

Time-Space Complexity of Quantum Search Algorithms in Symmetric Cryptanalysis

Panjin Kim, Kyung Chul Jeong^(✉), and Daewan Han

National Security Research Institute, Daejeon 34044, Korea
jeongkc@nsr.re.kr

Abstract. Performance of cryptanalytic quantum search algorithms is mainly inferred from *query* complexity which hides overhead induced by an implementation. To shed light on quantitative complexity analysis removing hidden factors, we provide a framework for estimating time-space complexity, with carefully accounting for characteristics of target cryptographic functions. Processor and circuit parallelization methods are taken into account, resulting in the time-space trade-offs curves in terms of *depth* and *qubit*. The method guides how to rank different circuit designs in order of their efficiency. The framework is applied to representative cryptosystems NIST referred to as a guideline for security parameters, reassessing the security strengths of AES and SHA-2.

Keywords: Post-quantum cryptography (PQC), Circuit model, Grover, Resource estimates, Time-space trade-offs, AES, SHA-2

1 Introduction

Quantum cryptanalysis is an area of study that has long been developed alongside the field of quantum computing, as many cryptosystems are expected to be directly affected by quantum algorithms. One of the quantum algorithms that would have an impact on symmetric cryptosystems is Grover's algorithm [28]. It had been widely known that many symmetric cryptosystem's security levels will be simply reduced by half due to the asymptotic behavior of the query complexity of Grover's algorithm under the oracle assumption. As the field has matured over decades, not mere asymptotic but more quantitative approaches to the cryptanalysis are also being considered recently [4,27,43,46,47]. These works have substantially improved the understanding of quantum attacks by systematically estimating quantum resources. Nevertheless, it is still noticeable that the existing works on resource estimates are more intended for suggesting exemplary quantum circuits (so that one can count the number of required gates and qubits explicitly) than fine-tuning of actual attack designs.

The importance of estimating costs of quantum search algorithms beyond pioneering works should be emphasized as it can be utilized to suggest practical security levels in the post-quantum era. NIST indeed suggested security levels based on the resistances of AES and SHA to quantum attacks in the PQC standardization call for proposals document [41]. In addition, the difficulty of

measuring the complexity of quantum attacks was questioned in the first NIST PQC standardization workshop [37]. The main purpose of this work is to formulate the time-space complexity of quantum search algorithm in order to provide reliable quantum security strengths of classical symmetric cryptosystems.

1.1 This work

There exist two noteworthy points overlooked in the previous works. First, the target function to be inverted is generally a pseudo-random function or a cryptographic hash function. Under the characteristics of such functions, bijective correspondence between input and output is not guaranteed. This makes Grover's algorithm *seemingly* inadequate due to the unpredictability of the number of targets. The second point is time-space trade-offs of quantum resources. Earlier works on quantitative resource estimates have implicitly or explicitly assumed a single quantum processor. Presuming that the *resource* in classical estimates includes the number of processors the adversary is equipped with, the single processor assumption is something that should be revised.

Being aware of the issues, we come up with a framework for analyzing the time-space complexity of cryptanalytic quantum search algorithms. The main consequences we presented in this paper are three folded:

Precise query complexity involving parallelization. The number of oracle queries, or equivalently Grover iterations, is first estimated as exactly as possible, reasonably accounting for previously overlooked points. Random statistics of the target function are carefully handled which lead to increase in iteration number compared with the case of unique target. Surprisingly, however, the cost of dealing with random statistics in this paper is not expensive compared with the previous work [27] under single processor assumption. Furthermore, when processor parallelization is considered, we observed that this extra cost gets even more negligible. It is also interesting to recognize that the parallelization methods could vary depending on the search problems. After taking the asymptotical big O notation off, the relation between time and space in terms of Grover iterations and number of processors, called *trade-offs curve*, is obtained. Apart from resource estimates, investigating the trade-offs curve of state-of-the-art collision algorithm in [18] with optimized parameters is one of our major concerns.

Qubit-depth trade-offs and circuit design tuning. In the next stage, time and space resources are defined in a way that they can be interpreted as physical quantities. Cost of quantum circuits for cryptanalytic algorithms can be estimated in units of *logical qubits* and *Toffoli-depths*. Taking the total number of gates as time complexity disturbs accurate estimates for the speed of quantum algorithms due to far different overheads introduced by various gates in real operation. With the definitions of quantum resources, the trade-offs curve now describes the relation between number of qubits and circuit depths. Since we are given a 'relation' between time and space, it is then possible to grade the various

quantum circuits in order of efficiency. In other words, the method described so far enables one to tell which attack design is more cost-effective.

By applying generic methodology newly introduced, time-space complexities of AES and SHA-2 against quantum attacks are measured in the following way. Various designs are constructed by assembling different circuit components with options such as reduced depth at the cost of the increase in qubits (or vice versa). Design candidates are then subjected to the trade-offs relation for comparison. The trade-offs coefficient of the most efficient design represents the hardness of quantum cryptanalysis. Compared with pre-existing circuit designs, we have improved the circuits by reducing required qubits and/or depths in various ways. However, we do not claim that we have found the optimal attacks for AES and SHA-2. The method enables us to select the ‘best’ one out of candidates at hand. It is remarked that the explicit circuit designs for quantum collision search algorithms is first introduced.

Revisiting the security levels of NIST PQC standardization. The procedure is applied to each primitive of security strength categories NIST specified in [41]. A new threshold that is required for the category classification, based on the cost metric proposed in this work, is provided in Fig. 1. It includes wide range of parameters and the quantum collision finding algorithms which do not outperform classical counterparts to explicitly recognize the quantum-side complexities of all the categories.

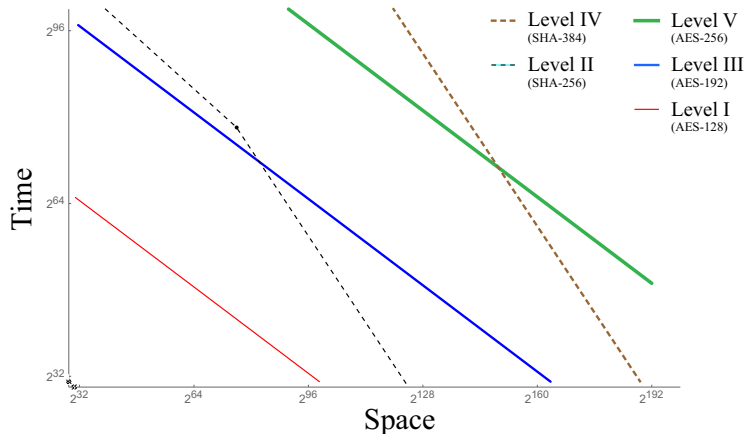


Fig. 1. Time-space cost, which has conditional ordering, of the quantum attacks on five security strength category representatives of NIST PQC standardization

We end this subsection with an important caveat. Use of classical resources appears in this paper, but we do not handle the complexity induced by it because of unclear comparison criteria for quantum and classical resources.

1.2 Organization

Next section covers the backgrounds including Grover's algorithm and its variants. In Sect. 3, the time-space complexity of relevant algorithms in a unit of Grover iteration is investigated. A basic unit of quantum computation is proposed in Sect. 4. Sections 5, 6 and 7 show the results of applying the time-space analysis to AES and SHA-2. In Sect. 8, based on the observations made in the previous sections, a comprehensive figure summarizing the quantum security strengths of AES and SHA-2 is drawn. Section 9 summarizes the paper.

2 Backgrounds

Grover's algorithm, the success probability, parallelization methods, and some generalizations or variants are explained briefly. A brief review of AES and SHA-2, and an introduction to related works on resource estimates are followed. We do not cover the basics of quantum computing, but leave the references [7,38] for interested readers. In this paper, every bra or ket state is normalized.

2.1 Grover's Algorithm

Throughout the paper, the target function is denoted by f and $N = 2^n$ for some $n \in \mathbb{N}$. Consider a set X of size N and a function $f: X \rightarrow \{0, 1\}$,

$$f(x) = \begin{cases} 1, & \text{if } x \in T, \\ 0, & \text{otherwise,} \end{cases}$$

where T of size t is a set of targets to be found.

Grover's algorithm [28] is an algorithm that repeatedly applies an operator $Q = -AS_0A^{-1}S_f$, called *Grover iteration*, to the initial state $|\Psi\rangle = A|0\rangle$, where

$$A = H^{\otimes n}, \quad S_0 = I - 2|0\rangle\langle 0|, \quad S_f = I - 2|\tau\rangle\langle \tau|, \quad (1)$$

where $H^{\otimes n}$ is a set of Hadamard operators and $|\tau\rangle$ is a target state which is an equal-phase and equal-weight superposition of $|x\rangle$ for all $x \in T$. The roles of S_0 and S_f are to swap the sign of zero and $|\tau\rangle$ states, respectively.

The operators S_f and $-AS_0A^{-1}$ are known as *oracle* and *diffusion* operators, respectively. By acting the oracle operator on a state, only the target state is marked through the sign change. The diffusion operator flips amplitudes around the average.

Success probability of measurement as a function of the number of iterations has been studied in [13], observing the optimal number of iterations that minimizes the ratio of the iterations to success rate. We introduce the results below with notation that is used throughout the paper.

By applying Q on the initial state i -times, the success probability of measuring one of the t solutions in the domain of size N , denoted by $p_{t,N}(i)$, becomes

$$p_{t,N}(i) = \sin^2 [(2i + 1)\theta_{t,N}], \quad (2)$$

where $\sin(\theta_{t,N}) = \sqrt{t/N}$ ($= \langle \Psi | \tau \rangle$). The number of repetitions of Q maximizing the success probability of measurement, denoted by $I_{t,N}^{\text{mp}}$, is estimated as

$$I_{t,N}^{\text{mp}} = (\pi/4) \cdot \sqrt{N/t} .^1 \quad (3)$$

When the measurement is made after i -repetitions of Q , the expected number of Grover iterations to find one of the targets can be expressed as a function of i . For t targets in the domain of size N , the function is denoted by $I_{t,N}(i)$ which reads $I_{t,N}(i) = i/p_{t,N}(i)$. The optimal number of iterations $i_{t,N}$ that minimizes $i/p_{t,N}(i)$ is found to be $i_{t,N} = 0.583 \dots \cdot \sqrt{N/t}$, and then the expected number of iterations, denoted by $I_{t,N}$ reads

$$I_{t,N} = I_{t,N}(i_{t,N}) = 0.690 \dots \cdot \sqrt{N/t} . \quad (4)$$

In some cases where the domain size is N , it is omitted such as $p_t(i)$ ($= p_{t,N}(i)$) or I_t ($= I_{t,N}$), for readability.

Parallelization of Grover's algorithm using multiple quantum computers has been investigated in applications to cryptanalysis [8,41,53]. Consideration of parallelization in a hybrid algorithm can be found in [9]. Asymptotically the execution time is reduced by a factor of the square root of the number of quantum computers. There are two straightforward parallelization methods having such property, called *inner* and *outer* parallelization.

Here we fix some notations for parameters related to parallelization. T_q and S_q stand for the number of sequential Grover iterations and the number of quantum computers, respectively. S_c stands for the amount of classical resources, such as the size of storage and/or the number of processors. Definitions of two parallelization methods can be given as follows.

Definition 1 (Inner Parallelization (IP)). *After dividing the entire search space into S_q disjoint sets, each machine searches one of the S_q sets for the target. The number of iterations can be reduced due to the smaller domain size.*

Definition 2 (Outer Parallelization (OP)). *Copies of Grover's algorithm on the entire search space is executed in S_q machines. Since it is successful if any of S_q machines finds the target, the number of iterations can be reduced.*

Parallelization is inevitable once the notion of MAXDEPTH [41] is applied.

2.2 Generalizations and Variants

Fixed-point [52] and quantum amplitude amplification (QAA) [14] algorithms are generalizations of Grover's algorithm. A brief review of QAA is given in this subsection which appears as a component of a collision finding algorithm

¹ Rounding function is not explicitly used in this paper for simplicity.

in later sections. We skip over the fixed-point algorithm as it has no advantage over Grover’s algorithm and QAA in this work.²

There exist a number of variants of Grover’s algorithm in application to collision finding. In [15], Brassard, Høyer and Tapp suggested a quantum collision finding algorithm (BHT) of $O(N^{1/3})$ query complexity using *quantum memory* amounting to $O(N^{1/3})$ classical data. A multi-collision algorithm using BHT was suggested in [30]. In this work however, we do not consider BHT as a candidate algorithm for the following reasons. One is that the algorithm entails a need for quantum memory where the realization and the usage cost are controversial [5], and the other is that we are unable to come up with any implementation restricted to use of elementary gates that do not exceed the total cost of $O(N^{1/2})$.

Apart from quantum circuits, algorithms primarily designed for other type of models such as measurement-based quantum computation also exist, for example quantum walk search [48,19] or element distinctness [1], but we do not cover them as state-of-the-art quantum architecture is targeting for circuit computation. Interested readers may further refer to [30] and related references therein for more information on quantum collision finding.

Bernstein analyzed quantum and classical collision finding algorithms in [10]. Quoting the work, no quantum algorithm with better time-space product complexity than $O(N^{1/2})$ which is achieved by the state-of-the-art classical algorithm [42] had not been reported. If Grover’s algorithm is parallelized with the distinguished point method, complexity of $O(N^{1/2})$ can be achieved. This is one of the examples of *immediate ways* to combine quantum search with the rho method as mentioned in [10]. We denote it as Grover with distinguished point (GwDP) algorithm in this paper.

In Asiacrypt 2017, Chailloux, Naya-Plasencia and Schrottenloher suggested a new quantum collision finding algorithm, called CNS algorithm, of $O(N^{2/5})$ query complexity using $O(N^{1/5})$ classical memory [18].

