions, generalizing beyond the scope of multivariate AND functions. Motivated by this, we present the first generic construction of optimal T-depth quantum circuits for arbitrary Boolean functions. Our approach generalizes and subsumes the methodology of [12], achieving the minimum possible Toffoli (and thus T) depth by exploiting the Algebraic Normal Form (ANF) of the Boolean function. While this construction incurs a considerable qubit overhead, it offers a compelling trade-off by substantially reducing T-depth. This makes our approach particularly valuable for efficient oracle construction, reversible circuit synthesis, and quantum implementations of cryptographic components such as S-boxes. We illustrate the practical relevance of our method by applying it to the quantum circuit design of AES, a widely used and structurally complex block cipher.

Numerous efforts have been made to reduce the quantum resource requirements for AES in last few years, including [13–19]. For example, [14, Table 7] and [16, Table VI] report a T-depth of 60 for AES and AES† combined (i.e., 30 per instance), while [19, Table 5] reports T-depths of 30, 36, and 42 for AES-128, AES-192, and AES-256, respectively. More recently, Huang et al. [18, Table 3] demonstrated an implementation of T-depth 3 for the 8-bit input 8-bit output AES S-box with a formal proof of its minimality.

In contrast, our approach is cipher-independent and provides a generic quantum circuit construction that achieves minimal T-depth for any Boolean function specified by its ANF, without putting any constraint on other resources. This construction also establishes a naive yet complete upper bound on ancillary resources, thereby setting benchmarks applicable to a broader class of ciphers beyond AES, or in general for any application that is understood by Boolean functions. In particular, our method generalizes the optimal T-depth MCT decomposition via Clifford plus Toffoli gates, as described in [12, Corollary 1]. All results are presented under the assumption that T-depth remains one of the most critical cost metrics for near-term quantum devices.

The structure of the paper and the contributions of each section are summarized below.

1.1. Organization and Contributions

Section 2 outlines the preliminaries. Section 3 is the prime contributory section in terms of theoretical result, which provides a worst-case resource bound for optimal T-depth quantum circuit implementation of any n-input, m-output Boolean function. We illustrate the result with examples and highlight its implications for S-box design, including a summary of optimal T-depth resource estimates for different standard S-boxes in Table 1. In Section 4, to explain with an example, we focus on AES, providing a step-by-step resource estimation for an optimal T-depth quantum implementation, and discuss the broader impact for a large class of block ciphers executed over multiple rounds. Section 5 explores the cryptanalytic implications of our results, particularly in the context of Grover's algorithm, and determines the optimal T-depth for a full-round Grover's search for a large class of block ciphers. Finally, Section 6 concludes the paper with a summary of contributions and future research directions.

2. Preliminaries

Let $\mathbb{F}_2 = \{0,1\}$ be the prime field of characteristic 2 and $\mathbb{F}_2^n \equiv \{\mathbf{x} = (x_1,\ldots,x_n) : x_i \in \mathbb{F}_2, 1 \le i \le n\}$ be the vector space of dimension n over $\mathbb{F}_2$. An n-input m-output Boolean function f is defined as a mapping $f : \mathbb{F}_2^n \to \mathbb{F}_2^m$. The set of all n-input m-output Boolean functions is denoted by $\mathcal{B}_n^m$. For $m = 1$, the set of all n-input, single-output Boolean functions is given by $\mathcal{B}_n^1 \equiv \mathcal{B}_n$.

The Algebraic Normal Form (ANF) of a Boolean function $f \in \mathcal{B}_n$ can be represented by a polynomial over $\mathbb{F}_2$ in n binary variables, given by

$$f(x_1,\ldots,x_n) = a_0 \oplus \bigoplus_{I \subseteq \{1,\ldots,n\}} a_I \prod_{i \in I} x_i$$

where $a_I \in \mathbb{F}_2$. The algebraic degree of a Boolean function $f \in \mathcal{B}_n$, denoted $\deg(f)$, is given by the maximum cardinality of an index set I, such that $a_I \neq 0$. Clearly, $\deg(f) \le n$. A Boolean function $f \in \mathcal{B}_n$ is nonlinear if $\deg(f) > 1$, affine if $\deg(f) = 1$, and constant if $\deg(f) = 0$.

For a Boolean function $f \in \mathcal{B}_n^m$, the output $f(\mathbf{x}) = \mathbf{a}$ does not uniquely determine the input $\mathbf{x}$, making f inherently irreversible. In contrast, quantum operations are reversible by nature (excluding measurement), and thus

the corresponding quantum circuit implementing f must be reversible. Such a circuit operates on $n+m$ qubits, where the first n qubits represent the input state $|\mathbf{x}\rangle$, and the remaining m qubits, initialized to $|\mathbf{y}\rangle$, store the functional output. The functioning of U_f is given by: $U_f : |\mathbf{x}\rangle\,|\mathbf{y}\rangle \rightarrow |\mathbf{x}\rangle\,|\mathbf{y} \oplus \mathbf{f}(\mathbf{x})\rangle$.

Given the ANF of a Boolean function $f \in \mathcal{B}_n$, a reversible quantum circuit implementing f can be constructed using CNOT gates for linear terms and multi-controlled Toffoli (MCT) gates for nonlinear terms, where the relevant input variable(s) act as control qubit(s) and the output qubit serves as the target. However, since MCT gates are not native to most quantum hardware, they must be decomposed into a universal gate set for practical implementation.

The Clifford+T gate set is the most widely adopted universal gate set in quantum computing due to its compatibility with fault-tolerant quantum computation. The Clifford group consists of the Hadamard, Phase, and CNOT gates. Augmenting it with the non-Clifford T gate yields a universal set capable of approximating any unitary operation to arbitrary precision.

