Improvements to low-qubit quantum resource estimates for quantum search

James H. Davenport and Benjamin Pring

University of Bath, Bath, BA2 7AY, UK
b.i.pring@bath.ac.uk

Abstract. In this paper we examine and solve an open problem regarding the results on the *Search with Two Oracles* (*STO*) problem by Kimmel et al. [14], which demonstrates how two oracles with different costs, which we interpret as circuit-complexity, can be used to provide lower quantum resource estimates for quantum search. Their solution to the *STO* problem relied upon exact knowledge of the number of targets that one of the oracles marks and the method can fail if this is unknown. We demonstrate how to adapt their solution to the realistic case where we only have knowledge concerning the probability distribution on the number of elements that that are marked by the cheaper quantum oracle. We apply these methods to both the single-target AES [12] and Multivariate Quadratic problem over $\mathbb{F}_2$ [23] preimage search problems to obtain a lower circuit-complexity for low qubit implementations of quantum search compared to solving these problems using Grover's algorithm [13].

Keywords: quantum search, AES, multivariate quadratic, cryptanalysis

1 Introduction

The solution to many problems in computer science and cryptanalysis can be reduced to the *unstructured search problem* (see Definition 1), which can be solved via classical computers for an expected cost of $O(\frac{2^n}{M_\chi})$ classical queries [1] or infamously by *Grover's algorithm* [13] for a cost of $O(\sqrt{\frac{2^n}{M_\chi}})$ quantum queries.

Definition 1 (Unstructured search problem).
Let $\chi : \{0,1\}^n \longrightarrow \{0,1\}$ and $M_\chi = |\chi^{-1}(1)|$. The unstructured search problem is, given only the ability to evaluate χ on elements of the search domain $\{0,1\}^n$ to find an element $x \in \{0,1\}^n$ such that $\chi(x) = 1$ or prove no such x exists.

This is an abstract definition of the search problem, as in many real-world examples we have access to the inner workings of $\chi : \{0,1\}^n \longrightarrow \{0,1\}$. An equivalent formulation of the search problem is the *preimage search problem*

Definition 2 (The preimage search problem).
Let $h : \{0,1\}^n \longrightarrow \{0,1\}^m$, $Y \subseteq \{0,1\}^m$ and $M_h = |h^{-1}(Y)|$. The preimage search problem is to find an $x \in \{0,1\}^n$ such that $h(x) \in Y$ or prove no such x exists. The single-target preimage search problem is the case where $|Y| = 1$.

By equivalent, we mean that any solution that solves the unstructured search problem can be used to solve the preimage search problem — and vice versa. It is clear that that the preimage search problem inherently has more structure that allows us to make more design choices, such as choosing the method that we use to perform the *set membership* test $h(x) \in Y$. Structure is fairly common in any search problem — though we can only exploit it to provide a polynomial reduction in the resources required to solve the search problem, compared to the superpolynomial reductions that are asymptotically interesting. Yet for real-world problems, a polynomial or constant reduction in resources is important — and in cryptography if we are extrapolating cryptographic parameters based upon the concrete hardness of solving certain problems, these asymptotically negligible reductions are of great importance.

The dangers of extrapolating key-sizes based upon the concrete resources required to implement quantum search for the Multivariate Quadratic problem has been previously studied by one of the authors [21] and in this paper we bring new optimisations of quantum search to light, which impact upon the resources required for quantum cryptanalysis of the Advanced Encryption Standard [22] in the logical quantum circuit-model of computation. This an important problem, as the current NIST post-quantum standardisation effort [25] ties the security of submissions to the resources required to break AES [26]. These quantum resources have previously been quantified in literature by examining the logical quantum circuit-complexity (in terms of the Clifford+T universal quantum gate set) by using Grover's quantum search algorithm [13] to solve this problem [12].

1.1 The impact of quantum search on cryptographic parameters

Grover's quantum search algorithm [13] solves the unstructured search problem (see Definition 1) and its execution cost $E_{\mathcal{G}}$ (the notation $E_{\mathcal{A}}$ can be thought of as denoting either circuit-depth or circuit-size of the quantum subroutine $\mathcal{A}$) can be easily derived by using the formula

$$E_{\mathcal{G}} = \underbrace{E_{H^{\otimes n}}}_{\text{Setup phase}} + \underbrace{\left\lfloor \frac{\pi}{4} \cdot \sqrt{\frac{2^n}{M_\chi}} \right\rfloor}_{\text{Query complexity}} \cdot \Bigg(\underbrace{E_{\mathcal{O}_\chi}}_{\text{Quantum oracle}} + \underbrace{2E_{H^{\otimes n}} + \mathcal{O}_{\bar{n}}}_{\text{Diffusion step}} \Bigg), \quad (1)$$

where the diffusion step is a standard piece of quantum circuitry requiring $O(n)$ quantum gates to implement. The quantum oracle can be implemented by a quantum circuit which implements a reversible implementation of the boolean function $\chi : \{0,1\}^n \longrightarrow \{0,1\}$ and quantum resource estimates often (if not always) estimate the cost of search problems by considering this methodology of quantum oracle construction, relying upon a simple conversion dictionary from a set of universal boolean gates to a set of quantum gates which implement reversible versions of these classical gates

$$\underbrace{\{\neg, \oplus, \wedge\}}_{\text{Classical gates}} \leftrightarrow \underbrace{\{X, \wedge_1(X), \wedge_2(X)\}}_{\text{Quantum gates}}. \quad (2)$$