QAA algorithm. Basic structure of QAA is the same as Grover’s original algorithm. Initial state $|\Psi\rangle = A|0\rangle$ is prepared, and then Grover iteration Q is repeatedly applied i times to get success probability (2). The only difference is that in QAA, the preparation operator A is not restricted to $H^{\otimes n}$ where $N = 2^n$, and so thus the search space can be arbitrarily defined. Detailed derivation is not covered here, but instead we describe the key feature in an example.

As a trivial example, let us assume we are given a quantum computer, and try to find a target bit-string 110011 in a set $N = \{x \mid x \in \{0,1\}^6 \text{ and two middle bits are } 0\}$. Domain size is not equal to 2^6 , and the initial state can be prepared by $A = H_1 H_2 H_5 H_6$ where H_r is Hadamard gate acting on r -th qubit. Remaining processes are to apply Grover iterations $Q = -AS_0A^{-1}S_f$ with A given by the

² There are two reasons. One is that Fixed-point search requires *two* oracle queries per iteration, and the other is $\log(2/\delta)$ factor in (3) in [52] which also increases the required number of iterations depending on the bounding parameter δ . Comparing these factors with the overhead in our method introduced by random statistics, we concluded that the fixed-point algorithm is not favored.

state preparation operator just mentioned. The search space examined is rather trivial, but QAA also works on arbitrary domain. Non-trivial domain can be given as something like $N = \{x \mid x \in \{0, 1\}^6, f(x) \neq 0\}$ for some given function f . It is a matter of preparing a state encoding appropriate search space, or in other words, that is to find an operator A . Once A is constructed, QAA works in the same way as in Grover's algorithm.

GwDP algorithm. GwDP algorithm is a parallelization of Grover's algorithm. Distinguished point (DP) can be defined by a function output whose d most significant bits are zeros, denoted by d -bit DP. We allow the notation DP to indicate a pair of DP and corresponding input or an input by itself.

For $S_q = S_c = 2^s$, use $(n - 2s)$ -bit DP. By running $T_q = O(2^{n/2-s})$ times of Grover iterations, DP is expected to be found on each machine. Storing $O(2^s)$ DPs sorted according to the output, a collision is found with high probability. The time-space product is always $T_q S_q = O(N^{1/2})$.

CNS algorithm. CNS algorithm consists of two phases, the list preparation and the collision finding. In the list preparation phase, a list of size 2^l of d -bit DPs is drawn up with the time complexity of $O(2^{l+d/2})$ and the classical storage of size $O(2^l)$. In the collision finding phase QAA algorithm is used. Each iteration of QAA algorithm consists of $O(2^{d/2})$ Grover iterations and $O(2^l)$ operations for the list comparison. After $O(2^{(n-d-l)/2})$ QAA iterations, a collision is expected to be found. In total, CNS algorithm has $O(2^{l+d/2} + 2^{(n-d-l)/2}(2^{d/2} + 2^l))$ time complexity and uses $O(2^l)$ classical memory. With the optimal parameters $l = d/2$ and $d = 2n/5$, a collision is found in $T_q = O(N^{2/5})$ with $S_c = O(N^{1/5})$.

If $S_q = 2^s$, time complexity becomes $O(2^{(n-d-l-s)/2}(2^{d/2} + 2^l) + 2^{l+d/2-s})$ for $s \leq \min(l, n - d - l)$. When $l = d/2$ and $d = 2/5\{n + s\}$, the complexities satisfy $(T_q)^5(S_q)^3 = O(N^2)$ and $T_q(S_c)^3 = O(N)$.

2.3 AES and SHA-2 algorithms

A brief review of AES and SHA-2 are given in this subsection. Specifically, AES-128 and SHA-256 algorithms are described which will form the main body of later sections.

AES-128. Only the encryption procedure of AES-128 which is relevant to this work will be shortly reviewed. See [39] for details.

Round. AES round consists of four elementary operations; SubBytes, ShiftRows, MixColumns, and AddRoundKey³. Each operation applies to internal state, which is represented by 4×4 array of bytes $S_{i,j}$, as shown in Fig. 2(a).

³ The first and the last rounds are different, but will not be covered in detail here.

- ShiftRows does cyclic shifts of the last three rows of the internal state by different offsets.
- MixColumns does a linear transformation on each column of the internal state that mixes the data.
- AddRoundKey does an addition of the internal state and the round key by an XOR operation.
- SubBytes does a non-linear transformation on each byte. SubBytes works as substitution-boxes (S-box) generated by computing a multiplicative inverse, followed by a linear transformation and an addition of S-box constant.

Key Schedule. AES key schedule consists of four operations; RotWord, SubWord, Rcon, and addition by XOR operation. The sequence of key scheduling is described in Fig. 2(b). Each operation applies to 32-bit word w_i , which is represented by 4×1 array of bytes $k_{i,j}^d$. First four words are given by original key which become the zeroth round key. More words— 40 in AES-128— are then generated by recursively processing previous words. Every Sixteen byte $k_{i,j}^d$ constitutes d -th round key. RotWord, SubWord, and Rcon only apply to every fourth word w_i , $i \in \{3, 7, 11, \dots, 39\}$.

- RotWord does a cyclic shift on four bytes.
- Rcon does an addition of the constant and the word by XOR operation.
- SubWord does an S-box operation on each byte in word.

SHA-256. For brevity, only SHA-256 hashing algorithm for one message block which is relevant to this work will be reviewed. Description of preprocessing including message padding, parsing, and setting initial hash value is also omitted here. See [40] for details.

Round. SHA-2 round consists of five operations; Ch , Maj , Σ_0 , Σ_1 , and addition modulo 2^{32} . Round operations apply on eight 32-bit working variables denoted by a, b, c, d, e, f, g, h . See Fig. 3(a) for procedures.

- $Ch(x, y, z) = (x \wedge y) \oplus (\neg x \wedge z)$,
 - $Maj(x, y, z) = (x \wedge y) \oplus (x \wedge z) \oplus (y \wedge z)$,
 - $\Sigma_0(x) = ROTR^2(x) \oplus ROTR^{13}(x) \oplus ROTR^{22}(x)$,
 - $\Sigma_1(x) = ROTR^6(x) \oplus ROTR^{11}(x) \oplus ROTR^{25}(x)$,
- where $ROTR^n(x)$ is circular right shift of x by n positions.

Message Schedule. SHA-2 message schedule consists of three operations; σ_0 , σ_1 , and addition modulo 2^{32} . The sequence of message scheduling is described in Fig. 3(b). Each operation applies to 32-bit word W_i . First 16 words are given by original message block which become the first 16 words fed to SHA-256 rounds. More words— 48 in SHA-256— are then generated by recursively processing previous words.

- $\sigma_0(x) = ROTR^7(x) \oplus ROTR^{18}(x) \oplus SHR^3(x)$,
 - $\sigma_1(x) = ROTR^{17}(x) \oplus ROTR^{19}(x) \oplus SHR^{10}(x)$,
- where $SHR^n(x)$ is right shift of x by n positions.

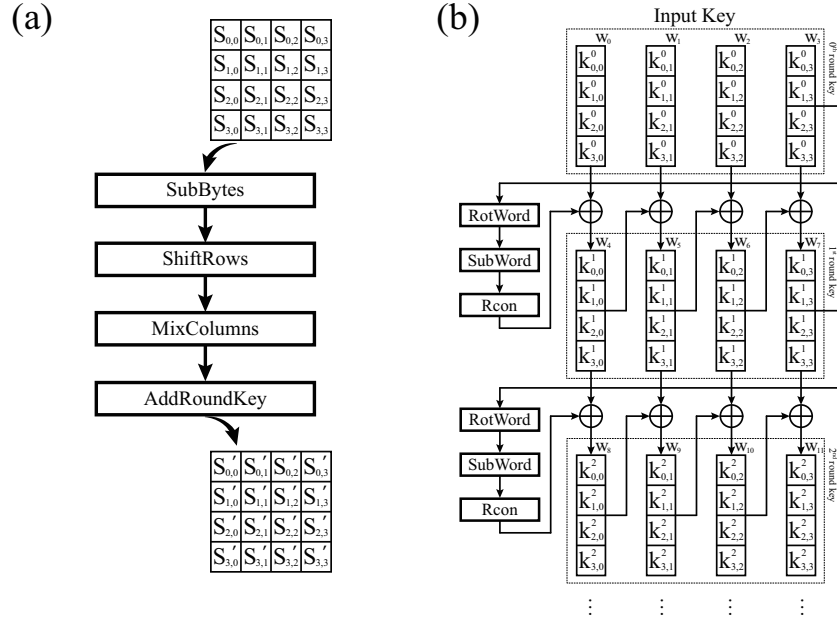


Fig. 2. (a) Round operations and (b) key schedule of AES-128 algorithm. Each square box accommodates one byte. In key schedule, 128-bit key is divided into four 32-bit words.

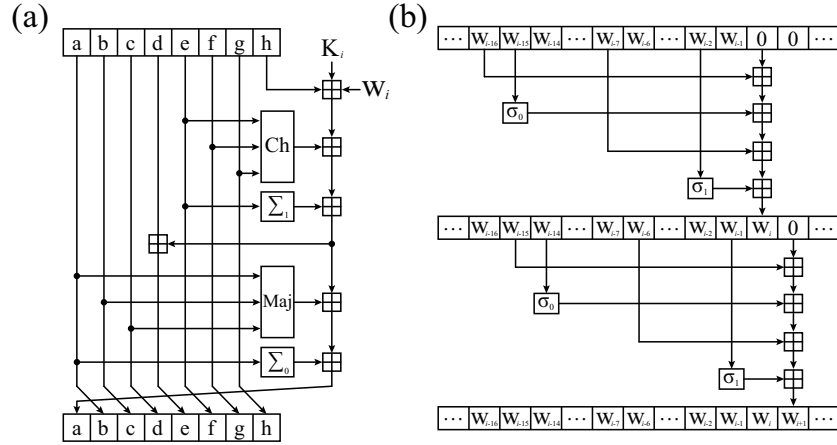


Fig. 3. (a) Round operations and (b) message schedule of SHA-256 algorithm. Each square box accommodates 32-bit word. The symbol \boxplus is addition modulo 2^{32} . Note that new word can be overwritten on existing word if the word has already been fed to round.

2.4 Quantum Resource Estimates

Quantum resource estimates of Shor’s period finding algorithm have long been studied in the various literature. See for example [45,46] and referenced materials therein. On the other hand, quantitative quantum analysis on cryptographic schemes other than period finding is still in its early stage. Partial list may include attacks on multivariate-quadratic problems [47], hash functions [4,43], and AES [27]. We introduce two of them which are the most relevant to our work.

AES Key Search. Grassl et al. reported the quantum costs of AES- k key search for $k \in \{128, 192, 256\}$ in the units of logical qubit and gate [27]. In estimating the time cost, the author’s focus was put on a specific gate called ‘T’ gate and its depth, although the overall gate count was also provided. Space cost was simply estimated as the total number of qubits required to run Grover’s algorithm.

There are two points we pay attention on. First is that the authors ensured a single target key. Since AES algorithm works like a random function, there is non-negligible probability that a plaintext end up with the same ciphertext when encrypted by two different keys. To avoid the cases, the authors encrypt r ($\in \{3, 4, 5\}$) plaintext blocks simultaneously to obtain r ciphertexts so that only the true key results in given ciphertexts. The procedure removes the ambiguity in the number of iterations. Note, however, that the removal of the ambiguity comes in exchange of at least tripling the space cost. The other point is that reversible circuit implementation of internal functions of AES was always aimed at reducing the number of qubits. One may see proposed circuit design as space-optimized.

SHA-2 and SHA-3 Pre-Image Search. Amy et al. reported the quantum costs of SHA-2 and SHA-3 pre-image search in the units of logical and physical qubit and gate [4]. The method considers an error-correction scheme called surface code. Time cost was set considering the scheme. Estimating the costs of T gates in terms of physical resources was one of the main results. One point we would like to address in the work is that random-like behavior of SHA function was not considered. It is assumed in the paper that the unique pre-image of a given hash exists.

3 Trade-offs in Query Complexity

The definitions of search problems involving random function are given. For each problem, the expected iteration number of the corresponding quantum search algorithm is calculated. Finally, the trade-offs equations between the number of iterations and the number of quantum machines are given.

3.1 Types of Search Problems

We assume that $f: X \rightarrow Y$ is random function which means f is selected from the set of all functions from X to Y uniformly at random. Useful statistics of random function can be found in [23]. The probabilities related to the number of pre-images are quoted below. When an element x is selected from a set X uniformly at random, it is denoted by $x \stackrel{\$}{\leftarrow} X$.

When $|Y| = N$ and $|X| = aN \in \mathbb{N}$ for some $a \in \mathbb{Q}$, an element $y \in Y$ is called a j -node if it has j pre-images, i.e., $|\{x \in X : f(x) = y\}| = j$. For $y \stackrel{\$}{\leftarrow} Y$, the probability of y to be a j -node, denoted by $q_{(aN)}(j)(j \geq 0)$, is approximately

$$q_{(aN)}(j) = \frac{1}{e^a} \cdot \frac{a^j}{j!} . \quad (5)$$

For $x \stackrel{\$}{\leftarrow} X$, the probability of $f(x)$ to be a j -node, denoted by $r_{(aN)}(j)(j \geq 1)$, is approximately

$$r_{(aN)}(j) = j \cdot q_{(aN)}(j) . \quad (6)$$

The target function in cryptanalytic search problems is usually modeled as a pseudo-random function (PRF) or a cryptographic hash function (CHF). The precise interpretation of this notions can be found in Sect. 3.5 and Sect. 5.5 of [33]. It can be assumed that PRF and CHF have similar statistic behaviors to random function.