Recently, substantial efforts have been directed toward optimizing the Clifford+T decomposition of the Toffoli gate. State-of-the-art techniques employ measurement-based uncomputation, achieving Toffoli decompositions with four T gates. Notably, Gidney's construction [20] attains an effective T-depth marginally above 1 without employing any ancilla qubit (see Figure 1a) by executing the initial T gates in parallel. In contrast, design provided in Jaques et al. [13] achieves a T-depth of 1 using a single reusable ancilla qubit (see Figure 1b). A comprehensive summary of optimized Toffoli decompositions across various resource metrics is provided in [12, Table 1].

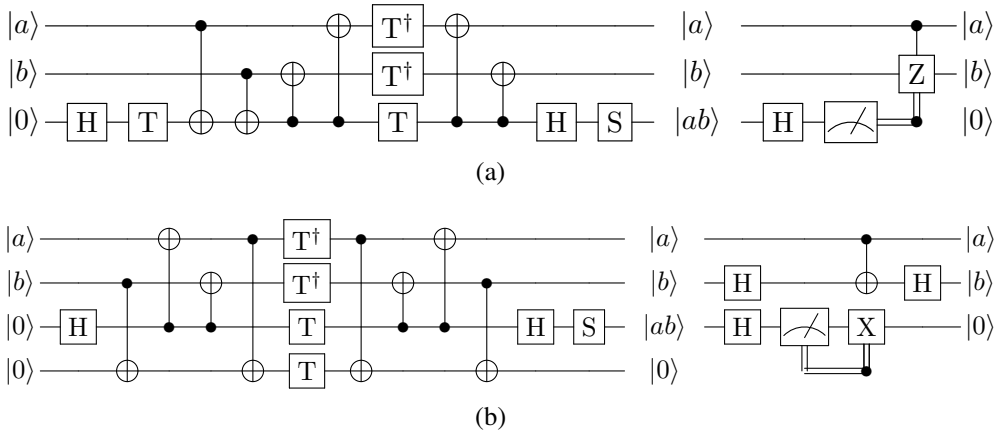

Figure 1. Toffoli decomposition with measurement-based uncomputation, using four T gates, (**a**) without using any ancilla qubit [20], (**b**) with a single reusable ancilla qubit [13].

In addition to optimizing basic Toffoli gates, significant progress has been made toward efficient implementations of multi-controlled Toffoli (MCT) gates. A recent work [12] demonstrates that employing a binary tree structure in the decomposition of an n-MCT gate achieves an optimal T-depth of $\lceil \log_2 n \rceil$, requiring at most $2(n-1)$ ancilla qubits and $4(n-1)$ T gates. In [12], the optimality was first established in terms of Toffoli depth, assuming the MCT gate is decomposed into Clifford and Toffoli gates. Subsequently, by applying the Toffoli gate decomposition shown in Figure 1b, the optimality result was extended to T-depth as well.

In this regard, one may note that, under a different set-up, the T-depth can be reduced to as low as one by employing different constructions, such as using a large number of ancilla qubits and CCZ resource states, and leveraging teleportation-based techniques, as outlined in [21]. However, such explorations are beyond the scope of this paper. In the present paper, we focus on a specific computational model and establish the optimality within the framework specified in [12].

A comprehensive summary of recent advancements in MCT circuit decompositions is presented in [12, Table II]. In the following section, we extend this line of work by presenting a generic construction for quantum circuits with optimal Toffoli (and hence T) depth for evaluating arbitrary Boolean functions. Our key insight generalizes the approach of [12] by utilizing the ANF framework, particularly the XOR ($\mathbb{F}_2$-addition) of AND ($\mathbb{F}_2$-multiplication) components, in conjunction with tree-based circuit synthesis. We detail this construction in the next section.

3. Optimal T-depth Quantum Resource Estimation

In this section, we present an optimal T-depth quantum circuit for implementing any arbitrary n-input, m-output Boolean function, thereby completing the extension proposed in [12]. Before proceeding, we revisit [12, Corollary 1] and formally provide the idea in Theorem 1 to make the paper self-contained.

In [12], the optimality of T-depth was established by structuring the MCT decomposition as a binary tree, where the Toffoli depth corresponds to the height of the tree, which is lower bounded by $\lceil \log_2 n \rceil$ for n leaf nodes. Each Toffoli gate was further decomposed into the Clifford+T gate set using the T-depth-1 construction shown in Figure 1b, yielding an overall T-depth of $\lceil \log_2 n \rceil$. This is the optimal T-depth for implementing an n-MCT, assuming the decomposition is via Clifford plus Toffoli gates. We refer to the T-depth optimality obtained from this specific framework, [12] as MCT-Toffoli-T optimality.

Throughout this paper, any reference to 'optimal T-depth' adheres to this MCT-Toffoli-T optimality, obtained via Clifford plus Toffoli decomposition.

Theorem 1. *([12, Corollary 1]) Following the Toffoli decomposition circuit proposed in [13] (see Figure 1b), the Clifford+T decomposition of an n-MCT gate, implemented via Clifford plus Toffoli decomposition, is lower bounded by $\lceil \log_2 n \rceil$, requiring $2n-2$ additional ancilla qubits and $4n-4$ T gates.*

As discussed earlier, implementing higher-degree terms in the ANF of a Boolean function requires multi-controlled Toffoli (MCT) gates, where the input variables of the monomial serve as control qubits and the corresponding output qubit as the target. Specifically, a degree-k monomial requires a k-MCT gate. According to Theorem 1, a k-MCT gate can be decomposed into the Clifford+T gate set (via Toffoli gates) with an optimal T-depth $\lceil \log_2 k \rceil$. Similarly, implementing the highest-degree term $x_1 x_2 \cdots x_n$ using a binary tree structure incurs a T-depth of $\lceil \log_2 n \rceil$.

Notably, the ANF of an arbitrary n-input, m-output Boolean function $f \in \mathcal{B}_n^m$ can contain up to $2^n - (n+1)$ unique nonlinear terms. More precisely, the ANF includes at most $\binom{n}{k}$ unique degree-k monomials, each requiring a k-MCT gate for its quantum implementation. If all k-MCT gates are applied in parallel, the MCT depth becomes 1, with the largest being an n-MCT contributing to a T-depth of $\lceil \log_2 n \rceil$. Once implemented, the nonlinear monomials can be XOR-ed to the appropriate output qubit using CNOT gates, without increasing the T-depth. In this manner, the complete ANF of any function $f \in \mathcal{B}_n^m$ can be implemented in a quantum circuit with an optimal T-depth of $\lceil \log_2 n \rceil$.