It is evident that a number of the submissions to the NIST post quantum standardisation effort [25] use this approach [9, 10], taking into account the cost of both the query-complexity and the cost of the quantum oracle in order to extrapolate the cost of quantum cryptanalysis via Grover's algorithm. Other submissions take a safer approach, acknowledging the true cost of quantum search but using lower-bounds on the resources required to perform cryptanalysis via quantum search [8]. The NIST call for proposals [26] itself ties the security of entries to the standardisation process to number of quantum gates that must be executed to perform cryptanalysis of the Advanced Encryption Standard [22] (AES) via quantum search, extrapolating the parameters from a previously performed quantum resource estimation using Grover's quantum search algorithm [12].

However, many optimisations and variants of quantum search exist in literature [15, 3] and Grover's algorithm itself is now viewed as simply a special case of the quantum subroutine of *amplitude amplification* [7]. In this paper we focus upon applying and improving a previously suggested method to improve quantum search when there exists structure in the problem [14].

Contributions In this paper we solve an open problem with regards to the *Search with Two Oracles* problem [14] for quantum search and demonstrate that meta-optimisations of quantum search can have an important impact upon the resources required to perform quantum cryptanalysis when we have both a bounded and unbounded number of qubits at our disposal. Whilst less important in terms of the NIST competition, a very real concern (for both industrial uses of quantum computing and in cryptanalysis) is the timeline of when quantum computers of certain sizes will be physically realisable — we demonstrate that certain problems suffer far less from the constraint of few qubits being available than previously thought.

Outline of this paper For reasons of space, we leave many details to the final paper. In Section 2 we briefly review of the costs involved in quantum search and amplitude amplification. In Section 3 we review the definition and existing solution to the *Search with Two Oracles* problem, whilst in Section 4 we provide our modification. In Section 5 we conclude with new quantum resource estimates for low qubit implementations of quantum search applied to cryptanalysis of AES and the Multivariate Quadratic problem over $\mathbb{F}_2$ and give our conclusions.

2 Amplitude amplification and quantum search

For reasons of space, we provide only a brief introduction to basic theory. Details may be found in standard resources [17]. Quantum algorithms act upon quantum states consisting of qubits. An n-qubit quantum state may be written

$$|\psi\rangle = \sum_{x \in \{0,1\}^n} \alpha_x |x\rangle \qquad \text{where} \quad \alpha_x \in \mathbb{C} \quad \text{and} \quad \sum_{x \in \{0,1\}^n} |\alpha_x|^2 = 1. \quad (3)$$

Measurement of this state results in a single bitstring $x \in \{0,1\}^n$ with probability $|\alpha_x|^2$. Quantum algorithms manipulate the *amplitudes* $\{\alpha_x : x \in \{0,1\}^n\}$ of the *computational basis states* $\{|x\rangle : x \in \{0,1\}^n\}$ so that measurement of the quantum state results in a bitstring which encodes useful information with high probability. Computational basis states can be interpreted as bitstrings and *reversible* boolean functions (permutations) that act upon these basis states can be implemented using a subset of *quantum gates* $\{X, \wedge_1(X), \wedge_2(X)\}$ (the X, CNOT and Toffoli gates), analogues of the universal boolean gate set $\{\neg, \oplus, \wedge\}$.

We additionally use the H (Hadamard), $\wedge_1(SWAP)$ (also known as the Fredkin or controlled SWAP gate) and $\wedge_k(X)$ gates (for $k > 2$). X and $\wedge_k(X)$ act upon 1 and $k+1$ qubit states by mapping

$$X|x_1\rangle \mapsto |x_1 \oplus 1\rangle \quad \text{and} \quad \wedge_k(X)|x_1 \ldots x_n\rangle|y\rangle \mapsto |x_1 \ldots x_n\rangle|y \oplus (x_1 \wedge \cdots \wedge x_n)\rangle. \quad (4)$$

Each of these quantum gates can be decomposed in terms of a *universal quantum gate set* [17]. In terms of the Clifford+T universal quantum gate set, the $X, \wedge_1(X)$ and H require a single quantum gate, whilst $\wedge_2(X)$ and $\wedge_1(SWAP)$ require 17 quantum gates [24, 2] and $\wedge_k(X)$ requires $72k - 84$ (for $k \geq 4$) quantum gates [16]. Full details will be in the final paper, but these gates are sufficient to implement *quantum phase oracles*, an integral component of amplitude amplification, if we possess a classical circuit to solve the unstructured search problem.