The formal definitions of search problems relevant to symmetric cryptanalysis can be described with random function. The way of generating the given information in each problem is carefully distinguished. The first is key search problem which comes from the secret key search problem using a pair of plaintext and ciphertext of an encryption algorithm.

Definition 3 (Key Search (KS)). For a fixed $f: X \rightarrow Y$, $y = f(x_0)$ generated from an element $x_0 \stackrel{\$}{\leftarrow} X$ is given. Key Search is to find the target x_0 .

The existence of the target x_0 in X is always ensured. However pre-images of y other than x_0 can be found, which is called a *false alarm*. The false alarms have to be resolved by additional information since no clue (that helps to recognize the real target) is given within the problem. Handling of false alarms is assumed not to consume quantum resources.

Definitions coming from the pre-image and the collision problems of CHF are given as follows.

Definition 4 (Pre-image Search (PS)). For a fixed $f: \{0, 1\}^* \rightarrow Y$, y is chosen at random, $y \stackrel{\$}{\leftarrow} Y$, or equivalently, $y = f(x_0)$ for $x_0 \stackrel{\$}{\leftarrow} \{0, 1\}^*$. Pre-image Search is to find any x satisfying $f(x) = y$ for given y .

There is no false alarm in pre-image search. However, the existence of a pre-image in a fixed subset of $\{0, 1\}^*$ cannot be ensured.

Definition 5 (Collision Finding (CF)). For a fixed $f: \{0, 1\}^* \rightarrow Y$, Collision Finding is to find any input pair (x_1, x_2) satisfying $f(x_1) = f(x_2)$.

$\{0, 1\}^*$ is imported since a domain of reasonable size including the original pre-image of a practically given hash value cannot be specified.

3.2 Trade-offs in Grover's Algorithm for Key Search

In this subsection, the expected iteration number and the parallelization trade-offs of Grover's algorithm are given. We assume that $|X| = |Y| = N$.

In key search problem, y becomes t -node with probability $r(t)$ of (6). The probability that one of the pre-images of y is found by the measurement after i -times Grover iterations becomes $p_t(i)$ of (2). Since only one target among t pre-images is the true key, the probability that the answer is correct is $1/t$. $P_{\text{rand}}^{\text{KS}}(i)$ denotes the success probability after i -times Grover iterations of key search problem. To emphasize that f is assumed to be a random function, the subscript 'rand' is specified. $P_{\text{rand}}^{\text{KS}}(i)$ is the summation over possible t 's,

$$P_{\text{rand}}^{\text{KS}}(i) = \sum_{t \geq 1} r(t) \cdot p_t(i) \cdot \frac{1}{t} . \quad (7)$$

Proposition about the optimal expected iterations follows.

Proposition 1. The optimal expected number, $I_{\text{rand}}^{\text{KS}}$ of Grover iterations for key search problem of random function becomes

$$I_{\text{rand}}^{\text{KS}} = 0.951 \dots \cdot \sqrt{N} .$$

Proof. This proof is similar to the one in Sect. 4 of [13].

If the measurement is made after i -times Grover iterations, the expected number of iterations can be expressed as a function of i , denoted by $I_{\text{rand}}^{\text{KS}}(i)$, which reads

$$I_{\text{rand}}^{\text{KS}}(i) = \frac{i}{P_{\text{rand}}^{\text{KS}}(i)} .$$

The optimal value, $I_{\text{rand}}^{\text{KS}}$ is the first positive local minimum value of $I_{\text{rand}}^{\text{KS}}(i)$. The first positive root of derivative of $I_{\text{rand}}^{\text{KS}}(i)$, denoted by $i_{\text{rand}}^{\text{KS}}$, can be calculated by a numerical analysis such as Newton's method. The result is $i_{\text{rand}}^{\text{KS}} = 0.434 \dots \cdot \sqrt{N}$ and $I_{\text{rand}}^{\text{KS}} = I_{\text{rand}}^{\text{KS}}(i_{\text{rand}}^{\text{KS}})$. \square

Comparing $I_{\text{rand}}^{\text{KS}}$ with I_1 , the expected iteration increases by 37.8...%.

The parallel trade-offs curve of key search problem is calculated in the rest of this subsection. If inner parallelization method is taken for $S_q \gg 1$, the number of pre-images of y in each divided space becomes only 0 or 1 for overwhelming probability, even though f is random function. Therefore the success probability after i -times iterations, denoted by $P_{\text{rand}}^{\text{KS:IP}}(i)$, reads

$$P_{\text{rand}}^{\text{KS:IP}}(i) = P_{\text{rand},N}^{\text{KS:IP}}(i) = p_{1,(N/S_q)}(i) , \quad (8)$$

from (2). The optimal expected iteration number is similar to (4) as

$$I_{\text{rand}}^{\text{KS:IP}} = 0.690 \dots \cdot \sqrt{(N/S_q)} . \quad (9)$$

In outer parallelization method, the success probability after i -times iterations becomes

$$P_{\text{rand}}^{\text{KS:OP}}(i) = 1 - (1 - P_{\text{rand}}^{\text{KS}}(i))^{S_q} ,$$

and then the optimal expected iteration number for $S_q \gg 1$ is given by

$$I_{\text{rand}}^{\text{KS:OP}} = 0.784 \dots \cdot \sqrt{(N/S_q)} . \quad (10)$$

As a result, inner parallelization is 11.9...% more efficient than outer method in key search problem. We denote the number of machines used in key search problem S_q^{KS} . The optimal expected number of iterations in key search problem, denoted by T_q^{KS} , can be considered as $I_{\text{rand}}^{\text{KS:IP}}$.

Proposition 2 (KS trade-offs curve). *For $S_q^{\text{KS}} \gg 1$, the parallelization trade-offs of Grover's algorithm for key search of random function is given by*

$$(T_q^{\text{KS}})^2 S_q^{\text{KS}} = 0.476 \dots \cdot N .$$

In the followings, the optimal expected number of iterations and trade-offs curves are defined and analyzed in the same way as in this subsection, but briefly.

3.3 Trade-offs in Grover's Algorithm for Pre-image Search

Let X be the domain of the function f , and assume $|X| = |Y| = N$. In pre-image search problem, there exist t pre-images of y with probability $q(t)$ in (5). The success probability of measuring one of the targets after i -times iterations is a summation of $q(t) \cdot p_t(i)$ over t as

$$P_{\text{rand}}^{\text{PS}}(i) = \sum_{t \geq 0} q(t) \cdot p_t(i) .$$

Since $p_0(i) = 0$ and $q(t) = r(t)/t$ for $t \geq 1$, it can be written as $P_{\text{rand}}^{\text{PS}}(i) = P_{\text{rand}}^{\text{KS}}(i)$. The important difference between the key search and the pre-image search is the existence of failure probability. If the domain of size N is used, the probability there is no pre-image of y in X is $q(0) = 1/e \approx 0.368 \dots$

Two resolutions can be sought for the failure. The first is to change the domain X in every execution of Grover's algorithm. In this case, the result on the optimal iteration number of pre-image search becomes the same as Proposition 1. The second is to expand the domain size, $|X| = aN \in \mathbb{N}$ for some $a > 1$. The success probability then reads $P_{\text{rand},(aN)}^{\text{PS}}(i) = \sum_{t \geq 1} q_{(aN)}(t) \cdot p_{t,(aN)}(i)$.

Proposition 3. *If $|X| \gg N$, the optimal expected number of iterations, denoted by $I_{\text{rand},(\gg N)}^{\text{PS}}$, for pre-image search problem is written as*

$$I_{\text{rand},(\gg N)}^{\text{PS}} = 0.690 \dots \cdot \sqrt{N} .$$

When $N = 2^{256}$, the proposition can be assumed to hold for $a \geq 2^{10}$. Subscript ‘ $\gg 1$ ’ specifies the assumption. The fact that $I_{\text{rand},(\gg N)}^{\text{PS}} \approx I_{1,N}$, i.e., better performance up to some converged value for larger domain size, is remarked. If a grows to 8, the inversely proportional failure probability decreases below $0.0004 \dots \approx 1/e^8$.

In the case of inner parallelization for $|X| = |Y|$, the pre-images of y are distributed to different divided spaces with overwhelming probability when $S_q \gg 1$. The success probability reads

$$P_{\text{rand},N}^{\text{PS:IP}}(i) = \sum_{t \geq 1} q(t) \cdot \left\{ 1 - (1 - p_{1,(N/S_q)}(i))^t \right\} .$$

and the optimal expected iteration number is written as

$$I_{\text{rand},N}^{\text{PS:IP}} = 0.981 \dots \cdot \sqrt{(N/S_q)} . \quad (11)$$

Since $P_{\text{rand}}^{\text{PS}}(i) = P_{\text{rand}}^{\text{KS}}(i)$, the behavior of outer parallelization of pre-image search is the same as in key search. The optimal expected iteration number is

$$I_{\text{rand}}^{\text{PS:OP}} = 0.784 \dots \cdot \sqrt{(N/S_q)} \quad (= I_{\text{rand}}^{\text{KS:OP}}) . \quad (12)$$

When $|X| = aN \in \mathbb{N}$ for some $a > 1$, if $S_q > a^2$, it can be assumed that all pre-images of y are separately distributed to the divided space in inner parallelization. For both of inner and outer parallelization, the optimal expected iteration converges to the value of (12) when $a \gg 1$ and $S_q > a^2$.

There are subtleties in comparing inner and outer parallelization which are inappropriate to be pointed out here. We conclude that it is always favored to enlarge the domain size, and then for large S_q , two parallelization methods show asymptotically the same performance. Denoting the optimal time and space complexities for pre-image search problem by T_q^{PS} and S_q^{PS} , the trade-offs curve is given as follows.

Proposition 4 (PS trade-offs curve). *For $S_q^{\text{PS}} \gg 1$, the parallelization trade-offs of Grover’s algorithm for pre-image search of random function is given by*

$$(T_q^{\text{PS}})^2 S_q^{\text{PS}} = 0.614 \dots \cdot N .$$

Note that while the inner parallelization is a better option in key search, both methods have similar behaviors in pre-image search.

3.4 Trade-offs in Quantum Collision Finding Algorithms

A collision could be found by using Grover's algorithm in the way of *second pre-image* search. This has the same result as Sect. 3.3 if the input of the given pair of 'first pre-image' is not included in the domain. Apart from Grover's algorithm, the optimal expected iterations and trade-offs curves for parallelizations of two collision finding algorithms, GwDP and CNS, are given in this subsection.

In collision finding algorithms, searching for a pre-image of large size set is required. Let $f: \{0, 1\}^* \rightarrow Y$, $X \subset \{0, 1\}^*$ be a set of size $aN \in \mathbb{N}$. For $y \xleftarrow{\$} Y$, the expected number of pre-images of y becomes $a = \sum_{j \geq 1} j \cdot q_{(aN)}(j)$. If the size of a set $A \subset Y$ is large enough, it can be assumed that the number of pre-images of A , $|\{x \in X : f(x) \in A\}| = a \cdot |A|$.

GwDP algorithm. Let $S_q = 2^s$ and $X \subset \{0, 1\}^*$ be a set of size N . In each quantum machine, a parameter $(n - 2s + 2)$ is used for the number of bits to be fixed in DPs. The parameter $(n - 2s + 2)$ is chosen as an optimal one only among integers in order to allow the easier implementation by quantum gates.

After i -times Grover iterations, the success probability of measuring a DP becomes $p_{(2^{2s-2})}(i)$ from (2). The expected number of DPs found is $2^s \cdot p_{(2^{2s-2})}(i)$ by measurements after i -times iterations on each machine. As a result of *birthday problem* (BP), known to be proposed by R. Mises in 1939, if there are k samples independently selected out of 2^{2s-2} DPs, the probability of at least one coincidence, denoted by $p_{(2^{2s-2})}^{\text{BP}}(k)$, is approximated as

$$p_{(2^{2s-2})}^{\text{BP}}(k) = 1 - \exp\left(\frac{-k^2}{2 \cdot 2^{2s-2}}\right) .$$

Details of approximation can be found in Sect. A.4 of [33]. The probability of finding at least one collision, denoted by $P_{\text{rand}}^{\text{GwDP}}(i)$, is then

$$P_{\text{rand}}^{\text{GwDP}}(i) = p_{(2^{2s-2})}^{\text{BP}}(2^s \cdot p_{(2^{2s-2})}(i)) .$$

The optimal expected iteration reads

$$I_{\text{rand}}^{\text{GwDP}} = 1.532 \dots \cdot \sqrt{N}/2^s . \quad (13)$$

Denoting the optimal time and space complexities by T_q^{GwDP} and S_q^{GwDP} for collision finding by GwDP algorithm, the trade-offs curve is given as follows.

Proposition 5 (GwDP trade-offs curve). *For $S_q^{\text{GwDP}} = 2^s \gg 1$, the trade-offs curve of GwDP algorithm for random function is given by*

$$T_q^{\text{GwDP}} S_q^{\text{GwDP}} = 1.532 \dots \cdot N^{\frac{1}{2}} .$$

Note that the algorithm also requires $S_c^{\text{GwDP}} = O(2^s)$ classical storage.