Additionally, the lower bound can be justified by noting that any Boolean circuit can be expressed as a composition of additions and multiplications. Since each layer of parallel multiplication increases the algebraic degree by at most one, the algebraic degree of the function is upper bounded by the maximum multiplicative depth.

To enable parallel MCT operations, multiple copies of the input variables and additional ancilla qubits are required to store intermediate results. These copies can be generated using CNOT gates, which, being Clifford operations, do not affect the T-depth. The result is summarized formally in the following theorem.

Theorem 2. *Let $f \in \mathcal{B}_n^m$ be an n-input, m-output Boolean function. Then, the Clifford+T decomposition of the quantum circuit implementing f can be realized with an optimal T-depth $\lceil \log_2 n \rceil$, using the following resources: (i) at most $2^{n-1}(3n-2) - 3n + m + 1$ reusable ancilla qubits, (ii) a total of $2^{n+1}(n-2) + 4$ T gates, (iii) a maximum of $2^{n-1}(11n + 2m - 18) - 4n - m + 9$ CNOT gates, with (iv) a CNOT depth $2^n + 2n + 9\lceil \log_2 n \rceil - 3$.*

Proof. There are $\binom{n}{k}$ degree-k monomials, each involving k variables. Computing all such nonlinear terms in parallel requires a total $\sum_{k=2}^{n} k\binom{n}{k}$ copies of the input variables. Since one copy of each input is already available, this entails an additional $\sum_{k=2}^{n} k\binom{n}{k} - n$ ancilla qubits and an equal number of CNOT gates for both computation and uncomputation. These introduce a CNOT depth of $2\lceil \log_2(\sum_{k=2}^{n} \binom{n-1}{k-1} - n) \rceil = 2(n-1)$. Additionally, $\sum_{k=2}^{n} \binom{n}{k}$ ancilla qubits are needed to store the outputs of these MCT gates.

After realizing all the unique nonlinear monomials from the ANF, any m-output Boolean function can be constructed by copying up to $\sum_{k=2}^{n} \binom{n}{k} + n$ monomials per function to m ancilla qubits. This requires at most $m(\sum_{k=2}^{n} \binom{n}{k} + n)$ CNOT gates and adds a maximum CNOT depth of $2^n - 1$.

Finally, by [12, Corollary 1], each k-MCT gate in its optimal T-depth decomposition uses $2(k-1)$ ancilla qubits and $9(k-1)$ CNOT gates. Ancilla qubits are made reusable via uncomputation. Since each T-depth layer corresponds to a CNOT depth of 9, the overall CNOT depth contribution is $9\lceil \log_2 n \rceil$. Hence, the overall resource requirement is given by:

- **#Ancilla qubits:**

$$\left[\sum_{k=2}^{n} k\binom{n}{k} - n\right] + \sum_{k=2}^{n}\binom{n}{k} + m + \sum_{k=2}^{n} 2(k-1)\binom{n}{k} = 2^{n-1}(3n-2) - 3n + m + 1,$$

- **#T gates:**

$$\sum_{k=2}^{n} 4(k-1)\binom{n}{k} = 2^{n+1}(n-2) + 4,$$

- **#CNOT gates:**

$$2\left[\sum_{k=2}^{n} k\binom{n}{k} - n\right] + m\left[\sum_{k=2}^{n}\binom{n}{k} + n\right] + \sum_{k=2}^{n} 9(k-1)\binom{n}{k} = 2^{n-1}(11n + 2m - 18) - 4n - m + 9,$$

- **CNOT depth:**

$$2(n-1) + (2^n - 1) + 9\lceil \log_2 n \rceil = 2^n + 2n + 9\lceil \log_2 n \rceil - 3.$$

□

Theorem 2 directly applies to the quantum circuit synthesis of S-boxes, where an n-bit S-box is modeled as an n-input, n-output Boolean function. The resulting circuit requires at most $2^{n-1}(3n-2) - 2n + 1$ reusable ancilla qubits, $2^{n+1}(n-2) + 4$ T gates, and up to $2^{n-1}(13n-18) - 5n + 9$ CNOT gates. The CNOT depth is bounded by $2^n + 2n + 9\lceil \log_2 n \rceil - 3$, while achieving an optimal T-depth of $\lceil \log_2 n \rceil$. To illustrate the construction, we present the 3-bit S-box used in LowMC [22] as an example.

Example 1. *The coordinate Boolean functions are given by:* $f_0 = x_0 \oplus x_1x_2$, $f_1 = x_0 \oplus x_1 \oplus x_0x_2$, *and* $f_2 = x_0 \oplus x_1 \oplus x_2 \oplus x_0x_1$. *Since the maximum algebraic degree is 2, the corresponding quantum circuit (see Figure 2) achieves the optimal T-depth* $\lceil \log_2 2 \rceil = 1$, *using 9 ancilla qubits, 12 T gates, and at most 33 CNOT gates. Notably, as the ANFs of these Boolean functions do not contain all possible monomials, the actual resource requirements are significantly lower than the worst-case estimates provided in Theorem 2. In practice, such sparsity often results in substantially reduced overheads compared to theoretical upper bounds.*

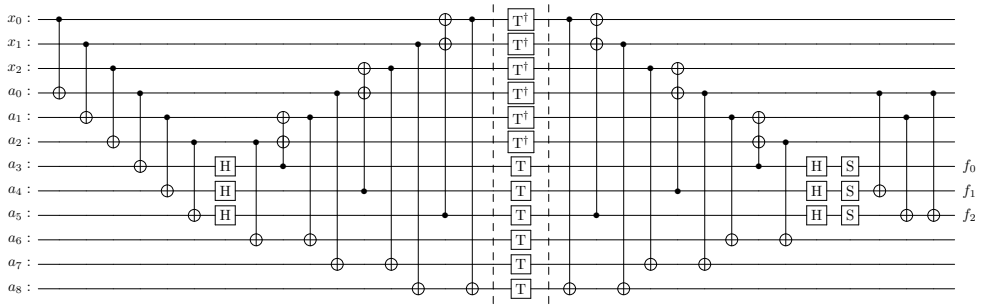

Figure 2. Quantum circuit implementing a 3-bit S-box (used in LowMC) with T-depth 1.