Definition 3 (Quantum phase oracle).
Let $\chi : \{0,1\}^n \longrightarrow \{0,1\}$. The quantum phase oracle $\mathcal{O}_\chi$ performs the following action upon the n-qubit computational basis state $|x\rangle$, where $x \in \{0,1\}^n$

$$\mathcal{O}_\chi|x\rangle \mapsto (-1)^{\chi(x)}|x\rangle \quad (5)$$

Quantum phase oracles can be realised by *quantum evaluations*.

Definition 4 (Quantum evaluation).
Let $h : \{0,1\}^n \longrightarrow \{0,1\}^m$. The quantum evaluation $\mathcal{E}_h$ performs the following action upon the $n + w + m$-qubit quantum state, where $b \in \{0,1\}^m$

$$\mathcal{E}_h|x\rangle|0^w\rangle|b\rangle \mapsto |x\rangle|g(x)\rangle|b \oplus h(x)\rangle, \quad (6)$$

where $g(x) \in \{0,1\}^w$ is the end state of the memory used to compute $h(x)$.

A quantum phase oracle can be implemented by executing $\mathcal{E}_\chi$ on the state $|x\rangle|0^w\rangle|0\rangle$ and using a single qubit Z gate on the last qubit, which maps $Z|\chi(x)\rangle \mapsto (-1)^{\chi(x)}|\chi(x)\rangle$. The quantum evaluation can then be uncomputed by executing $\mathcal{E}_\chi^\dagger$, which is $\mathcal{E}_\chi$ in reverse.

2.1 The cost of implementing amplitude amplification

We will use the notation $E_\mathcal{A}$ to represent the execution cost of an arbitrary quantum algorithm or gate $\mathcal{A}$. All costs denoted this way will be components that must be executed in serial, hence the notation $E_\mathcal{A}$ can represent either circuit-size, circuit-complexity or communication overheads between qubits. The statement and cost of amplitude amplification can now be given.

Definition 5 (Success probability of a quantum algorithm).
The success probability of a measurement-free quantum algorithm $\mathcal{A}$ relative to the boolean function $\chi : \{0,1\}^n \longrightarrow \{0,1\}$ is the probability that measuring the quantum state $\mathcal{A}\left|0^n\right\rangle$ results in a bitstring $x \in \{0,1\}^n$ such that $\chi(x) = 1$.

Hence if we are searching for the unique $y \in \{0,1\}^n$ such that $\chi(x) = 1$ and

$$\mathcal{A}\left|0^n\right\rangle \mapsto \sum_{x\in\{0,1\}^n} \alpha_x \left|x\right\rangle, \tag{7}$$

then the success probability of $\mathcal{A}$ relative to χ is $|\alpha_y|^2$.

Theorem 1 (Amplitude amplification [7]).
Let $\mathcal{A}$ be any measurement-free quantum algorithm and the quantum oracles $\mathcal{O}_\chi$ and $\mathcal{O}_{\bar{0}}$ be defined by $\chi, \bar{n} : \{0,1\}^n \longrightarrow \{0,1\}$ where $\bar{n}(x) \mapsto 1$ iff $x \neq 0^n$.
Let $a > 0$ be the success probability of $\mathcal{A}$ relative to χ and $k \in \mathbb{N}_0$. Then there exists a quantum algorithm $\mathcal{B}(k)$ for which the success probability of $\mathcal{B}(k)$ relative to χ is $\sin^2\left((2k+1)\cdot \arcsin\sqrt{a}\right)$. The quantum algorithm $\mathcal{B}(k)$ has a cost of

$$E_{\mathcal{A}} + k\left(E_{\mathcal{O}_\chi} + 2E_{\mathcal{A}} + E_{\mathcal{O}_{\bar{n}}}\right). \tag{8}$$

Knowledge of a gives us the quadratic advantage in query-complexity for Grover.

Theorem 2 (Quadratic speedup [7, 6]).
Let the conditions be as Theorem 1 and $a \leq \frac{1}{2}$. Then choosing $k = \left\lfloor \frac{\pi}{4\arcsin\sqrt{a}} \right\rfloor$ gives $\mathcal{B}(k)$ a success probability of at least $1 - a$ relative to χ.

Grover's algorithm [13] is simply Theorem 1 with $\mathcal{A} := H^{\otimes n}$ (the Hadamard transform) which applied to the state $\left|0^n\right\rangle$ produces the *uniform superposition*

$$\frac{1}{2^{n/2}} \sum_{x\in\{0,1\}^n} \left|x\right\rangle. \tag{9}$$

This means that if $M_\chi = |\chi^{-1}(1)| = 1$, then the success probability of $H^{\otimes n}$ relative to χ is $\frac{1}{2^n}$. Theorem 2 then defines the cost and success probability. The *Search with Two Oracles* (*STO*) problem in Section 3 uses *exact amplitude amplification* to create a deterministic algorithm to solve the *STO* problem.