CNS algorithm. In the list preparation phase, a list L of size 2^l , a subset of d -bit DPs, is to be made. Set $X_1 \subset \{0, 1\}^*$ of size $N = 2^n$ and the function

$$f_{DP}(x) = \begin{cases} 1 & , \text{ if } f(x) \text{ is DP} & , \\ 0 & , \text{ otherwise} & . \end{cases}$$

Let $f_{DP}|_{X_1}$ be the restriction of f_{DP} on X_1 . Grover iteration in this phase is defined by $Q_{DP} = -A_{DP}S_0A_{DP}^{-1}S_{f_{DP}|_{X_1}}$ where the oracle operator $S_{f_{DP}|_{X_1}}$ is a quantum implementation of the function $f_{DP}|_{X_1}$ and A_{DP} is a usual state preparation operator $H^{\otimes n}$.

Since there are about 2^{n-d} ($= |X_1|/2^d$) DPs in X_1 , the expected number of Grover iterations to find a DP is the same as $I_{2^{n-d}} = 0.690 \dots \cdot 2^{d/2}$ of (4). The expected number of Grover iterations to build L is $0.690 \dots \cdot 2^{l+d/2}$. A classical storage of size $O(2^l)$ is required in addition.

In the collision finding phase, let $X_2 \subset \{0, 1\}^*$ be a set of size N such that $X_1 \cap X_2 = \emptyset$. Let the state $|\psi\rangle$ be an equal-phase and equal-weight superposition of states encoding all the DPs in X_2 . State preparation operator A_L such that $|\psi\rangle = A_L|0\rangle$ is explicitly

$$A_L = (-A_{DP}S_0A_{DP}^{-1}S_{f_{DP}|_{X_2}})^{(\pi/4) \cdot 2^{d/2}} A_{DP} & ,$$

which is essentially Grover iterations similar to Q_{DP} with repetition number $I_{2^{n-d}}^{\text{mp}}$ of (3). The function $f_L: X_2 \rightarrow \{0, 1\}$ is defined as

$$f_L(x) = \begin{cases} 1 & , \text{ if } f(x) \in L & , \\ 0 & , \text{ otherwise} & . \end{cases}$$

To realize the oracle operator S_{f_L} – a quantum implementation of f_L – without a need for quantum memory, the authors have suggested a computational method taking $O(2^l)$ elementary operations per quantum f_L query.

QAA iteration Q_L of the collision finding phase consists of two steps. The first is acting of the oracle operator S_{f_L} . Let t_L be the ratio of the time cost of S_{f_L} per list element of L to that of Grover iteration. The second step is acting of the diffusion operator $-A_L S_0 A_L^{-1}$.

The success probability of QAA algorithm is known to have the same behaviors of Grover's algorithm [14]. Since there are about 2^{n-d} ($= 2^n \cdot (1/2^d)$) DPs encoded in the state with equal probabilities and about 2^l ($= 2^{n-d} \cdot (2^l/2^{n-d})$) pre-images of L in $|\psi\rangle$, by applying Q_L operator $I_{(2^l), (2^{n-d})} = 0.690 \dots \cdot 2^{(n-d-l)/2}$ times on $|\psi\rangle$, the algorithm is expected to find a collision. The time cost of the collision finding phase reads $0.690 \dots \cdot 2^{(n-d-l)/2} \cdot \{2 \cdot (\pi/4) \cdot 2^{d/2} + t_L \cdot 2^l\}$. Note that the time cost of S_0 in collision finding phase and the initial A_L are negligible. The time cost of CNS algorithm in terms of Grover iterations denoted by $I_{\text{rand}}^{\text{CNS}}(d, l)$ reads

$$I_{\text{rand}}^{\text{CNS}}(d, l) = 0.690 \dots \cdot 2^{l+\frac{d}{2}} + 0.690 \dots \cdot 2^{\frac{n-d-l}{2}} \left(\frac{\pi}{2} \cdot 2^{\frac{d}{2}} + t_L \cdot 2^l \right) . \quad (14)$$

The optimal value $I_{\text{rand}}^{\text{CNS}}$ is given as follows.

Proposition 6. *The optimal expected number of Grover iterations in CNS algorithm for collision finding of random function reads*

$$I_{\text{rand}}^{\text{CNS}} = 3.150 \dots \cdot t_L^{\frac{1}{5}} \cdot 2^{\frac{2}{5}n} ,$$

when $l = d/2 + \log_2(\pi/(2t_L))$, and $d = 2/5\{n + \log_2((2t_L)^3/\pi)\}$.

Using $S_q = 2^s$ quantum machines, natural parallelization of the list preparation phase is finding 2^{l-s} elements on each machine. Outer parallelization of QAA algorithm in the collision finding phase has the same expected iterations as (12). The expected number of Grover iterations, denoted by $I_{\text{rand}}^{\text{CNS:OP}}(d, l)$, where $s < \min(l, n - d - l)$, is written as

$$I_{\text{rand}}^{\text{CNS:OP}}(d, l) = 0.690 \dots \cdot 2^{l+\frac{d}{2}-s} + 0.784 \dots \cdot 2^{\frac{n-d-l-s}{2}} \left(\frac{\pi}{2} \cdot 2^{\frac{d}{2}} + t_L 2^l \right) .$$

When $l = d/2 + \log_2(\pi/(2t_L))$, and $d = 2/5\{n + s + \log_2(1.291 \dots \cdot (2t_L)^3/\pi)\}$, the optimal expected number of iterations reads

$$I_{\text{rand}}^{\text{CNS:OP}} = 3.488 \dots \cdot t_L^{\frac{1}{5}} 2^{\frac{2}{5}n} 2^{-\frac{3}{5}s} . \quad (15)$$

We denote the optimal time and space complexities by T_q^{CNS} and S_q^{CNS} for collision finding by CNS algorithm. T_q^{CNS} can be considered as $I_{\text{rand}}^{\text{CNS:OP}}$. The trade-offs curve of CNS algorithm is then given as follows.

Proposition 7 (CNS trade-offs curve). *For $S_q^{\text{CNS}} \gg 1$, the parallelization trade-offs curve of CNS algorithm for random function is given by*

$$(T_q^{\text{CNS}})^5 (S_q^{\text{CNS}})^3 = (3.488 \dots)^5 \cdot t_L \cdot N^2 .$$

The algorithms also requires the classical resource $S_c^{\text{CNS}} = O(N^{1/5}(S_q^{\text{CNS}})^{1/5})$. If the constant t_L is determined, the time-space complexity of CNS algorithm could be derived from this trade-offs curve.

4 Depth-Qubit Cost Metric

Universal quantum computers are capable of carrying out elementary logic operations such as Pauli X, Hadamard, CNOT, T, and so on. See [38] for details on quantum gates. Implementation of any cryptographic operation in this paper is restricted such that it can only be realized by using these gates. One may think of the restriction as a quantum version of software implementation in classical computing. Quantum security of symmetric cryptosystems can then be estimated in units of elementary logic gates.

It is generally known that each elementary gate has different physical implementation time. Considering various aspects of quantum computing, we suggest to simplify a measure of computation time and to ignore all the other factors or gates that complicates the analysis of quantum algorithms.

Two primary resources in quantum computing, circuit depth and qubit, can be exchanged to meet a certain attack design criteria. Time-space complexity investigated in the previous section can be used to give an attribute ‘efficiency’ to each and every design. To further quantify *depth-qubit complexity* and to be able to rank the efficiency, we briefly cover the time-space trade-offs of quantum resources in this section.

4.1 Cost Measure

Difficulties often arise when it comes to setting quantum complexity measures that are physically interpretable. There exists a number of factors making it complicate, for example totally different architecture each experimental group is pursuing. A qubit or a certain gate may costs differently in each architecture. It is therefore hardly possible to accurately assess operational time of each type of gate in general, and to estimates overall run time. Despite the notable difficulty in quantifying the basic unit cost of quantum computation, a number of groups have attempted to estimate the algorithm costs in various applications [4,27,43]. The cost metric varies depending on author’s viewpoint. For example, one considering the fault-tolerant computation would estimate the cost involving specific hardware implementations or error-correction schemes. On the other hand, one that is not to impose constraints on hardware or error-correction scheme would estimate the cost in logical qubits and gates. The latter approach is adopted in this work. Readers should keep in mind that this approach ignores the overheads introduced by fault-tolerance⁴.

High-level circuit description of Grover iteration involves not only elementary gates but also larger gates such as C^k NOT. It is very unlikely that such gates can be directly operated in any realistic universal quantum computers. Decomposition of those gates into smaller ones is thus required in practical estimates.

Determining the unit time cost is a subtle matter. We would like to address that the simplest, yet justified time cost measure involves Toffoli gate.

Definition 6. *A unit of quantum computational time cost is the time required to operate a non-parallelizable Toffoli gate.*

In other words, Toffoli-depth will be the time cost of the algorithm. Three justifications can be given for the distinctness of Toffoli gates. First, Toffoli (and single) gates are universal [21,49,50]. Second, circuits consisting only of Clifford gates are not advantageous over classical computing, implying that the use of non-Clifford gates such as Toffoli is essential for quantum benefit [26]. Third, logical Toffoli gates are the main source of time bottleneck [4,24,32] due to the magic state distillation process for T gates [17] which comes only from a decomposition of Toffoli gates in this work. See for example in [4], the ratio of execution time in all Clifford gates to all T gates is about 0.0001 . . . in SHA-256 when fault-tolerance is considered. To sum up we adopt universal Toffoli gate as

⁴ Fault-tolerant cost could be in general huge, but we expect that logical cost to fault-tolerant cost conversion would be more or less uniform.

the only non-Clifford gate, which is responsible for quantum speedup as well as main time bottleneck of circuits presented in this paper. Moreover it is plausible to assume that multiple Toffoli gates can be applied to qubits simultaneously as long as their input/output qubits are independent, justifying Definition 6.

Space cost is estimated as a total number of logical qubits required to perform the quantum search algorithm.

Definition 7. *Quantum computational space cost is the number of logical qubits required to run the entire circuit.*

Decomposition of a high-level circuit component into smaller ones often entails a need for additional qubits, which sometimes turn into garbage bits or get cleaned after certain operations. Overall space cost mainly comes from these qubits. To avoid confusion caused by terminology, we clarify five kinds of qubits.

1. **Data qubits** are qubits of which the space is searched by the quantum search algorithm. For example in AES-128, the size of the key space is 2^{128} which requires 128 data qubits.
2. **Work qubits** are initialized qubits those assist certain operation. Whether it stays in an initialized value or gets written depends on the operation.
3. **Garbage qubits** are previously initialized work qubits, which then get written unwanted information after a certain operation.
4. **Output qubits** are previously initialized work qubits, which then get written the output information of a certain operation.
5. **Oracle qubit** is a single qubit used for phase kick-back (sign change) in oracle and diffusion operators.

There is one more type of qubit not falling into above categories; a borrowed qubit [29]. The concept of the borrowed qubit is not considered in this work. Garbage and output qubits must be re-initialized before the diffusion of Grover iteration to be disentangled from data qubits.

4.2 Toffoli and T Gates

Commonly acknowledged universal quantum gate set consists of Clifford gates and T gate. As stated in the previous subsection, operational time costs of different gates may vary depending on architecture. However, it is less disputable that physical implementation of T gate (or preparation of the magic state) is important, difficult, and generally more expensive than Clifford gates. There are communities dedicated to better implementation of T [16,17,31,36] and reducing the number of the gates applied [2,3,12,25,46], as it is time bottleneck in fault-tolerant quantum computing.

Toffoli gate is a non-Clifford gate that is composed of a few T and Clifford gates. Taking Toffoli gate over T gate as a basic unit of time resource has its merits and demerits. We cautiously compare the relation between Toffoli and T to the one between high- and low-level languages. Example of implementation of a two-bit addition in terms of Toffoli and T gates are given in Fig. 4.




Fig. 4. Addition of two bits a and b in terms of (a) Toffoli and (b) T gates (Fig. 7(d) in [3]). The third qubit (output qubit) is written a carry. The third and the second qubits save the binary representation of $a + b$ as $ab \cdot 2^1 + (a \oplus b) \cdot 2^0$.

Being reminded that Toffoli and CNOT operate as

$$\text{TOF}|a\rangle|b\rangle|0\rangle = |a\rangle|b\rangle|0 \oplus ab\rangle, \quad \text{CNOT}|a\rangle|b\rangle = |a\rangle|a \oplus b\rangle,$$

respectively, it is immediately noticeable from Fig. 4(a) that the circuit works as a two-bit addition operator. The same operation realized by depth-optimized Clifford+T set [3] is described in Fig. 4(b). Assuming that a given quantum computer can only perform gates in Clifford+T set, this circuit enables more transparent expectation of runtime.

Typically in previous studies a quantum algorithm is first implemented in Toffoli-level, and then the circuit undergoes a kind of ‘compilation’ process that looks for an elementary-level circuit [4,46]. Finding an optimal compiling method is very complicated and worth researching [43]. At this stage however, it is hardly possible to find true optimal elementary-level circuit from compiling huge high-level circuit. In this work therefore, we stay in Toffoli-level implementation conforming the purpose of providing a general framework.

4.3 Time-Space Trade-offs

Readers those are familiar with quantum circuit model can safely skip over this subsection as it covers some general facts about depth-qubit trade-offs. In quantum circuit model, it is often possible to sacrifice efficiency in qubits for better performance in time and vice versa. Quantum version of such time-space trade-offs forms a main body of Sect(s). 5 and 6. As preliminary we give an example to introduce the general concept of trade-offs in quantum circuits.