In Table 1, we summarize the quantum resource requirements for synthesizing various standard S-boxes with optimal T-depth following Theorem 2.

Table 1. Optimal T-depth quantum circuit synthesis for various standard S-boxes with corresponding resource estimates using Theorem 2. *A detailed analysis for the AES S-box is provided in Section 4.

S-box	#Variables	Ancilla count	CNOT count	CNOT depth	T-count	T-depth
LowMC [22]	3	9	39	13	12	1
DEFAULT [23]	4	17	79	27	24	2
GIFT [24]	4	14	76	27	24	2
PRESENT [25]	4	19	105	30	32	2
PRINCE [26]	4	24	128	28	40	2
ASCON [27]	5	27	120	22	32	1
AES [28]*	8	596	3647	184	984	3

From Theorem 2, the optimal T-depth Clifford+T decomposition of an n-input, single-output Boolean function $f \in \mathcal{B}_n$ requires at most $2^{n-1}(3n-2)-3n+2$ reusable ancilla qubits, $2^{n+1}(n-2)+4$ T gates, and up to $2^{n-1}(11n-16)-4n+8$ CNOT gates, with a CNOT depth of $2^n+2n+9\lceil\log_2 n\rceil-3$. We illustrate the circuit construction with a representative example.

Example 2. *Let $f = x_0x_2 \oplus x_1x_3 \oplus x_0x_1x_2x_3$ be a Boolean function consisting of three nonlinear terms. Since $\deg(f) = 4$, its Clifford+T decomposition yields a T-depth of $\lceil\log_2 4\rceil = 2$. The resulting quantum circuit requires 12 ancilla qubits, 20 T gates, and 46 CNOT gates, with a CNOT depth of 12 (see Figure 3). Notably, as f does not contain all possible monomials, the resource overhead is significantly lower than the theoretical worst-case bound.*

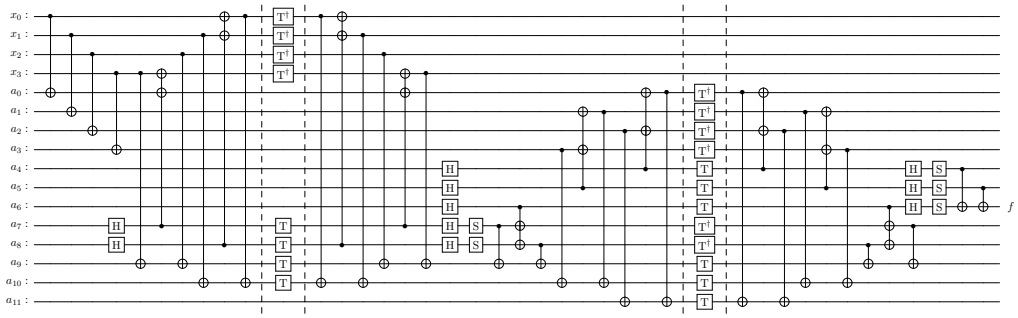

Figure 3. Quantum circuit implementing $f(x_0,x_1,x_2,x_3) = x_0x_2 \oplus x_1x_3 \oplus x_0x_1x_2x_3$ with T-depth 2.

While our construction achieves optimal T-depth, it incurs an exponential overhead in ancilla qubits and CNOT gates with respect to the number of variables. A more practical alternative is to allow a slight increase in T-depth in exchange for a substantial reduction in ancilla and CNOT counts. For instance, replacing the T-depth-1 Toffoli decomposition from [13] (Figure 1b) with the logical-AND-based decomposition from [20] (Figure 1a) in Theorem 2 reduces the ancilla count by $2^{n-1}(n-2)+1$ and the CNOT count by $3\cdot 2^{n-1}(n-2)+3$, while marginally increasing the T-depth from $\lceil\log_2 n\rceil$ to $\lceil\log_2 n\rceil+1$. This trade-off highlights a promising direction for future research, where our construction can serve as a T-depth benchmark while optimizing other quantum resources, even at the cost of slight T-depth sub-optimality, rather than prioritizing depth reduction alone.

This observation also clarifies the optimal T-depth for ANF-based implementations. In particular, for quantum circuits designed for cryptanalytic purposes, this optimal depth, viewed as a benchmark, has not been systematically explored as a performance metric. In the next section, we illustrate this gap through concrete examples, focusing on the widely studied AES block cipher. The idea naturally extends to any standard block ciphers in general.

4. Optimal T-depth Quantum Circuit of AES

This section presents an optimal T-depth quantum circuit for the AES algorithm. We first construct an optimal T-depth implementation of the AES S-box and subsequently extend the construction to the full AES algorithm for key sizes of 128, 192, and 256 bits, corresponding to 10, 12, and 14 rounds, respectively.

AES is a symmetric-key block cipher standardized by NIST, operating on 128-bit data blocks. The internal state is represented as a 4×4 matrix $S \in \mathbb{F}_{2^8}^{4\times 4}$, where each element $S_{i,j} \in \mathbb{F}_{2^8}$ corresponds to a byte.