Theorem 3 (Exact amplitude amplification [7]).
Let the conditions be as in Theorem 1 and the success probability of $\mathcal{A}$ relative to χ be known. Then there exists a quantum algorithm $\mathcal{B}(k)$ with approximately the same cost formula as (8) and for which $k' = \left\lceil \frac{\pi}{4\arcsin\sqrt{a}} \right\rceil$ gives a success probability of $\mathcal{B}(k')$ relative to χ of exactly 1.

Theorem 3 will be used to sketch the previously proposed solution [14] to the *Search with Two Oracles* (see Definition 11) problem in Section 3.1. We will use only amplitude amplification (Theorem 1) for our adaptation as we will not know the exact success probabilities involved at several stages, hence the minor extra cost that exact amplitude amplification incurs will serve no purpose.

3 The Search with Two Oracles problem

In this section we examine the original formulation of the Search with Two Oracles (*STO*) problem by Kimmel et al. [14] dealing with single-target search. We recall that $E_{\mathcal{A}}$ denotes the cost (either circuit-size or circuit-depth) of the quantum algorithm $\mathcal{A}$ and so $E_{\mathcal{O}_\chi}$ is the cost of the quantum phase oracle $\mathcal{O}_\chi$.

Returning to the cost formula for Grover's algorithm given by (1), the cost of the single-target search case can be written as

$$\approx \frac{\pi}{4} \cdot 2^{n/2} \cdot \left(E_{\mathcal{O}_\chi} + 2E_{H^{\otimes n}} + E_{\mathcal{O}_{\bar{n}}} \right) \tag{10}$$

where the $\frac{\pi}{4} \cdot 2^{n/2}$ terms comes from the query-complexity of Grover's algorithm and the multiplicative overhead comes from the requirement that we must implement the quantum oracle and diffusion step. It is clear that if the quantum oracle $E_{\mathcal{O}_\chi}$ is expensive then this will add a large overhead to the cost executing Grover's algorithm. Whilst Grover's algorithm considers the *unstructured search problem* (see Definition 1), the *STO* problem examines how we can reduce this overhead when there exists structure in the problem.

Definition 6 (Search with Two Oracles (*STO*) [14]).
Let $f_*, f_S : \{0,1\}^n \longrightarrow \{0,1\}$ be two boolean functions with the property that

$$f_*^{-1}(1) \subseteq f_S^{-1}(1) \quad \text{where} \quad M_* = |f_*^{-1}(1)| \in \{0,1\} \quad \text{and} \quad M_S = |f_S^{-1}(1)| \tag{11}$$

and which respectively define the quantum oracles $\mathcal{O}_*, \mathcal{O}_S$. The Search with Two Oracles (STO) problem is to locate an element $x \in \{0,1\}^n$ such that $f_*(x) = 1$ or prove that no such element exists. It is given that $E_{\mathcal{O}_*} \geq E_{\mathcal{O}_S}$.

This is a realistic scenario in many cases. For a concrete example in classical computing, we consider the problem of finding a solution to a system of m equations in n variables over the finite field $\mathbb{F}_2$ by classical exaustive search. In this case, evaluation of a subset of $r < m$ equations on a potential solution $(x_1, \ldots, x_n) \in (\mathbb{F}_2)^n$ corresponds to a cheap test and evaluation of all m equations corresponds to the expensive test. It is clear that the set of solutions to the full set of m equations is a subset of the solutions to $r < m$ equations. We therefore need only perform an expensive test on a potential solution if it has passed the cheap test — in essence we can perform a filtering strategy.

This strategy is employed in the *Fast Exhaustive Search* (FES) algorithm [4, 5] for enumeration of solutions to systems of multivariate quadratic equations over $\mathbb{F}_2$ to reduce the dependence of the complexity of the search process upon m. Crucially such strategies impact only upon the cost of testing, not the total number of elements that we test, hence the query-complexity remains unchanged.

The cost of implementing quantum oracles for the problem of solving degree-two equations over the finite field $\mathbb{F}_2$ has previously been quantified [23, 21], demonstrating that both expensive and cheap quantum phase oracles exist.

Kimmel et al. [14] assume that we know M_S and if $E_{\mathcal{O}_*} = E_{\mathcal{O}_S}$, suggest we simply use Grover's algorithm with the quantum oracle $\mathcal{O}_*$ and ignore $\mathcal{O}_S$. Otherwise, the procedure in the following section is suggested.