Consider a function f that carries out binary multiplications of k single bit values. At the end of this subsection we will deal with general k , but for now, let us explicitly write down the description with $k = 2$, the multiplication of two bits a and b as $f(a, b) = ab$. It is noticeable that the function f can be implemented by AND gate in a classical circuit. However in a quantum circuit where only unitary operations are allowed, similar implementation is prohibited since AND operation is not unitary as the input information a, b cannot be retrieved by knowing ab only. Simple resolution can be found by keeping the input information all the way such that

$$U_f|a\rangle|b\rangle|0\rangle = |a\rangle|b\rangle|0 \oplus ab\rangle, \quad (16)$$

where $|a\rangle$ and $|b\rangle$ are quantum states encoding a and b , and U_f is the quantum implementation of the function f . Previously zeroed qubit represented by the state $|0\rangle$ on the left-hand side holds the result after the operation. There exists a quantum gate that exactly performs the operation by U_f called a k -fold controlled-NOT (C^k NOT) with $k = 2$, or better known as Toffoli gate. Figure 5(a) illustrates the graphical representation of C^2 NOT gate achieving (16). General C^k NOT gates read k input bits carried by wires intersecting with black dots and change a target bit carried by a wire intersecting with exclusive-or symbol. In this case, the gate works as NOT on target bit if $a = b = 1$ and identity otherwise.




Fig. 5. (a) C^2 NOT (Toffoli) gate and (b) C^4 NOT gate

Similarly, multiplications of four bits can be implemented by using C^4 NOT gate as shown in Fig. 5(b). C^4 NOT gate carries out NOT operation on target bit if $a = b = c = d = 1$ and nothing otherwise.

Now assume we are to split up a C^4 NOT gate into multiple Toffoli gates with the help of a few extra qubits. Decomposing a large gate into smaller gates is a typical task one confront in compilation [43]. There can be various ways to achieve the goal, and one of the immediate designs is the one in Fig. 6(a).




Fig. 6. Decomposition of C^4 NOT gate into (a) five Toffoli gates and (b) ten Toffoli gates. In (a), the third and the fifth zeroed qubits from the top are work qubits whereas in (b), only the fifth arbitrary-valued qubit is a work qubit.

Let us examine the action of each Toffoli gate on the register one-by-one,

$$\begin{array}{lcl}
& |a\rangle|b\rangle|0\rangle|c\rangle|0\rangle|d\rangle|0\rangle & \xrightarrow{\textcircled{1}} |a\rangle|b\rangle|ab\rangle|c\rangle|0\rangle|d\rangle|0\rangle \\
\textcircled{2} \xrightarrow{\quad} & |a\rangle|b\rangle|ab\rangle|c\rangle|abc\rangle|d\rangle|0\rangle & \xrightarrow{\textcircled{3}} |a\rangle|b\rangle|ab\rangle|c\rangle|abc\rangle|d\rangle|abcd\rangle \\
\textcircled{4} \xrightarrow{\quad} & |a\rangle|b\rangle|ab\rangle|c\rangle|0\rangle|d\rangle|abcd\rangle & \xrightarrow{\textcircled{5}} |a\rangle|b\rangle|0\rangle|c\rangle|0\rangle|d\rangle|abcd\rangle \quad ,
\end{array} \tag{17}$$

where the circled number above the mapping arrow indicates the corresponding Toffoli gate in Fig. 6(a). The result actually comes out after $\textcircled{3}$, but we further perform a kind of un-computation with two extra Toffoli gates to re-initialize the work qubits. It is up to users to decide whether the procedure should stop just after $\textcircled{3}$ at the cost of two garbage qubits being generated, or go all the way to the end of the circuit. As one can notice, it is already the trade-offs.

A less straightforward decomposition can be found in Fig. 6(b). It makes use of twice as many Toffoli gates as Fig. 6(a) but requires only a single arbitrary qubit⁵. Each Toffoli gate transforms the state as follows.

$$\begin{array}{lcl}
& |a\rangle|b\rangle|c\rangle|d\rangle|w\rangle|0\rangle & \xrightarrow{\textcircled{1}\textcircled{2}} |a\rangle|b\rangle|c\rangle|d\rangle|w\rangle|ab\rangle \\
\textcircled{3} \xrightarrow{\quad} & |a\rangle|b\rangle|c\rangle|d\rangle|w \oplus abc\rangle|ab\rangle & \xrightarrow{\textcircled{4}} |a\rangle|b\rangle|c\rangle|d\rangle|w \oplus abc\rangle|0\rangle \\
\textcircled{5} \xrightarrow{\quad} & |a\rangle|b\rangle|c\rangle|d\rangle|w \oplus abc\rangle|abcd \oplus dw\rangle & \\
\textcircled{6} \xrightarrow{\quad} & |a\rangle|b\rangle|c\rangle|d\rangle|w \oplus abc \oplus abcd \oplus cdw\rangle|abcd \oplus dw\rangle & \tag{18} \\
\textcircled{7} \xrightarrow{\quad} & |a\rangle|b\rangle|c\rangle|d\rangle|w \oplus abc \oplus abcd \oplus cdw\rangle|abcd \oplus dw \oplus ab\rangle & \\
\textcircled{8} \xrightarrow{\quad} & |a\rangle|b\rangle|c\rangle|d\rangle|w\rangle|abcd \oplus dw \oplus ab\rangle & \\
\textcircled{9} \xrightarrow{\quad} & |a\rangle|b\rangle|c\rangle|d\rangle|w\rangle|abcd \oplus dw\rangle & \xrightarrow{\textcircled{10}} |a\rangle|b\rangle|c\rangle|d\rangle|w\rangle|abcd\rangle \quad .
\end{array}$$

Both designs work as desired. In fact for general k , time-efficient design as in Fig. 6(a) requires $k - 2$ zeroed work qubits within depth $2k - 3$, whereas space-efficient design as in Fig. 6(b) uses only one arbitrary qubit within depth $8k - 24$ (for $k \geq 5$) [6]. We denote time- and space-efficient designs lower-depth and less-qubit $C^k\text{NOT}$, respectively.

Bit multiplication is one of examples qubit and depth are mutually exchangeable. In Sect(s). 5 and 6 we will compare multiple circuits that do the same job with a different number of qubits, and examine the consequence of each design when parallelized.

5 Complexity of AES-128 Key Search

This section presumes that readers are familiar with standard AES-128 encryption algorithm [39]. We assume that a quantum adversary is given a plaintext-ciphertext pair and asked to find the key used for the encryption. Since AES-128 works as a PRF, it is possible that multiple keys lead to the same ciphertext,

$$\text{AES}(k_0, p) = \text{AES}(k_1, p) = \dots \quad ,$$

⁵ The first Toffoli gate in Fig. 6(b) is redundant in this case, but needed if one wants to carry out $z \oplus abcd$, where z is the initial value of the last qubit.

where $k_i \in \{0, 1\}^{128}$ are different keys and p is a given plaintext. The term pre-image will be used to denote each key k_i that generates given ciphertext upon the encryption of given plaintext.

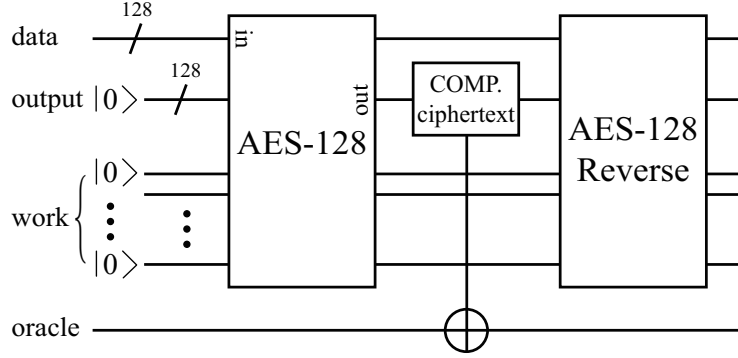


Fig. 7. Oracle circuit for key search attack on AES-128

The idea of applying Grover's algorithm to exhaustive attack on AES-128 is as follows. Linearly superposed 2^{128} input keys encoded in 128 data qubits are fed as an input to an AES box shown in Fig. 7. AES box contains a reversible circuit implementation of AES-128 encryption algorithm. The AES box encrypts the given plaintext, outputting superposed ciphertexts encoded in output qubits. Superposed ciphertexts are then compared with given ciphertext via C^{128} NOT gate to mark the target. After marking is done, every qubit except the oracle qubit is passed on to a reverse AES box to disentangle the data qubits from other qubits.

5.1 Circuit Implementation Cost

AES-128 encryption internally performs SubBytes, MixColumns, ShiftRows, AddRoundKey, SubWord, RotWord and Rcon. Quantum circuits for these operations are mostly adopted from [27] with improvements and fixes.

MixColumns, ShiftRows and RotWord are linear operations acting on 32 bits that do not require any work qubit nor Toffoli gate. Among them, last two are simple bit permutations which require no quantum gates (by re-wiring) or at most SWAP gates. MixColumns needs to be treated more carefully as it is not a bit permutation. Treating each four-byte column of the internal state as a length-four vector, MixColumns is expressed as a matrix multiplication,

$$\begin{pmatrix} s'_{0,j} \\ s'_{1,j} \\ s'_{2,j} \\ s'_{3,j} \end{pmatrix} = \begin{pmatrix} 02 & 03 & 01 & 01 \\ 01 & 02 & 03 & 01 \\ 01 & 01 & 02 & 03 \\ 03 & 01 & 01 & 02 \end{pmatrix} \begin{pmatrix} s_{0,j} \\ s_{1,j} \\ s_{2,j} \\ s_{3,j} \end{pmatrix}, \quad \text{for } 0 \leq j \leq 3, \quad (19)$$

where 01, 02, 03 are sub-matrices when each byte $s_{i,j}$ is treated as a length-eight vector, written as

$$01 = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}, \quad 02 = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}, \quad 03 = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}.$$

Since an explicit form of transformation matrix is given in (19), the quantum circuit implementation of the matrix can be found by methods given in [11,44].

AddRoundKey and Rcon are XOR-ings of fixed-size strings which can also be efficiently realized by CNOT or X gates only.

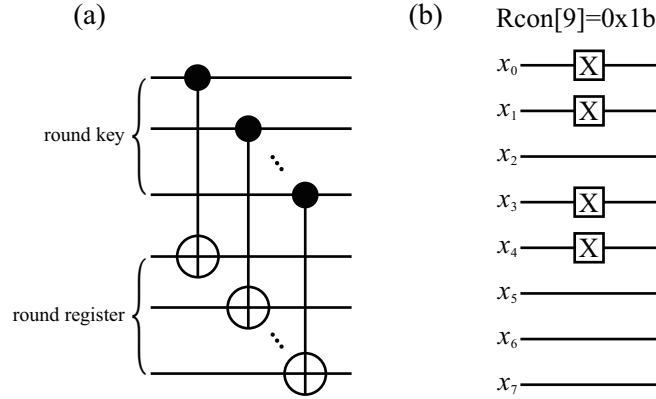


Fig. 8. Quantum circuit implementation of (a) AddRoundKey and (b) Rcon. In (b), Rcon[9] circuit is given as an example.

SubBytes and SubWord are the only operations which require quantum resources. Since SubBytes and SubWord consist of 16 and 4 S-boxes, the S-box is the only operation to be carefully discussed.

Classically, S-box can be implemented as a look-up table. However, a quantum counterpart of such table should involve the notion of the quantum memory aforementioned in Sect. 2.2. Therefore in this work, S-box is realized by explicitly calculating multiplicative inverse followed by GF-linear mapping and addition of the S-box constant as described in Sect. 3.2.1 of [27].

S-box is realized by calculating multiplicative inverse followed by GF-linear mapping and addition of S-box constant. By treating a byte as an element in $\text{GF}(2^8) = \text{GF}(2)[x]/(x^8 + x^4 + x^3 + x + 1)$, GF-linear mapping and addition of

S-box constant are summarized as the equation

$$\begin{pmatrix} x'_0 \\ x'_1 \\ x'_2 \\ x'_3 \\ x'_4 \\ x'_5 \\ x'_6 \\ x'_7 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 1 & 1 & 0 & 0 & 0 & 1 & 1 & 1 \\ 1 & 1 & 1 & 0 & 0 & 0 & 1 & 1 \\ 1 & 1 & 1 & 1 & 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} x_0 \\ x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \\ x_7 \end{pmatrix} + \begin{pmatrix} 1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 1 \\ 1 \\ 0 \end{pmatrix}, \quad (20)$$

where addition is XOR operation and x_i are coefficients of polynomial of order x^7 . No work qubit nor Toffoli gate is required in this step. While XOR operation is simply done by applying X gates to relevant qubits, implementing a transformation matrix in (20) is not trivial. See [11,44] for general methods of realizing linear transformations.