The AES algorithm can be abstracted as a Boolean function $f : \mathbb{F}_2^{m+k} \to \mathbb{F}_2^c$, where $m = 128$, is the message length, $k \in \{128, 192, 256\}$ is the key length, and $c = 128$, is the ciphertext length. From Theorem 2, the optimal T-depth for AES, when viewed as a monolithic Boolean circuit, is given by $\lceil \log_2 256 \rceil = 8$, $\lceil \log_2 320 \rceil = 9$, and $\lceil \log_2 384 \rceil = 9$, depending on the key size. However, modeling AES as a single combinational Boolean circuit is both impractical and analytically intractable due to its iterative structure. While these T-depth values offer theoretical lower bounds, they do not serve as practical benchmarks for circuit design.

In practice, AES is implemented in a round-wise. Hence, a more feasible approach is to design the quantum circuit for each round individually and estimate the overall T-depth based on actual implementations. Each AES round (except the final one) consists of four transformations, described as follows.

- **SubBytes**: Applies the AES S-box to each element $S_{i,j} \in \mathbb{F}_{2^8}$ of the internal state. This is the only nonlinear operation and the primary contributor to the T-depth, requiring 16 parallel S-box evaluations: $S_{i,j} \leftarrow \text{S-box}(S_{i,j})$.
- **ShiftRows**: Performs a cyclic left shift on each row of the state matrix. For row i, the transformation is defined as $S_{i,j} \leftarrow S_{i,(j+i) \bmod 4}$. As this is essentially a SWAP operation, it is implemented via rewiring in the quantum circuit and does not incur additional quantum resources.
- **MixColumns**: Applies a linear transformation to each column $\mathbf{c} \in \mathbb{F}_{2^8}^4$ of the internal state S using an MDS matrix $M \in \mathbb{F}_{2^8}^{4\times 4}$: $\mathbf{c} \leftarrow M \cdot \mathbf{c}$. Since this transformation is linear, it can be realized using only CNOT gates. Note that the final round of AES omits the MixColumns operation.
- **AddRoundKey**: The AES key schedule expands the primary key K into round keys $K_0, K_1, \ldots, K_r \in \mathbb{F}_{2^8}^{4\times 4}$, each used in a specific round. In this step, the internal state S is XORed with the current round key K_i, defined as $S \leftarrow S \oplus K_i$.

Since only the SubBytes transformation introduces nonlinearity via S-boxes, specifically through multi-controlled Toffoli (MCT) gates contributing to the T-depth, the overall T-depth of an AES implementation is dominated by the S-box operations. As 16 S-boxes are evaluated in parallel per round, the total T-depth of the AES circuit can be estimated as r times the T-depth of a single S-box, where r denotes the number of rounds. We now present a detailed quantum resource estimation for the AES S-box with optimal T-depth, as derived from Theorem 2.

The coordinate Boolean functions of the AES S-box have a maximum algebraic degree of 7. By Theorem 2, an optimal T-depth quantum circuit for the AES S-box can thus be constructed with a T-depth of $\lceil \log_2 7 \rceil = 3$. Notably, the ANF of all coordinate functions collectively includes all monomials up to degree 7. We construct these using three layers of parallel Toffoli gates.

- **First Layer:** We begin by constructing all $\binom{8}{2} = 28$ degree-2 monomials. This requires 7 copies of each input variable, 56 copies in total, 8 of which already exist. Thus, 48 ancilla qubits and 48 CNOT gates are needed to copy the inputs, with a CNOT depth of 3. Storing the 28 quadratic terms requires 28 additional ancilla and 28 parallel Toffoli gates, contributing a Toffoli depth of 1. In total, the first layer uses $48 + 28 = 76$ ancilla qubits, 48 CNOT gates (CNOT depth 3), and 28 Toffoli gates.
- **Second Layer:** This layer constructs all $\binom{8}{3} = 56$ degree-3 and $\binom{8}{4} = 70$ degree-4 monomials by combining previously computed terms. To create 56 cubic terms, we combine degree-2 terms with single-variable terms. As we have 56 single-variable qubits and 28 degree-2 terms, we need to duplicate the quadratic terms using 28 ancilla and 28 CNOT gates (depth 1). Storing the 56 cubic terms adds 56 ancilla. To construct 70 quartic terms, we need five additional copies of all degree-2 monomials (140 ancilla and 140 CNOT gates), and 70 more ancilla to store the outputs. The second layer thus requires a maximuma of $(28 + 56 + 140 + 70) = 294$ ancilla, $(28 + 140) = 168$ CNOT gates (CNOT depth 3), and $(56 + 70) = 126$ Toffoli gates.
- **Third Layer:** At this stage, we have: 70 quartic, 56 cubic, 196 quadratic, and 56 single-variable qubits. We now build all degree-5, 6, and 7 monomials. The 8 septics ($\binom{8}{7} = 8$) are computed by combining eight degree-4 and eight degree-3 terms. The 56 quintics ($\binom{8}{5}$) are obtained by combining 56 degree-4 and degree-1 terms. The 28 sextics ($\binom{8}{6}$) require combinations of four degree-4 and degree-2, and 24 degree-3 and degree-3 terms. In total, constructing and storing these 92 higher-degree monomials requires $(56 + 28 + 8) = 92$ ancilla qubits and 92 Toffoli gates.

Across all eight coordinate functions, 1001 monomials appear in total, requiring 1001 CNOT gates and 8 ancilla qubits for output storage. The CNOT depth is dominated by the coordinate function with the most monomials, which is 145.

Using the T-depth-1 Toffoli-to-T decomposition (see Figure 1b), implementing 246 Toffoli gates require $(246 \times 4) = 984$ T gates and $(246 \times 9) = 2214$ CNOT gates with a CNOT depth of $(9 \times 3) = 27$. Since Toffoli gates are executed in three different layers (with a maximum of 126 in the second layer), this incurs an additional (reusable) ancilla count of 126. Although, the Toffoli gates can be uncomputed without extra cost, uncomputing intermediate qubits requires $(48 + 168) = 216$ CNOT gates.

Hence, the optimal T-depth implementation of the AES S-box requires a total $(76 + 294 + 92 + 8 + 126) = 596$ ancilla qubits, $(48 + 168 + 1001 + 2214 + 216) = 3647$ CNOT gates, with a CNOT depth of $2(3+3) + 27 + 145 = 184$, and 984 T gates, achieving the T-depth 3.