3.1 A cost effective solution for *STO*

Define the quantum algorithm $\mathcal{A} = H^{\otimes n}$ (the Walsh-Hadamard transform) and use exact amplitude amplification to create a quantum algorithm $\mathcal{B}$ with success probability 1 relative to the function f_S. By Theorem 3, $\mathcal{B}$ requires $k_1 = \left\lceil \frac{\pi}{4 \arcsin \sqrt{\frac{M_S}{2^n}}} \right\rceil \approx \frac{\pi}{4}\sqrt{\frac{2^n}{M_S}}$ applications of $\mathcal{O}_S$ and $2k_1 + 1$ applications of $\mathcal{A}$. It is plain the the success probability of $\mathcal{B}$ relative to f_* is $b = \frac{1}{M_S}$.

We can then define a second quantum algorithm, $\mathcal{C}$, by using exact amplitude amplification with the quantum algorithm set to be $\mathcal{B}$, which has a success probability relative to f_* of $\frac{1}{M_S}$. By Theorem 1, we can create a quantum algorithm $\mathcal{C}$ with a success probability of 1 relative to f_*. By Theorem 3, $\mathcal{C}$ requires $k_2 = \left\lceil \frac{\pi}{4 \arcsin \sqrt{\frac{1}{M_S}}} \right\rceil \approx \frac{\pi}{4}\sqrt{M_S}$ applications of $\mathcal{O}_*$ and $2k_2 + 1$ applications of $\mathcal{B}$. This has an approximate cost of $\frac{\pi}{4}\sqrt{M_S} \cdot E_{\mathcal{O}_*} + \frac{\pi^2}{8}\sqrt{2^n} \cdot E_{\mathcal{O}_S} + \frac{\pi^2}{4}\sqrt{2^n} \cdot E_{H^{\otimes n}}$ and so if $E_{H^{\otimes n}} + E_{\mathcal{O}_S} \ll E_{\mathcal{O}_S}$ then we have an efficient solution. If we guess that $M_S = M'_S$, then the probability of success of $\mathcal{C}$ can be shown to be c, where

$$c = \sin^2\left(\left(2\hat{k}_2 + 1\right) \cdot \arcsin\sqrt{z \cdot \frac{M'_S}{M_S} \cdot \sin^2\left(\frac{\pi}{4\hat{k}_2 + 2}\right)}\right) \cdot \left(\frac{b_g - b \cdot b_g}{b_g - b \cdot \hat{b}_g}\right) + \frac{b \cdot b_g - b \cdot \hat{b}_g}{b_g - b \cdot \hat{b}_g} \tag{12}$$

where $b_g = \frac{1}{M'_S}$, $\hat{k}_2 = \left\lceil \frac{\pi}{4 \arcsin \sqrt{b_g}} \right\rceil$, $\hat{b}_g = \sin^2\left(\frac{\pi}{4\hat{k}_2 + 2}\right)$, $b = \frac{z}{M_S}$ and where

$$z = \sin^2\left(\left(2\hat{k}_1 + 1\right) \cdot \arcsin\sqrt{\frac{M_S}{M'_S} \cdot \sin^2\left(\frac{\pi}{4\hat{k}_1 + 2}\right)}\right) \cdot \left(\frac{a_g - a \cdot a_g}{a_g - a \cdot \hat{a}_g}\right) + \frac{a \cdot a_g - a \cdot \hat{a}_g}{a_g - a \cdot \hat{a}_g}. \tag{13}$$

where $a_g = \frac{M'_S}{2^n}$, $\hat{k}_1 = \left\lceil \frac{\pi}{4 \arcsin \sqrt{a_g}} \right\rceil$, $\hat{a}_g = \sin^2\left(\frac{\pi}{4\hat{k}_1 + 2}\right)$ and $a = \frac{M_S}{2^n}$.

Two errors are introduced which stem from the ratio $M_S : M'_S$. This leads to the algorithm potentially terminating with a high probability of failure if this ratio is large. We suggest a modification to control this ratio in the next section.

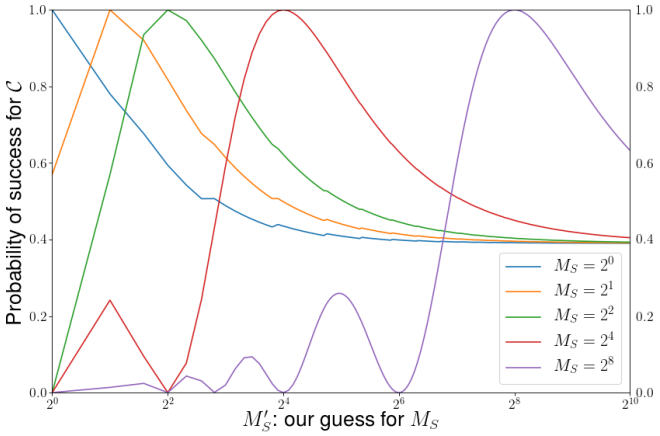

Fig. 1: Success of $\mathcal{C}$ with $n = 128$, $M_* = 1$ and varied M_S.