Resource estimate of quantum AES-128 encryption has been narrowed down to estimate the cost of finding multiplicative inverse of the element α in $GF(2^8)$. In [27], multiplicative inverse of α is calculated by using two arithmetic circuits; Maslov et al.'s modular multiplier [35] and in-place squaring [27]. Slight modification of previous method is found in this work with seven multipliers being used, verified by the quantum circuit simulation by matrix product state [51], with the seven multipliers being used as following sequences,

$$\begin{array}{ccccccc} |\alpha\rangle & & |\alpha\rangle & & |\alpha\rangle & & |\alpha\rangle \\ |0\rangle & & |0\rangle & & |0\rangle & & |0\rangle \\ |0\rangle & \xrightarrow{\text{CNOTs}} & |\alpha\rangle & \xrightarrow{\text{Sq}} & |\alpha^2\rangle & \xrightarrow{\text{Mul}} & |\alpha^2\rangle & \xrightarrow{\text{Sq}^{-1}} & |0\rangle & \xrightarrow{\text{Sq} \times 2} & |0\rangle \\ |0\rangle & & |0\rangle & & |0\rangle & & |\alpha^3\rangle & & |\alpha^3\rangle & & |\alpha^{12}\rangle \\ |0\rangle & & |0\rangle & & |0\rangle & & |0\rangle & & |0\rangle & & |0\rangle \end{array}$$

$$\begin{array}{ccccccc} |\alpha\rangle & & |\alpha\rangle & & |\alpha\rangle & & |\alpha^{64}\rangle & & |\alpha^{64}\rangle & & |\alpha\rangle \\ |0\rangle & & |0\rangle & & |0\rangle & & |0\rangle & & |\alpha^{254}\rangle & & |\alpha^{254}\rangle \\ |0\rangle & \xrightarrow{\text{CNOTs}} & |\alpha^{60}\rangle & \xrightarrow{\text{Sq}^{-1} \times 2} & |\alpha^{60}\rangle & \xrightarrow{\text{Mul}} & |\alpha^{60}\rangle & \xrightarrow{\text{Sq}} & |\alpha^{60}\rangle & \xrightarrow{\text{Sq}^{-1} \times 6} & |\alpha^{60}\rangle \\ |\alpha^{12}\rangle & & |\alpha^{12}\rangle & & |\alpha^3\rangle & & |\alpha^3\rangle & & |\alpha^3\rangle & & |\alpha^3\rangle \\ |\alpha^{48}\rangle & & |\alpha^{12}\rangle & & |0\rangle & & |\alpha^{63}\rangle & & |\alpha^{63}\rangle & & |0\rangle \end{array}$$

$$\begin{array}{ccccccc} |\alpha\rangle & & |\alpha\rangle & & |\alpha\rangle & & |\alpha\rangle & & |\alpha\rangle & & |\alpha\rangle \\ |\alpha^{254}\rangle & \xrightarrow{\text{Sq} \times 2} & |\alpha^{254}\rangle & \xrightarrow{\text{Sq} \times 2} & |\alpha^{254}\rangle & \xrightarrow{\text{Sq}^{-1} \times 2} & |\alpha^{254}\rangle & \xrightarrow{\text{CNOTs}} & |\alpha^{254}\rangle & \xrightarrow{\text{Sq}} & |\alpha^{254}\rangle \\ |\alpha^{60}\rangle & & |0\rangle & & |0\rangle & & |\alpha\rangle & & |\alpha\rangle & & |\alpha^2\rangle \\ |\alpha^{12}\rangle & & |\alpha^{12}\rangle & & |\alpha^{12}\rangle & & |\alpha^3\rangle & & |0\rangle & & |0\rangle \\ |\alpha^{12}\rangle & & |\alpha^{48}\rangle & & |0\rangle & & |0\rangle & & |0\rangle & & |0\rangle \end{array}, \quad (21)$$

where each state ket represents eight-bit register, Sq and Mul denote modular squaring and multiplication operations, and CNOTs implies eight CNOT gates copying the string. Seven multipliers including reverse operations have been used as can be seen from (21).

As squaring in $\text{GF}(2^8)$ is linear, it does not involve the use of Toffoli nor work qubits. Therefore it is only required to estimate the cost of multipliers. Table 1 summarizes the elementary operation costs in AES-128. Two distinct multipliers are considered in this work; Maslov et al.’s design [35] and Kepley and Steinwandt’s design [34].

Table 1. Costs of elementary operations in AES-128. A quarter of work qubits needed in S-box turn into garbage qubits.

	Less-qubit		Lower-depth	
	Multiplier	S-box	Multiplier	S-box
Toffoli-depth	18	126	8	56
Work qubits	8	32	27	108

First four multiplications in S-box are aimed at computing the multiplicative inverse. Remaining three (reverse) multiplications are then used to clean garbage qubits produced by previous multiplications. At the end of S-box, a quarter of total work qubits needed in S-box turn into garbage qubits.

5.2 Design Candidates

Four main trade-offs points are considered.

First point, that has an impact on the overall design, is to determine whether key schedule and AES rounds are carried out in parallel. As S-box is used both in key schedule and AES round, schedule-round parallel implementation would require more work qubits. This option is denoted by *serial/parallel schedule-round*.

Second, AES round functions can be reversed in the middle of encryption process to save work qubits. The idea of reverse AES round was suggested in Sect. 3.2.3 in [27]. Since each run of round function produces garbage qubits, forward running of 10 rounds accumulates ≥ 1280 garbage qubits. Putting reverse rounds in between forward rounds reduce a large amount of work qubits at the cost of longer Toffoli-depth. This option is denoted by *reverse round* when applied.

Thirdly, a choice of multiplier could make an important trade-offs point. Less-qubit and lower-depth multipliers are two options. For simplicity, we do not consider adaptive use of both multipliers although it is possible to improve the efficiency by using appropriate multiplier in different part of circuit. This option is denoted by *less-qubit/lower-depth multiplier*.

Fourth, to present the extremely depth-optimized circuit design, the cleaning process in S-box could be skipped leaving every work qubit used in S-box garbage. This option is denoted by *S-box un-cleaning* when applied.

In total, there exist 16 ($=2^4$) different circuit designs. We only take six of them into account as others seem to be flawed compared with the six. Six designs are denoted as follows.

- AES-C1: Serial schedule-round, reverse-round, less-qubit multiplier
- AES-C2: Serial schedule-round, reverse-round, lower-depth multiplier
- AES-C3: Parallel schedule-round, less-qubit multiplier
- AES-C4: Parallel schedule-round, lower-depth multiplier
- AES-C5: Parallel schedule-round, less-qubit multiplier, S-box un-cleaning
- AES-C6: Parallel schedule-round, lower-depth multiplier, S-box un-cleaning

5.3 Comparison

Toffoli-depth and total number of qubits are carefully estimated for each design. Costs of quantum AES-128 encryption circuit and entire Grover’s algorithm on a single quantum processor is summarized in Table 2. Estimates for single Grover iteration is omitted from the table as it can easily be calculated from costs of AES-128 encryption circuit;

$$\text{cost}(\text{Grover iteration}) = 2 \cdot \text{cost}(\text{AES-128}) + 2 \cdot \text{cost}(C^{128}\text{NOT}) ,$$

where $\text{cost}(C)$ is Toffoli-depth of a circuit C . Note that full Toffoli-depth of the entire Grover’s algorithm is estimated considering $I_{\text{rand}}^{\text{KS}}$ in Proposition 1.

Table 2. Costs of AES-128 encryption circuit and entire attack circuit on a single quantum processor. MAXDEPTH is not considered.

	AES-128		Grover	
	Toffoli-depth	Qubits	Toffoli-depth	Qubits
AES-C1	11088	984	$1.360 \dots \times 2^{78}$	985
AES-C2	4928	3017	$1.290 \dots \times 2^{77}$	3018
AES-C3	1260	2208	$1.405 \dots \times 2^{75}$	2209
AES-C4	560	7148	$1.510 \dots \times 2^{74}$	7149
AES-C5	720	6654	$1.808 \dots \times 2^{74}$	6655
AES-C6	320	21854	$1.064 \dots \times 2^{74}$	21855

Table 3. Comparison of time-space complexity of different AES-128 circuit designs without the oracle assumption. The smallest $c_{\#}^{\text{KS}}$ is found by AES-C4 with $c_4^{\text{KS}} = 1.048 \dots \times 2^{33}$. Other values are divided by c_4^{KS} for easier comparison.

	AES-C1	AES-C2	AES-C6	AES-C5	AES-C3	AES-C4
$c_{\#}^{\text{KS}}/c_4^{\text{KS}}$	28.606...	19.705...	1.519...	1.333...	1.070...	1

Proposition 2 basically sets up the criterion for a comparison of circuit designs. Here we replace T_q^{KS} and S_q^{KS} by $\mathcal{T}_q^{\text{KS}}$ and $\mathcal{S}_q^{\text{KS}}$, respectively, denoting

Toffoli-depth and total number of qubits in key search problem, i.e.,

$$(\mathcal{T}_q^{\text{KS}})^2 \mathcal{S}_q^{\text{KS}} = c_{\#}^{\text{KS}} N , \quad (22)$$

where $c_{\#}^{\text{KS}}$ varies depending on circuit designs. Now a parameter $c_{\#}^{\text{KS}}$ is the only ‘yardstick’ that tells us which design is better. When parallelized for large S_q , the expected iteration number converges to the one given in (9). Taking the converged value, the $c_{\#}^{\text{KS}}$ for each circuit design is summarized in Table 3. Assuming the MAXDEPTH is capped at some fixed value smaller than \sqrt{N} , the table indicates that for example AES-C1 requires about 28.6... times as many qubits as AES-C4.

5.4 Comparison to Ensured Single Target

It is possible to guarantee an existence of a single target by using multiple plaintext-ciphertext pairs. To ensure a single target, the oracle now performs r AES encryptions simultaneously. In [27], $r = 3$ is chosen for AES-128. Each AES box encrypts different plaintext with the same superposed input keys. As a result, for example for $r = 3$ in AES-128, the probability that two pre-images exist is the same as for k_1 to exist such that

$$\text{AES}(k_0, p_1) \parallel \text{AES}(k_0, p_2) \parallel \text{AES}(k_0, p_3) = \text{AES}(k_1, p_1) \parallel \text{AES}(k_1, p_2) \parallel \text{AES}(k_1, p_3) ,$$

where \parallel is concatenation, k_0 is the true key and p_i are distinct plaintexts. The cost of guaranteeing a single target is more or less multiplying the total number of qubits by r .

Table 4. Comparison of attack design with and without a single target. Toffoli-depth of encryption circuit is the same in both, because the same AES design is implemented. It is noticeable that the full Toffoli-depth of *Unique Key* is not far different from that of AES-C4, although the number of qubits is nearly doubled.

	AES-128		Grover	
	Toffoli-depth	Qubits	Toffoli-depth	Qubits
Unique Key	560	14296	$1.269 \dots \times 2^{74}$	14297
AES-C4	560	7148	$1.510 \dots \times 2^{74}$	7149

It is now natural to ask if the oracle operator with a single target is more cost-efficient than the random function oracle with less qubits. Assuming $r = 2$ guarantees a single target, we compare a design dubbed *Unique Key* with AES-C4. *Unique Key*’s encryption circuit design is chosen to be the same as AES-C4, meaning that the difference in efficiency solely comes from ensuring a single target. Results are summarized in Table 4. Full Toffoli-depth of *Unique Key* is estimated considering I_1 in (4). With a guaranteed single target, Toffoli-depth is

expected to be shortened compared with AES-C4 at the cost of doubling qubits. Although ensuring single target can be regarded as the optimization point when using single processor, it *strictly* cannot be an option in parallel attack since the inner parallelization removes a penalty of random characteristics as in (8).

6 Complexity of SHA-256 Pre-image Search

This section presumes that readers are familiar with standard SHA-256 [40]. The idea of applying Grover's algorithm to pre-image attack on SHA-256 is as follows.




Fig. 9. Oracle circuit for pre-image attack on SHA-256

A message block consisting of α bits of message and $512 - \alpha$ bits of padding are input to the hash box as shown in Fig. 9. Hash box contains a reversible circuit implementation of SHA-256 to permit superposed input. The input of linearly superposed 2^α messages are then passed on to the hash box resulting in superposed corresponding hashes. Processed hashes are then compared with the given hash via C^{256} NOT gate. After the target is marked, the entire qubits except the oracle qubit are further processed through the reverse hash box as in Fig. 9. The quantum state of the data qubits at the end of Fig. 9 reads

$$|\psi\rangle = \frac{1}{\sqrt{2^\alpha}} (|00 \dots 0\rangle + |00 \dots 1\rangle + \dots + |t_i\rangle + \dots + |11 \dots 1\rangle) \otimes |\text{padding}\rangle ,$$

where each ket state encodes a message and t_i 's are pre-images of the given hash value. The number of targets probabilistically varies depending on α which is capped at $447 (= 512 - 64 - 1)$.

6.1 Circuit Implementation Cost

SHA-256 internally performs five elementary operations, $\sigma_{0(1)}$, $\Sigma_{0(1)}$, Ch , Maj , and $ADDER$ (modular addition) [40].

Among internal operations carried out in SHA-256, $\Sigma_{0(1)}$ consists only of XOR-ings of bit permutations. Results of three $ROTR$ operations are written on 32-bit output register, with being successively XOR-ed. Only CNOT gates are involved in implementation with 32 work qubits.

Similarly, $\sigma_{0(1)}$ is implemented with one difference from $\Sigma_{0(1)}$, that is SHR . SHR itself is not linear, but writing a result of SHR on 32-bit output register is possible. Therefore, $\sigma_{0(1)}$ is also efficiently realized by CNOT gates with 32 work qubits.

Ch and Maj are bit-wise operations that do require Toffoli gates. We adopt Amy et al.'s design where Ch and Maj require one and two Toffoli gates, respectively. See Figs. 4 and 5 in [4].

Serial schedule-round implementation of SHA-256 is illustrated in Fig. 10.

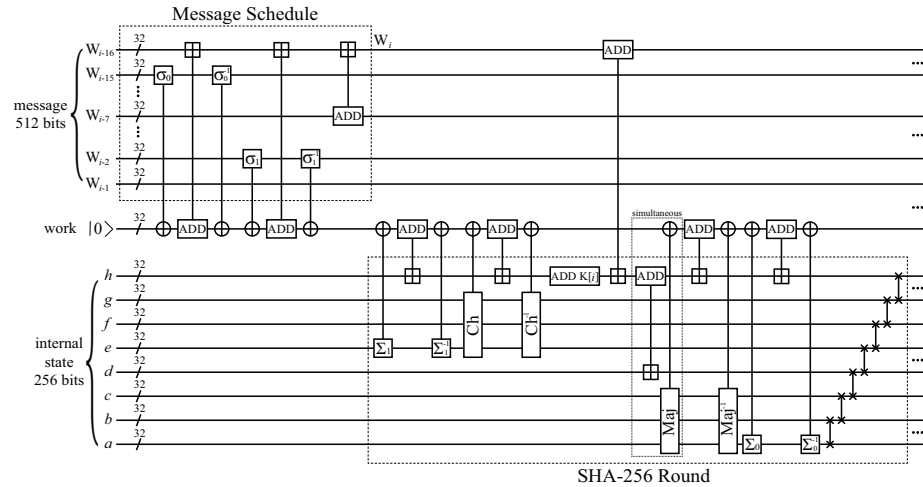


Fig. 10. Reversible circuit for serial implementation of SHA-256 message schedule and round function. The message block consisting of 16 words are recursively updated in-place. Note that it is straightforward to make message schedule and round functions work in parallel by expanding the work space. Seven two-qubit gates at the end of round are SWAP gates. The symbol \boxplus is addition modulo 2^{32} .