Alternatively, replacing the T-depth-1 design with Gidney's logical-AND decomposition (see Figure 1a) reduces the ancilla count by 126 and the CNOT count by $(246 \times 3) = 738$, at the cost of increasing the T-depth to $\lceil \log_2 7 \rceil + 1 = 4$. In this case, each Toffoli gate has CNOT depth 6, resulting in a total CNOT depth of $2 \times (3+3) + 145 + (3 \times 6) = 175$. Table 2 summarizes the quantum resource requirements for implementing the AES S-box.

Table 2. Quantum circuits for AES S-box with corresponding resource estimates.

Toffoli-to-T	Ancilla count	CNOT count	CNOT depth	T-count	T-depth
Using Figure 1a	470	2909	175	984	4
Using Figure 1b	596	3647	184	984	3

In the SubBytes transformation, 16 AES S-boxes are executed in parallel, resulting in a 16-fold increase in the number of ancilla qubits, CNOT gates, and T gates compared to a single S-box, while the circuit depths remain unchanged. The ShiftRows transformation introduces no additional quantum resource overhead. According to [29, Table 5], the MixColumns transformation can be implemented in-place (i.e., without additional ancilla) using 98 CNOT gates with a CNOT depth of 13. As MixColumns is applied in parallel to the four columns of the internal state, the total CNOT count becomes $4 \times 98 = 392$, while the CNOT depth remains 13. Finally, the bit-wise XOR with the round key requires 128 parallel CNOT gates, contributing a CNOT depth of 1. The overall quantum resource requirements for a single round of AES are summarized in Table 3.

Table 3. Quantum implementation of a single round of AES with corresponding resource estimates.

Toffoli-to-T	Ancilla count	CNOT count	CNOT depth	T-count	T-depth
Using Figure 1a	7520	47064	189	15744	4
Using Figure 1b	9536	58872	198	15744	3

Assume that AES operates for $r \in \{10, 12, 14\}$ rounds, depending on the key length. Ancilla qubits used in the S-boxes of each round are reclaimed for subsequent rounds through uncomputation, except for the 8 qubits required to store the functional output. As a result, the ancilla count increases by $8 \times 16 = 128$ per round, yielding a total ancilla requirement of $9536 + 128(r-1)$. Furthermore, since the final round omits the MixColumns operation, the total CNOT count is given by $58872r - 392$, and the CNOT depth is $198r - 13$. Both the T-count and T-depth scale linearly with the number of rounds.

Table 4 presents the resource estimates for complete quantum implementations of AES with different key sizes, using our optimal T-depth construction (see Figure 1b). The T-depth achieved in this construction is provably optimal (from Theorem 1), as no quantum circuit for AES, when the rounds are implemented sequentially, can attain a lower T-depth, when decomposed via Toffoli plus Clifford gates, irrespective of the number of ancilla qubits or other resource overheads.

This additional resource overhead in Table 4 demonstrates the practical viability of our optimal T-depth quantum circuit constructions for AES. While earlier implementations have reported fewer ancilla qubits, lower gate counts, and in some cases, comparable T-depths (e.g., [14, Table 7] and [16, Table VI] report a T-depth of 60 for AES and AES† combined, i.e., 30 per instance, and [19, Table 5] reports T-depths of 30, 36, and 42 for AES-128, AES-192, and AES-256, respectively), none of these works formally establish the optimality of their constructions. Given that presently AES is one of the most popular ciphers in symmetric-key cryptography, numerous heuristic and brute-force efforts

have been made to optimize its quantum implementation. The T-depth reductions observed in prior works primarily derived from those efforts, rather than from systematic constructions with provable optimality over a general class of block ciphers.

Table 4. Optimal T-depth quantum implementation of full round AES with corresponding resources.

Key size	No. of round	Ancilla count	CNOT count	CNOT depth	T-count	T-depth
128	10	10688	588328	1967	157440	30
192	12	10944	706072	2363	188928	36
256	14	11200	823816	2759	220416	42

In this context, we acknowledge the recent work of Huang et al. [18], which also reports a T-depth 3 implementation of the AES S-box (see [18, Table 3]) and establishes its minimality, similar in spirit to our own findings. However, we emphasize that our work is independent and presents a generic quantum circuit construction method that achieves minimal the T-depth for any arbitrary Boolean function derived from its Algebraic Normal Form. This construction provides a naive yet complete upper bound on ancillary quantum resources, thereby setting new benchmarks applicable to a broad class of S-boxes beyond AES. Furthermore, our approach is a generalization of the optimal T-depth MCT decomposition (via Clifford plus Toffoli gates), as outlined in [12, Corollary 1]). Consequently, optimization strategies such as logical-AND and conditionally clean ancilla techniques can also be integrated to reduce quantum resource overheads significantly, at the cost of a marginal increase in T-depth beyond the optimal bound.

It is important to note that a T-depth 3 implementation of the AES S-box does not necessarily imply the optimality for the full-round AES circuit. Algebraic manipulations across multiple rounds may reduce the combined T-depth below the additive bound and must be analyzed individually for each block cipher. In contrast, we propose a generic, step-by-step construction framework that guarantees optimal T-depth for a broad class of block ciphers beyond AES, establishing a definitive benchmark for quantum cryptographic implementations.

For reference, Table 5 summarizes prior AES implementations focused on T-depth reduction, along with associated resource estimates.

Table 5. Optimal T-depth quantum implementation of full round AES: A comparison with prior results.