4 Adapting *STO* to problem instances

Essentially, our modification uses only amplitude amplification (see Theorem 1) in the same way as the original *STO* algorithm uses exact amplitude amplification (see Theorem 3), but tames the ratio $M_S : M'_S$ by artifically introducing new targets into the intermediate search space, so that the ratio becomes $(M_S + 2^t) : (M'_S + 2^t)$, which approaches 1 as t increases. This introduction of artificial targets can be implemented by a relatively cheap operation.

Theorem 4 (A modified solution to the *STO* problem).
Let $f_S, f_* : \{0,1\}^n \longrightarrow \{0,1\}$, $M_* = |f_*^{-1}(1)| = 1$ and $M_S = |f_S^{-1}(1)|$ define an instance of the STO problem, where M_S is unknown. Let $M'_S \in \mathbb{N}$ and $\epsilon \in [0,1]$ be such that $Pr[M_S \geq M'_S] \leq \epsilon$ and let $t \in \mathbb{N}$ be such that $0 \leq t \leq n$. Let $\mathcal{O}_{\bar{n}}$ be as in Theorem 1. Then there exists a quantum algorithm with a success probability relative to f_* of at least

$$(1-\epsilon) \cdot \sin^2\left(\left(\frac{\pi}{2\arcsin\sqrt{\frac{1}{2^t}}} - 2\right) \cdot \sqrt{\frac{b'}{M'_S + 2^t}}\right) \tag{14}$$

where

$$b' = \sin^2\left(\left(\frac{\pi}{2\arcsin\sqrt{\frac{M'_S+2^t}{2^n}}} - 2\right) \cdot \sqrt{\frac{2^t}{2^n}}\right). \tag{15}$$

The algorithm has a total execution cost of

$$\begin{aligned} & k_2\left(E_{\mathcal{O}_{f_*}} + E_{\bar{n}}\right) \\ & + (2k_1+1)(2k_2+1)E_{H^{\otimes n}} \\ & + (2k_2+1)k_1(E_{\mathcal{O}_{f_S}} + E_{\wedge_{n-t}(X)} + E_{\wedge_{n-t+1}(X)} + E_{\wedge_1(X)} + E_{\bar{n}}) \end{aligned} \tag{16}$$

where

$$k_1 = \left\lfloor \frac{\pi}{4\arcsin\sqrt{\frac{M'_S+2^t}{2^n}}} - \frac{1}{2} \right\rfloor \qquad \textit{and} \qquad k_2 = \left\lfloor \frac{\pi}{4\arcsin\sqrt{\frac{1}{2^t}}} - \frac{1}{2} \right\rfloor. \tag{17}$$

Whilst somewhat unwieldy, this formulation allows us to easily derive the success probabilities and costs for particular problem instances via numerical simulation.

Proof: Let $S = f_S^{-1}(1)$ and define $Z_t = \{1^{n-t}\|x' : x' \in \{0,1\}^t\}$ to be the set of 2^t bitstrings of length n whose first $n-t$ values are 1.
Let $f_{S\cup Z_t} : \{0,1\}^n \longrightarrow \{0,1\}$ where $f_{S\cup Z_t}(x) \mapsto 1$ if and only if $f_S(x) = 1$ or $x \in Z_t$. Let $M_{S\cup Z_t} = |f_{S\cup Z_t}^{-1}(1)|$. This defines the quantum oracle

$$\mathcal{O}_{S\cup Z_t}|x\rangle \mapsto \begin{cases} -|x\rangle & \text{if } (f_s(x) = 1) \vee (x \in Z_t) \\ |x\rangle & \text{otherwise.} \end{cases} \tag{18}$$

Using amplitude amplification (see Theorem 1) with the quantum algorithm $\mathcal{A} = H^{\otimes n}$, the quantum oracle $\mathcal{O}_{S \cup Z_t}$ and setting $k_1 = \left\lfloor \frac{\pi}{4 \arcsin \sqrt{\frac{M'_S + 2^t}{2^n}}} - \frac{1}{2} \right\rfloor$ results in a quantum algorithm $\mathcal{B}(k_1)$ with success probability b, where

$$b = \sin^2 \left(2 \left\lfloor \frac{\pi}{4 \arcsin \sqrt{\frac{M'_S + 2^t}{2^n}}} - \frac{1}{2} \right\rfloor + 1 \right) \cdot \arcsin \sqrt{\frac{M_{S \cup Z_t}}{2^n}} \Big) \tag{19}$$

$$> \sin^2 \left(\left(\frac{\pi}{2 \arcsin \sqrt{\frac{M'_S + 2^t}{2^n}}} - 2 \right) \cdot \sqrt{\frac{2^t}{2^n}} \right) = b'. \tag{20}$$

We therefore have the that success probability of $\mathcal{B}(k_1)$ relative to $f_{S \cup Z_t}$ is $b > b'$. The success probability of $\mathcal{B}(k_1)$ relative to f_* is therefore $\frac{b}{M_{S \cup Z_t}}$ and a lower bound on this success probability is $\frac{b'}{M'_S + 2^t}$. It is plain that $\frac{b'}{M'_S + 2^t} < \frac{b}{M_{S \cup Z_t}} \leq \frac{1}{2^t}$.