Low-level circuit design for each function in this work is mostly adopted from [4] except $ADDER$ choice and totally re-designed message schedule. A few options are available for $ADDER$ circuits one can adopt (see for example, [45]). For our purpose of comparing various circuit designs, we choose two versions of adders; a poly-depth $ADDER$ [20] and a log-depth $ADDER$ [22]. Table 5 summarizes resource costs of elementary operations in SHA-256.

Table 5. Costs of elementary operations in SHA-256. Work qubits in *ADDER* columns get cleaned within the respective Toffoli-depth. Outputs of $\sigma_{0(1)}$, $\Sigma_{0(1)}$, *Ch*, and *Maj* are written on work qubits.

	<i>ADDER</i> (poly)	<i>ADDER</i> (log)	$\sigma_{0(1)}, \Sigma_{0(1)}$	<i>Ch</i>	<i>Maj</i>
Toffoli-depth	61	22	0	1	2
Work qubits	1	53	32	32	32

6.2 Design Candidates

Three optimization points are considered.

First point, that has an impact on the overall design, is to determine whether message schedule and round functions are carried out in parallel. Figure 10 shows a serial circuit implementation of SHA-256. In the algorithm description, i -th round function is fed by i -th word from the schedule meaning that parallel implementation is possible if enough work qubits are given. This option is denoted by *serial/parallel schedule-round*.

Second point is to determine which *ADDER* is to be used. Use of the poly-depth *ADDER* is better in saving work space whereas the log-depth *ADDER* could shorten the execution time. For simplicity, we do not consider adaptive use of both *ADDER*s although it is possible to improve the efficiency by using appropriate *ADDER* in different part of circuit. This option is denoted by *poly-depth/log-depth ADDER*.

Lastly, it is now optional to decide how many work qubits are to be used to implement C^{256} NOT gate for marking the targets (hash comparison). As discussed in Sect. 4.3, C^k NOT gate can be one of the trade-offs points. However in AES-128, we do not need to consider C^{128} NOT as an optimization point seriously since the encryption process accompanies enough number of work qubits that can be reused in lower-depth C^{128} NOT gate. Situation is different in SHA-256. It is noticeable that hashing process of SHA-256 does not involve as many work qubits as AES-128, meaning that the lower-depth C^{256} NOT gate cannot be implemented unless more qubits are introduced solely for hash comparison. Toffoli-depth and work qubits required for lower-depth (less-qubit) C^{256} NOT gate are 509 (2024) and 254 (1), respectively. Note that lower-depth and less-qubit C^k NOT gates present here are only two extreme exemplary designs. This option is denoted by *less-qubit/lower-depth C^{256} NOT*.

In total, there exist 8 ($=2^3$) distinct circuit designs. We only analyze six of them since others do not seem to have merits. Six designs are denoted as follows.

- SHA-C1: Serial schedule-round, poly-depth *ADDER*, less-qubit C^{256} NOT
- SHA-C2: Serial schedule-round, log-depth *ADDER*, less-qubit C^{256} NOT
- SHA-C3: Serial schedule-round, log-depth *ADDER*, lower-depth C^{256} NOT
- SHA-C4: Parallel schedule-round, poly-depth *ADDER*, less-qubit C^{256} NOT
- SHA-C5: Parallel schedule-round, log-depth *ADDER*, less-qubit C^{256} NOT
- SHA-C6: Parallel schedule-round, log-depth *ADDER*, lower-depth C^{256} NOT

6.3 Comparison

Toffoli-depth and total number of qubits are carefully estimated for each design. The number of data qubits α has to be determined at this point. In our numerical calculation, $\alpha = 266$ seems to safely achieve the optimal expected iteration number given by Proposition 3 and to remove the failure probability. Costs of quantum SHA-256 hashing circuit and the entire Grover’s algorithm on a single quantum processor is summarized in Table 6. Estimates for single Grover iteration is omitted from the table as it can easily be calculated from costs of SHA-256 circuit;

$$\text{cost}(\text{Grover iteration}) = 2 \cdot \text{cost}(\text{SHA-256}) + \text{cost}(C^{256}\text{NOT}) + \text{cost}(C^{266}\text{NOT}) .$$

Table 6. Costs of SHA-256 hashing circuit and entire attack circuit on a single quantum processor. MAXDEPTH is not considered.

	SHA-256		Grover	
	Toffoli-depth	Qubits	Toffoli-depth	Qubits
SHA-C1	36368	801	$1.586 \dots \times 2^{143}$	802
SHA-C2	13280	853	$1.227 \dots \times 2^{142}$	854
SHA-C3	13280	853	$1.163 \dots \times 2^{142}$	1023
SHA-C4	27584	834	$1.216 \dots \times 2^{143}$	835
SHA-C5	10112	938	$1.919 \dots \times 2^{141}$	939
SHA-C6	10112	938	$1.792 \dots \times 2^{141}$	1023

Table 7. Comparison of trade-offs coefficients of different SHA-256 circuit designs without the oracle assumption. The smallest $c_{\#}^{\text{PS}}$ is found by SHA-C6 with $c_6^{\text{PS}} = 1.034 \dots \times 2^{38}$. Other values are divided by c_6^{PS} for easier comparison.

	SHA-C1	SHA-C4	SHA-C3	SHA-C2	SHA-C5	SHA-C6
$c_{\#}^{\text{PS}}/c_6^{\text{PS}}$	9.830...	6.015...	1.685...	1.565...	1.053...	1

Proposition 4 establishes the criterion for the comparison. Similar to (22), we replace T_q^{PS} and S_q^{PS} by $\mathcal{T}_q^{\text{PS}}$ and $\mathcal{S}_q^{\text{PS}}$, respectively, denoting Toffoli-depth and total number of qubits, i.e.,

$$(\mathcal{T}_q^{\text{PS}})^2 \mathcal{S}_q^{\text{PS}} = c_{\#}^{\text{PS}} N , \quad (23)$$

where $c_{\#}^{\text{PS}}$ varies depending on efficiency of circuits. When parallelized for large S_q , the expected iteration number converges to the one given in (12). Taking the converged value, $c_{\#}^{\text{PS}}$ for each design is summarized in Table 7. If MAXDEPTH is capped at some fixed value smaller than \sqrt{N} , the table indicates that for example SHA-C1 requires about 9.8... times as many qubits as SHA-C6.

7 Complexity of SHA-256 Collision Finding

Costs of two collision finding algorithms, GwDP and CNS, are to be estimated in this section. We adopt SHA-C6 which also turn out to be the most efficient in time-space complexity in GwDP and CNS algorithms⁶.

7.1 GwDP Algorithm

Estimating the cost of GwDP algorithm is straightforward. Basically this algorithm constructs a set of DPs by running multiple instances of Grover’s algorithm so that there occurs collision in the set. By using (13), costs of GwDP algorithm for selected number of machines are summarized in Table 8.

Table 8. Costs of GwDP algorithm for various number of machines. Note that the algorithm also requires classical memory of size $O(S_q)$.

S_q	Toffoli-depth	Qubits
2^2	$1.986 \dots \times 2^{141}$	4084
2^4	$1.985 \dots \times 2^{139}$	16272
2^8	$1.984 \dots \times 2^{135}$	258304
2^{16}	$1.981 \dots \times 2^{127}$	$6.508 \dots \times 10^7$
2^{32}	$1.975 \dots \times 2^{111}$	$4.127 \dots \times 10^{12}$
2^{64}	$1.963 \dots \times 2^{79}$	$1.732 \dots \times 10^{22}$

If T_q^{GwDP} and S_q^{GwDP} in Proposition 5 are replaced by Toffoli-depth $\mathcal{T}_q^{\text{GwDP}}$ and number of qubits $\mathcal{S}_q^{\text{GwDP}}$, the trade-offs curve reads

$$\mathcal{T}_q^{\text{GwDP}} \mathcal{S}_q^{\text{GwDP}} = c^{\text{GwDP}} \cdot N^{\frac{1}{2}}, \quad (24)$$

where c^{GwDP} is found to be $1.802 \dots \times 2^{25}$ by using $S_q = 2^{64}$ case.

7.2 CNS Algorithm

Proposition 6 suggests the optimal expected number of iterations in terms of t_L . The only extra work need to be done here is to determine t_L explicitly. From

⁶ Details on circuit comparisons in GwDP and CNS algorithms are dropped from the main text. An interesting point worth noticing is that SHA-C5 has small advantageous range of $S_q (< 2^8)$ over SHA-C6. The reason is that while SHA-C5 requires zero additional qubit in hash comparison, SHA-C6 needs $(256 - d - 2)$ qubits in comparison where d is the number of fixed bits in DP. Since d grows as S_q increases, there occurs crossover point. It is also noticeable that SHA-C6 cannot exactly fit into Proposition 5 for the same reason just mentioned, but deviation is small.

the definition of t_L , it reads

$$\begin{aligned}
t_L &= \frac{\text{cost}(S_{f_L})}{2^l \cdot \text{cost}(G)} , \\
\text{cost}(S_{f_L}) &= 2 \cdot \text{cost}(\text{SHA-256}) + 2^l \cdot \text{cost}\left(\text{C}^{(256-d)}\text{NOT}\right) , \\
\text{cost}(G) &= 2 \cdot \text{cost}(\text{SHA-256}) + \text{cost}\left(\text{C}^d\text{NOT}\right) + \text{cost}\left(\text{C}^{256}\text{NOT}\right) , \\
l &= \frac{d}{2} + \log_2\left(\frac{\pi}{2t_L}\right) , \quad d = \left\lfloor \frac{512}{5} + \frac{2}{5} \log_2\left(\frac{(2t_L)^3}{\pi}\right) \right\rfloor , \quad (25)
\end{aligned}$$

where G is Grover iteration. Numerical approach was taken to find t_L , d and l , which came out to be $0.015182\dots$, 96 and $54.538\dots$, respectively. By substituting these values for parameters in (14), the expected number of iterations becomes $I_{\text{rand}}^{\text{CNS}} = 1.856\dots \times 2^{102}$. Note that this value is somewhat different from that of Proposition 6 as d has been rounded off. Finally by multiplying $I_{\text{rand}}^{\text{CNS}}$ and the time cost of G , we obtain the total Toffoli-depth of CNS algorithm as

$$I_{\text{rand}}^{\text{CNS}} \cdot \text{cost}(G) = 1.184\dots \times 2^{117} . \quad (26)$$

Quantum space cost is cheaper than SHA-C6 because $\text{C}^{(256-d)}\text{NOT}$ gate used for list-comparison requires less work qubits than C^{256}NOT in pre-image search. It is estimated to be 939 qubits in total.

Table 9. Parameter values and costs of CNS algorithm for various number of machines. Note that the algorithm also requires $O(N^{1/5}S_q^{1/5})$ classical resources.

S_q	l	d	t_L	Toffoli-depth	Qubits
2^2	55.155...	97	0.015064...	$1.353\dots \times 2^{116}$	3756
2^4	55.558...	98	0.014987...	$1.203\dots \times 2^{115}$	15024
2^8	56.364...	99	0.014834...	$1.729\dots \times 2^{112}$	240384
2^{16}	57.976...	102	0.014527...	$1.960\dots \times 2^{107}$	$6.154\dots \times 10^7$
2^{32}	61.201...	109	0.013914...	$1.352\dots \times 2^{98}$	$4.033\dots \times 10^{12}$
2^{64}	67.654...	121	0.012692...	$1.100\dots \times 2^{79}$	$1.732\dots \times 10^{22}$

When parallelized, t_L slightly changes since l and d depend on $S_q (= 2^s)$, the number of machines. Modified l and d reads

$$l = \frac{d}{2} + \log_2\left(\frac{\pi}{2t_L}\right) , \quad d = \left\lfloor \frac{512 + 2s}{5} + \frac{2}{5} \log_2\left(\frac{1.291\dots (2t_L)^3}{\pi}\right) \right\rfloor ,$$

where t_L , $\text{cost}(S_{f_L})$ and $\text{cost}(G)$ are the same as in (25). We have estimated the quantum resource costs of CNS algorithm for a few S_q values as summarized in Table 9. Note that estimated time complexities are different from ones given by (15) as the equation is obtained for large S_q , and d here has been rounded off to the nearest integer. Due to the bound $s < \min(l, n - d - l)$, $S_q = 2^{66}$ is almost the maximum number of quantum machines Proposition 7 holds.

8 Security Strengths of AES and SHA-2

Based on the results of previous sections, quantum security strengths of AES and SHA-2 are drawn in this section. Three MAXDEPTH parameters, 2^{40} , 2^{64} , and 2^{96} , are adopted from [41]. Note that using these values of MAXDEPTH in our analysis is a conservative approach as our estimates only count Toffoli gates as time resources whereas NIST has counted all gates. Security strength of SHA-2 is determined by collision finding problem, not by pre-image search problem.