Key size	References	Ancilla count	CNOT count	CNOT depth	T-count	T-depth
128	[13, Table 4]	4244	284420	NA	54400	120
128	[30, Table 2]	9384	NA	NA	33600	50
128	[15, Table 13]	3689	132376	NA	27200	40
128	[14, Table 7]	5576	285393	NA	62400	30
128	[16, Table VI]	NA	228020	NA	52800	30
128	[18, Table 8]	NA	176580	NA	33600	30
128	[19, Table 5]	6128	120812	NA	117984	30
128	**[This paper]**	10688	588328	1967	157440	30
192	[13, Table 4]	4564	321021	NA	60928	144
192	[30, Table 2]	10456	NA	NA	37632	60
192	[15, Table 13]	3945	149256	NA	30464	48
192	[19, Table 5]	6448	136812	NA	132960	36
192	**[This paper]**	10944	706072	2363	188928	36
256	[13, Table 4]	4884	393534	NA	75072	168
256	[30, Table 2]	46368	NA	NA	12704	70
256	[15, Table 13]	4457	187128	NA	38080	56
256	[19, Table 5]	6768	168548	NA	165264	42
256	**[This paper]**	11200	823816	2759	220416	42

5. Cryptanalytic Implications

We know that the Shor's factoring algorithm [5] poses the most critical threat to existing classical public key cryptographic protocols due to its implications in attacking the RSA and the discrete logarithm-based schemes. On the other hand, the impact of Grover's algorithm [3] can be temporarily mitigated by doubling the key size. Still there exist numerous instances where the Grover's search, often in conjunction with other quantum algorithms, has demonstrated significant cryptanalytic potential. For example, in [31], Grover's algorithm is combined with Simon's hidden shift technique to break the security of Even-Mansour constructions. Needless to mention that the resource requirements for a full-scale key recovery attack on symmetric ciphers using Grover's algorithm remain infeasible in the near term. Still, substantial work has been conducted on Grover-based cryptanalysis, particularly on AES [13–19]. In this context, we present the following result on the optimality of T-depth required for full-round cryptanalysis using Grover's search on a block cipher **B**, where the only source of nonlinearity generates from a single S-box.

Suppose **B** is a block cipher whose only nonlinearity arises from an n-bit S-box **S**. Then, from Theorem 2, implementing the n-bit S-box **S** requires an optimal T-depth of $\lceil \log_2 n \rceil$. Since the S-box is the sole source of nonlinearity in the cipher, each round incurs an optimal T-depth of $\lceil \log_2 n \rceil$. Assuming the cipher runs for r rounds and produces a ciphertext of length m, the quantum implementation of the full-round cipher **B** has a total T-depth of $r\lceil \log_2 n \rceil$.

To perform an exhaustive key search using Grover's algorithm, one must implement **B**, compare the output with a known m-bit ciphertext using an m-MCT gate (which requires an optimal T-depth of $\lceil \log_2 m \rceil$), and then apply the inverse of the full-round cipher, $\mathbf{B}^\dagger$, which adds another $r\lceil \log_2 n \rceil$ to the T-depth. Therefore, a single Grover iteration requires an optimal T-depth of $2r\lceil \log_2 n \rceil + \lceil \log_2 m \rceil$. Since Grover's algorithm requires $2^{k/2}$ iterations for a key of length k, the exhaustive key recovery attack using sequential Grover's search requires an overall T-depth of $(2r\lceil \log_2 n \rceil + \lceil \log_2 m \rceil)2^{k/2}$.

As an immediate corollary, an exhaustive key recovery attack on AES using Grover's algorithm can be executed with an optimal T-depth of $(2r\lceil \log_2 8 \rceil + \lceil \log_2 128 \rceil)2^{k/2}$, which evaluates to 67×2^{64}, 79×2^{96}, and 91×2^{128} for key lengths of 128, 192, and 256 bits, respectively. In this regard, we make the following remarks based on different execution scenarios of Grover's algorithm.

Remark 1.

- *If the block cipher **B** is implemented out-of-place, then the inverse full-round cipher can be realized via measurement-based uncomputation, which does not incur any additional T-depth. Consequently, Grover's search can be executed with an optimal T-depth of* $(r\lceil \log_2 n \rceil + \lceil \log_2 m \rceil)2^{k/2}$. *The corresponding T-depth requirements for AES would be* 37×2^{64}, 43×2^{96}, *and* 49×2^{128} *for key lengths of 128, 192, and 256 bits, respectively.*
- *Furthermore, if Grover's search is executed in parallel, i.e., all the* 2^k *copies of the oracle* **B** *is running concurrently, the T-depth for the parallel Grover's search becomes* $(r\lceil \log_2 n \rceil + \lceil \log_2 m \rceil)$. *Accordingly, for AES, the corresponding minimal T-depth would be 37, 43, and 49 for key lengths of 128, 192, and 256 bits, respectively.*

6. Conclusion

Given the ANF of an arbitrary n-input, m-output Boolean function f having algebraic degree k, this work presents the construction of an optimal T-depth quantum circuit for f, along with a complete resource estimation. The primary focus of this paper is on minimizing T-depth, an essential metric in quantum circuit design due to its direct impact on circuit latency and coherence time. While the overall resource usage may be high, we argue that establishing the benchmark T-depth should be the first step in any quantum circuit synthesis process, after which other resource parameters may be optimized. This work conclusively settles the minimum achievable T-depth for Boolean functions and demonstrates its practical relevance through block cipher constructions, as well as in cryptanalysis.

Author Contributions
All authors have contributed equally. All authors have read and agreed to the published version of the manuscript.

Funding
No funding was received to assist with the preparation of this manuscript.

Conflicts of Interest Statement
The authors declare no conflicts of interest relevant to the content of this article.

Data Availability Statement
No data were created or analyzed in this study.

Acknowledgments
We sincerely thank the editor and the reviewers for their meticulous reviews, valuable insights, and constructive suggestions, which have substantially improved both the technical and editorial quality of this paper. A.C. acknowledges support through MoE AcRF Tier 1 Award No. RT10/23. S.M. acknowledges financial support from the Information Security Education and Awareness (ISEA) Project Phase-III, Cluster Cryptography, an initiative of MeitY, Grant No. L-14017/1/2022-HRD.