We then use amplitude amplification again with the quantum algorithm $\mathcal{B}(k_1)$, the quantum oracle $\mathcal{O}_{f_*}$ and set $k_2 = \left\lfloor \frac{\pi}{4 \arcsin \sqrt{\frac{1}{2^t}}} - \frac{1}{2} \right\rfloor$ to obtain a quantum algorithm $\mathcal{C}(k_2)$ with a success probability relative to f_* of c, where

$$c = \sin^2 \left(\left(2 \left\lfloor \frac{\pi}{4 \arcsin \sqrt{\frac{1}{2^t}}} - \frac{1}{2} \right\rfloor + 1 \right) \cdot \arcsin \sqrt{\frac{b}{M_{S \cup Z_t}}} \right) \tag{21}$$

$$> \sin^2 \left(\left(\frac{\pi}{2 \arcsin \sqrt{\frac{1}{2^t}}} - 2 \right) \cdot \sqrt{\frac{b'}{M'_S + 2^t}} \right) = c'. \tag{22}$$

We therefore have a good lower bound on the probability of success for our algorithm assuming that $M_S \leq M'_S$, the probability of which occuring is $1 - \epsilon$. This allows us to compute a firm computational lower bound on the success rate of our solution to the *STO* problem. For intuitive purposes, if we assume the condition that $M_S < M'_S \ll 2^t \ll 2^n$ and use the approximation $\arcsin(x) \approx x$, then it is easily seen that $b' \approx 1$ and $c' \approx 1$. We now examine the costs involved.

Implementing $\mathcal{O}_{f_{S \cup Z_t}}$. We assume that we possess a circuit to compute the quantum evaluation $\mathcal{E}_{f_S}$. The identity $A \vee B \equiv A \oplus B \oplus (A \wedge B)$ implies we can implement $f_{S \cup Z_t} : \{0,1\}^n \longrightarrow \{0,1\}$ via computing

$$f_{S \cup Z_t}(x) = f_S(x) \oplus (x \in Z_t) \oplus (f_S(x) \wedge (x \in Z_t)), \tag{23}$$

and $\mathcal{O}_{f_{S \cup Z_t}}$ can be implemented via one $\wedge_{n-t+1}(X)$, one $\wedge_{n-t}(X)$ and one $\wedge_1(X)$ gate for a cost of $O(n-t)$ quantum gates, a cost usually dominated by $E_{\mathcal{E}_{f_S}}$.

Derivation of the cost $E_{\mathcal{B}}$ and therefore $E_{\mathcal{C}}$ can be derived by considering the cost equation (8) from Theorem 1 applied to the procedure described above. □

5 Applications to low qubit quantum oracles

We highlight two applications to cryptanalysis that have previously been studied in literature — single-target cryptanalysis of symmetric-key encryption systems and the Multivariate Quadratic ($\mathcal{MQ}$) problem over $\mathbb{F}_2$. Both problems share an essential structure — there exists a basic decomposition of the boolean function $\chi : \{0,1\}^n \longrightarrow \{0,1\}$ defining the search problem so that

$$\chi(x) \mapsto \chi_1(x) \wedge \cdots \wedge \chi_k(x) \qquad \text{where} \qquad \chi_i\{0,1\}^n \longrightarrow \{0,1\}. \quad (24)$$

In this way, $\chi(x) = 1$ if and only if $\chi_1(x) = \cdots = \chi_k(x) = 1$. In relation to the *STO* problem we can construct the function $f_S(x) \mapsto \chi_{i_1}(x) \wedge \cdots \wedge \chi_{i_r}(x)$ for some $r < k$ and $f_*(x) \mapsto \chi_1(x) \wedge \cdots \wedge \chi_k(x)$. The optimal choice of r indices that define f_S will be problem-specific and depend both upon the individual cost of each E_{χ_i} and the expected probability distribution of the value $|\chi_{i_1}^{-1}(1) \cap \cdots \cap \chi_{i_r}^{-1}(1)|$.

5.1 Low qubit implementations

We focus upon optimising the quantum circuit-size of low-qubit implementations in this paper and assume that we only have access to $n + w + 2$ qubits, where n is the number of bits required to represent the elements of the search space, $w = \max\{w_1, \ldots, w_k\}$ (where w_i is the number of qubits that each $\mathcal{O}'_{\chi_i}$ requires for working memory), one qubit is to allow efficient realisation of the $\wedge_n(X)$ gate and one qubit is kept in the state $|-\rangle$ to enable conditional phase inversion.