Table 10. Trade-offs coefficients of AES- k key search problem for $k \in \{128, 192, 256\}$ and SHA- m collision finding problem for $m \in \{256, 384, 512\}$. Coefficients c_k^{KS} and c_m^{CF} are divided by their respective minimal values $c_{128}^{\text{KS}} = c_4^{\text{KS}}$ and $c_{256}^{\text{CF}} = c^{\text{GwDF}}$.

	AES-128	AES-192	AES-256
$c_k^{\text{KS}} / c_{128}^{\text{KS}}$	1	1.560...	2.586...
	SHA-256	SHA-384	SHA-512
$c_m^{\text{CF}} / c_{256}^{\text{CF}}$	1	3.837...	3.940...

Resource estimates for AES-128 key search problem with circuit AES-C4 is extended to AES-192 and AES-256, and similarly that of SHA-256 collision finding problem with circuit SHA-C6 is applied to SHA-384 and SHA-512. Since the depth-qubit trade-offs curves (22) and (24) must hold for larger key and message digest sizes, we only compare their trade-offs coefficients in Tables 10. There is a tendency that the values of coefficients grow as the key or message digest sizes get larger. Increasing coefficient values reflect various complexity factors added; more rounds, longer schedule, larger word size, and so on. Especially in hash, size of the message block in SHA-384 is doubled compared with SHA-256 leading to large gap between c_{256}^{CF} and c_{384}^{CF} . In contrast, c_{384}^{CF} and c_{512}^{CF} do not show much difference as SHA-384 and SHA-512 algorithms are identical except truncation and initial values. The result of Sect. 7.2 is also extended to SHA-384 and reflected in Fig. 11.

Once trade-offs coefficients are obtained, we are able to draw the security strength of each algorithm in terms of required qubits as a function of Toffoli-depth. Note that somewhere between MAXDEPTH = 2^{64} and 2^{96} , security strengths of SHA-256 (SHA-384) and AES-192 (AES-256) are reversed in order, due to their different trade-offs curve behaviors. One minor note is that for large MAXDEPTH (for example, 2^{96}), Proposition 2 does not exactly hold since the size of the domain is larger than that of the codomain in AES-192 and AES-256. This factor is handled in a conservative way and reflected in Fig. 11.

Figure 11 summarizes the results which can be interpreted as another threshold to be used, for the security strength classification of proposed schemes in NIST PQC standardization process.

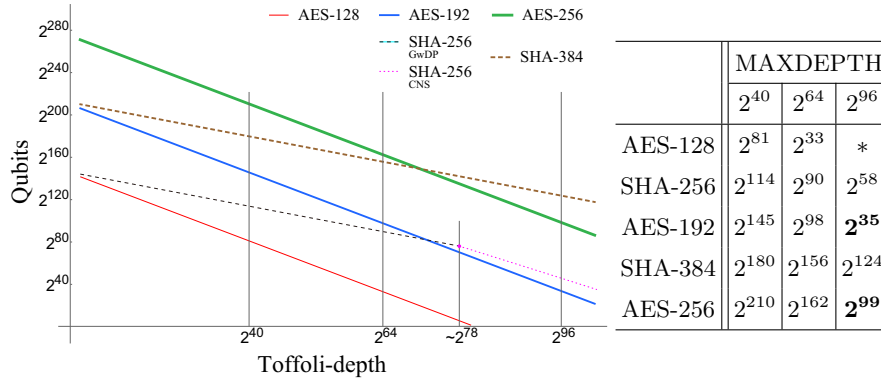


Fig. 11. Security strengths of AES and SHA-2. Values in the table are approximated number of qubits required to run the respective algorithm for given MAXDEPTH.

9 Summary

Instead of conventional query complexity, we have examined the time-space complexity of Grover’s algorithm and its variants. Three categories of cryptographic search problems and their characteristics are carefully considered in conjunction with probabilistic nature of quantum search algorithms.

To relate the time-space complexity with physical quantity, we have proposed a way of quantifying the computational power of quantum computers. Despite its simplicity, counting the number of sequential Toffoli gates reflects the reliable time complexity in estimating security levels of symmetric cryptosystems. With simplified cost measure, one can estimate the quantum complexity of a cryptosystem concisely by counting (and focusing) relevant operations only. It is worth noting that the above scheme is general for quantum resource estimates in symmetric cryptanalysis.

The scheme has been applied to resource estimates for AES and SHA-2. When multiple quantum trade-offs options are given, the time-space complexity provides clear criteria to tell which is more efficient. Based on the trade-offs observations made in AES and SHA-2, security strengths of respective systems are investigated with the MAXDEPTH assumption.

Acknowledgement. We are grateful to Brandon Langenberg, Martin Roetteler, and Rainer Steinwandt for helpful discussion and sharing details of their previous work which has motivated us.

References

1. Ambainis, A.: Quantum walk algorithm for element distinctness. *SIAM Journal on Computing* 37(1), 210–239 (2007)
2. Amento, B., Rötteler, M., Steinwandt, R.: Efficient quantum circuits for binary elliptic curve arithmetic: Reducing t-gate complexity. *Quantum Info. Comput.* 13(7-8), 631–644 (2013)
3. Amy, M., Maslov, D., Mosca, M., Roetteler, M.: A meet-in-the-middle algorithm for fast synthesis of depth-optimal quantum circuits. *IEEE Trans. Comput. Des. Integr. Circuits Syst.* 32(6), 818–830 (2013)
4. Amy, M., Di Matteo, O., Gheorghiu, V., Mosca, M., Parent, A., Schanck, J.: Estimating the cost of generic quantum pre-image attacks on SHA-2 and SHA-3. In: *SAC 2016*. pp. 317–337 (2017)
5. Arunachalam, S., Gheorghiu, V., Jochym-O’Connor, T., Mosca, M., Srinivasan, P.V.: On the robustness of bucket brigade quantum RAM. *New Journal of Physics* 17(12), 123010 (2015)
6. Barenco, A., Bennett, C.H., Cleve, R., DiVincenzo, D.P., Margolus, N., Shor, P., Sleator, T., Smolin, J.A., Weinfurter, H.: Elementary gates for quantum computation. *Phys. Rev. A* 52, 3457–3467 (1995)
7. Bernstein, D.J., Buchmann, J., Dahmen, E.: *Post Quantum Cryptography*, 1st edn. Springer Publishing Company, Incorporated (2008)
8. Bernstein, D.J., Lange, T.: Post-quantum cryptography. *Nature* 549(7671), 188–194 (2017)
9. Bernstein, D.J., Yang, B.Y.: Asymptotically faster quantum algorithms to solve multivariate quadratic equations. In: *PQCrypto 2018*. pp. 487–506 (2018)
10. Bernstein, D.: Cost analysis of hash collisions: Will quantum computers make SHARCS obsolete? In: *SHARCS ’09* (2009)
11. Beth, T., Rötteler, M.: *Quantum Algorithms: Applicable Algebra and Quantum Physics*, pp. 96–150. Springer Berlin Heidelberg (2001)
12. Bocharov, A., Svore, K.M.: Resource-optimal single-qubit quantum circuits. *Phys. Rev. Lett.* 109, 190501 (2012)
13. Boyer, M., Brassard, G., Høyer, P., Tapp, A.: Tight bounds on quantum searching. *Fortschritte der Physik* 46(4-5), 493–505 (1998)
14. Brassard, G., Høyer, P., Mosca, M., Tapp, A.: Quantum amplitude amplification and estimation. *AMS Contemp. Math.* 305, 53–74 (2002)
15. Brassard, G., Høyer, P., Tapp, A.: Quantum cryptanalysis of hash and claw-free functions. In: *LATIN ’98*. pp. 163–169 (1998)
16. Bravyi, S., Haah, J.: Magic-state distillation with low overhead. *Phys. Rev. A* 86, 052329 (2012)
17. Bravyi, S., Kitaev, A.: Universal quantum computation with ideal Clifford gates and noisy ancillas. *Phys. Rev. A* 71, 022316 (2005)
18. Chailloux, A., Naya-Plasencia, M., Schrottenloher, A.: An efficient quantum collision search algorithm and implications on symmetric cryptography. In: *ASIACRYPT 2017*. pp. 211–240 (2017)
19. Childs, A.M., Goldstone, J.: Spatial search by quantum walk. *Phys. Rev. A* 70, 022314 (2004)
20. Cuccaro, S.A., Draper, T.G., Kutin, S.A., Moulton, D.P.: A new quantum ripple-carry addition circuit. *arXiv preprint quant-ph/0410184* (2004)
21. Deutsch, D.: Quantum computational networks. *Proceedings of the Royal Society of London A* 425(1868), 73–90 (1989)

22. Draper, T.G., Kutin, S.A., Rains, E.M., Svore, K.M.: A logarithmic-depth quantum carry-lookahead adder. *Quantum Info. Comput.* 6(4 & 5), 351–369 (2006)
23. Flajolet, P., Odlyzko, A.M.: Random mapping statistics. In: EUROCRYPT '89. pp. 329–354. Springer Berlin Heidelberg (1990)
24. Fowler, A.G., Mariantoni, M., Martinis, J.M., Cleland, A.N.: Surface codes: Towards practical large-scale quantum computation. *Phys. Rev. A* 86, 032324 (2012)
25. Gosset, D., Kliuchnikov, V., Mosca, M., Russo, V.: An algorithm for the T-count. *Quantum Info. Comput.* 14(15-16), 1261–1276 (2014)
26. Gottesman, D.: The Heisenberg representation of quantum computers. In: GP22: ICGTMP '98. pp. 32–43 (1998)
27. Grassl, M., Langenberg, B., Roetteler, M., Steinwandt, R.: Applying Grover's algorithm to AES: Quantum resource estimates. In: PQCrypto 2016. pp. 29–43 (2016)
28. Grover, L.K.: Quantum mechanics helps in searching for a needle in a haystack. *Phys. Rev. Lett.* 79, 325–328 (1997)
29. Häner, T., Roetteler, M., Svore, K.M.: Factoring using $2n + 2$ qubits with Toffoli based modular multiplication. *Quantum Info. Comput.* 17(7 & 8), 673–684 (2017)
30. Hosoyamada, A., Sasaki, Y., Xagawa, K.: Quantum multicollision-finding algorithm. In: ASIACRYPT 2017. pp. 179–210 (2017)
31. Jones, C.: Multilevel distillation of magic states for quantum computing. *Phys. Rev. A* 87, 042305 (2013)
32. Jones, N.C., Van Meter, R., Fowler, A.G., McMahon, P.L., Kim, J., Ladd, T.D., Yamamoto, Y.: Layered architecture for quantum computing. *Phys. Rev. X* 2, 031007 (2012)
33. Katz, J., Lindell, Y.: *Introduction to Modern Cryptography*, 2nd edn. Chapman & Hall/CRC (2007)
34. Kepley, S., Steinwandt, R.: Quantum circuits for \mathbb{F}_2^n -multiplication with sub-quadratic gate count. *Quantum Information Processing* 14(7), 2373–2386 (2015)
35. Maslov, D., Mathew, J., Cheung, D., Pradhan, D.K.: An $O(m^2)$ -depth quantum algorithm for the elliptic curve discrete logarithm problem over $\text{GF}(2^m)^a$. *Quantum Info. Comput.* 9(7), 610–621 (2009)
36. Meier, A.M., Eastin, B., Knill, E.: Magic-state distillation with the four-qubit code. *Quantum Info. Comput.* 13(3-4), 195–209 (2013)
37. Moody, D.: Let's get ready to rumble - the NIST PQC "competition". PQCrypto 2018 invited presentation (2018)
38. Nielsen, M.A., Chuang, I.L.: *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press (2010)
39. NIST: Advanced Encryption Standard (AES). FIPS PUB 197 (2001)
40. NIST: Secure Hash Standard (SHS). FIPS PUB 180-4 (2015)
41. NIST: Post-Quantum Cryptography- Call for Proposals (2017)
42. van Oorschot, P.C., Wiener, M.J.: Parallel collision search with cryptanalytic applications. *Journal of Cryptology* (1999)
43. Parent, A., Roetteler, M., Svore, K.M.: Reversible circuit compilation with space constraints. arXiv preprint arXiv:1510.00377 (2015)
44. Patel, K.N., Markov, I.L., Hayes, J.P.: Optimal synthesis of linear reversible circuits. *Quantum Info. Comput.* 8(3), 282–294 (2008)
45. Pavlidis, A., Gizopoulos, D.: Fast quantum modular exponentiation architecture for Shor's factoring algorithm. *Quantum Info. Comput.* 14(7 & 8), 649–682 (2014)
46. Roetteler, M., Naehrig, M., Svore, K.M., Lauter, K.: Quantum resource estimates for computing elliptic curve discrete logarithms. In: ASIACRYPT 2017. pp. 241–270 (2017)

47. Schwabe, P., Westerbaan, B.: Solving binary \mathcal{MQ} with Grover's algorithm. In: SPACE 2016. pp. 303–322 (2016)
48. Shenvi, N., Kempe, J., Whaley, K.B.: Quantum random-walk search algorithm. *Phys. Rev. A* 67, 052307 (2003)
49. Shi, Y.: Both Toffoli and controlled-NOT need little help to do universal quantum computation. arXiv preprint quant-ph/0205115 (2002)
50. Toffoli, T.: Reversible computing, pp. 632–644. Springer Berlin Heidelberg (1980)
51. Vidal, G.: Efficient classical simulation of slightly entangled quantum computations. *Phys. Rev. Lett.* 91, 147902 (2003)
52. Yoder, T.J., Low, G.H., Chuang, I.L.: Fixed-point quantum search with an optimal number of queries. *Phys. Rev. Lett.* 113, 210501 (2014)
53. Zalka, C.: Grover's quantum searching algorithm is optimal. *Phys. Rev. A* 60, 2746–2751 (1999)