References

1. R. P. Feynman (1986). Quantum mechanical computers. *Foundations of Physics*, 16, pp. 507–531.
2. D. Deutsch and R. Jozsa (1992). Rapid solution of problems by quantum computation, Proceedings of the Royal Society A: Mathematical. *Physical and Engineering Sciences*, 439 (1907), pp. 553–558.
3. L. K. Grover (1996). A fast quantum mechanical algorithm for database search, *Proceedings of the 28th Annual ACM Symposium on Theory of Computing (STOC '96)*, pp. 212–219.
4. D. R. Simon (1997). On the power of quantum computation, *SIAM Journal on Computing*, 26(5), pp. 1474–1483.
5. P. W. Shor (1997). Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5), pp. 1484–1509.
6. C. Jones (2013). Low-overhead constructions for the fault-tolerant Toffoli gate. *Phys. Rev. A*, 87(2), 022328.
7. S. Bravyi and A. Kitaev (2005). Universal quantum computation with ideal Clifford gates and noisy ancillas. *Phys. Rev. A*, 71(2), 022316.
8. P. Selinger (2013). Quantum circuits of T-depth one. *Phys. Rev. A*, 87(4), 042302.
9. M. Amy, D. Maslov, M. Mosca and M. Roetteler (2013). A meet-in-the-middle algorithm for fast synthesis of depth-optimal quantum circuits, *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 32(6), pp. 818–830.
10. C. Gidney and M. Ekera (2021). How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits. *Quantum*, 5, 433.
11. C. Gidney (2025). How to factor 2048 bit RSA integers with less than a million noisy qubits, arXiv: 2505.15917.
12. S. Dutta, S. Wang, A. Baksi, A. Chattopadhyay and S. Maitra (2025). Exact space-depth trade-offs in multicontrolled Toffoli decomposition, *Phys. Rev. A*, 111, 052611.
13. S. Jaques, M. Naehrig, M. Roetteler and F. Virdia (2020). Implementing Grover oracles for quantum key search on AES and LowMC, Advances in Cryptology – EUROCRYPT 2020. *LNCS*, vol. 12106, pp. 280–310.
14. Z. Huang and S. Sun (2022). Synthesizing quantum circuits of AES with lower T-depth and less qubits, Advances in Cryptology – ASIACRYPT 2022. *LNCS*, vol. 13793, pp. 614–644.
15. Q. Liu, B. Preneel, Z. Zhao and M. Wang (2023). Improved quantum circuits for AES: reducing the depth and the number of qubits, Advances in Cryptology – ASIACRYPT 2023. *LNCS*, vol. 14440, pp. 67–98.
16. M. Zhang, T. Shi, W. Wu and H. Sui (2024). Optimized quantum circuit of AES with interlacing-uncompute structure. *IEEE Transactions on Computers*, vol. 73, no. 11, pp. 2563–2575.
17. H. Shi and X. Feng (2024). Quantum circuits of AES with a low-depth linear layer and a new structure, Advances in Cryptology – ASIACRYPT 2024. *LNCS*, vol. 15491, pp. 358–395.
18. Z. Huang, F. Zhang and D. Lin (2025). Constructing quantum implementations with the minimal T-depth or minimal width and their applications, Advances in Cryptology - EUROCRYPT 2025. *LNCS*, vol. 15601, pp. 155–185.
19. K. Jang, A. Baksi, H. Kim, G. Song, H. Seo and A. Chattopadhyay (2025). Quantum analysis of AES, *IACR Communications in Cryptology*, 2(1).
20. C. Gidney (2018). Halving the cost of quantum addition. *Quantum*, 2, pp. 74–79.
21. C. Gidney and A. G. Fowler (2019). *Flexible layout of surface code computations using autoCCZ states*, arXiv: 1905.08916.
22. M. R. Albrecht, C. Rechberger, T. Schneider, T. Tiessen and M. Zohner (2015). Ciphers for MPC and FHE, Advances in Cryptology - EUROCRYPT 2015. *LNCS*, vol. 9056, pp. 430–454.
23. A. Baksi, S. Bhasin, J. Breier, M. Khairallah, T. Peyrin, S. Sarkar and S. M. Sim (2021). DEFAULT: cipher level resistance against differential fault attack, Advances in Cryptology – ASIACRYPT 2021. *LNCS*, vol. 13091, pp. 124–156.

24. S. Banik, S. K. Pandey, T. Peyrin, Y. Sasaki, S. M. Sim and Y. Todo (2017). GIFT: a small present, CHES 2017. *LNCS*, vol. 10529, pp. 321–345.
25. A. Bogdanov, L. R. Knudsen, G. Leander, C. Paar, A. Poschmann, M. J. B. Robshaw, *et al.* (2007). PRESENT: an ultra-lightweight block cipher, CHES 2007. *LNCS*, vol. 4727, pp. 450–466.
26. J. Borghoff, A. Canteaut, T. Güneysu, E. B. Kavun, M. Knezevic, L. R. Knudsen, *et al.* (2012). PRINCE – a low-latency block cipher for pervasive computing applications, Advances in Cryptology – ASIACRYPT 2012. *LNCS*, vol. 7658, pp. 208–225.
27. C. Dobraunig, M. Eichlseder, F. Mendel and M. Schläffer (2021). Ascon v1.2: lightweight authenticated encryption and hashing. *J Cryptol*, 34, 33.
28. J. Daemen and V. Rijmen (2002). *The design of Rijndael*, Springer, Information Security and Cryptography.
29. Y. Yuan, W. Wu, T. Shi, L. Zhang and Y. Zhang (2024). A framework to improve the implementations of linear layers, *IACR Transactions on Symmetric Cryptology*, 2024(2), pp. 322–347.
30. T. Häner and M. Soeken (2022). Lowering the T-depth of quantum circuits via logic network optimization. *ACM Transactions on Quantum Computing*, 3(2), art. no. 6, pp. 1–15.
31. G. Leander and A. May (2017). Grover meets Simon – quantumly attacking the FX-construction, Advances in Cryptology – ASIACRYPT 2017. *LNCS*, vol. 10625, pp. 161–178.