Schwabe and Westerbaan suggested a low-qubit strategy for evaluating the $\mathcal{MQ}$ problem over $\mathbb{F}_2$ [23] — this is easily extended to the above decomposition (24) and details may be found their original paper. For reasons of space we simply note that computing $\mathcal{O}_\chi$ in this inherently serial manner uses $n+w+\lfloor \log_2 k \rfloor + 3$ qubits and has an execution cost (circuit-depth or circuit-size) of at most

$$\sum_{i=1}^{k} 4E_{\chi'_i} - 2\max\left\{E_{\chi'_i}\right\} + 2\lfloor \log_2(k) \rfloor (E_{\wedge_1(SWAP)} + E_{\wedge_2(X)}) + E_{\wedge_{\lfloor \log_2 k \rfloor + 1}(X)}, \quad (25)$$

where $E_{\chi'_i}$ is the cost of $\mathcal{O}'_{\chi_i}$. Assuming that the $E_{\chi'_i}$ terms dominate the cost, the counter-based approach is roughly a factor of 2 times more costly in terms of circuit-size than an approach where we are not limited by the number of qubits and simply evaluate each $\mathcal{O}'_{\chi_i}$ in parallel and use the a compute, output $|\chi(x)\rangle$ using a $\wedge_k(X)$ gate and run each $\mathcal{O}'_{\chi_i}$ in reverse to uncompute the garbage bits.

5.2 A note on probability distributions

In order to choose M'_S we can simply apply Markov's inequality. This gives us that $\Pr\left[M_s \geq \mathbb{E}[M_S]\epsilon^{-1}\right] \leq \epsilon$, letting us derive a choice for M'_S based upon our chosen ϵ and problem-specific $\mathbb{E}[M_S]$. As both the AES and $\mathcal{MQ}$ problems can be modelled as the preimage search problem for a pseudorandom function [12, 11], the expected number of solutions to these problems can be easily computed.

For the $\mathcal{MQ}$ problem involving m equations in n variables over $\mathbb{F}_2$, this function is $F : \{0,1\}^n \longrightarrow \{0,1\}^m$, where m is the number of equations we are evaluating. For AES-k, where $k \in \{128, 192, 256\}$, we have $F : \{0,1\}^k \longrightarrow \{0,1\}^{128 \cdot r}$, where r is chosen such that $k < 128r$ to ensure that we have a single-target and r corresponds to the number of quantum circuits of AES-k we have to evaluate.

When the domain is smaller than the co-domain we cannot guarantee that there will exist only one solution, hence AES-k is often overdefined in cryptanalysis to ensure that there is only a single element — the key that encrypts r known plaintexts into r known ciphertexts. In the case of the $\mathcal{MQ}$ problem evaluating only $l < m$ equations leads to the function $F : \{0,1\}^n \longrightarrow \{0,1\}^l$, which allows us to create an underdefined problem where there exist many solutions.

This allows us to optimise the low-qubit version of quantum search applied to the binary $\mathcal{MQ}$ problem, where we use the low-qubit implementation from [23] in conjunction with Theorem 4. We use the suggested parameters for the Gui cryptosystem as a benchmark [20, 18, 19], where λ means that it should require at least 2^λ quantum gates to find a preimage of two $\mathcal{MQ}$ systems. We were only able to break the case of $\lambda = 256$ by using our method with just the counter-based oracle, but a hybrid approach that uses an intermediate number of qubits and only the counter-based oracle for the more expensive oracle was able to break all parameters. Full details will be in the final paper.

λ	$n = m$	[23]	[23] (counter)	[21]	Our method	Our method (counter)	Our method (hybrid)
80	117	$2^{80.9}/237/1$	$2^{81.9}/127/1$	$2^{78.6}/230/1$	$2^{79.7}/237/0.9999$	$2^{80.8}/127/0.9999$	$2^{79.9}/153/0.9999$
128	209	$2^{129.4}/421/1$	$2^{130.4}/220/1$	$2^{126.3}/415/1$	$2^{127.5}/421/0.9999$	$2^{128.5}/220/0.9999$	$2^{127.6}/246/0.9999$
256	457	$2^{256.7}/915/1$	$2^{257.7}/468/1$	$2^{252.9}/905/1$	$2^{253.8}/915/0.9999$	$2^{254.8}/468/0.9999$	$2^{253.9}/497/0.9999$

Table 1: Comparison quantum circuit-size/qubits/probability of success for various approaches to quantum search applied to cryptanalysis of Gui [20, 18].

We apply the same methodology to cryptanalysis of the Advanced Encryption Standard (AES) using a previously studied implementation of a quantum oracle for this problem [12]. We note that for the single-target case, we only require $r = 2$ plaintext-ciphertext pairs to ensure that we uniquely identify a single user's key for AES-128 and AES-192 and $r = 3$ plaintext-ciphertext pairs to uniquely identify a user's key for AES-256. We give our results in Table 2 below. In this case the more expensive oracle evaluates r plaintexts with a key, whilst the cheaper oracle evaluates only 1 plaintext with that key.

AES-k	[12] ($r = 2/3$)	Our method ($r = 10$)	Our method (counter) ($r = 10$)
128	$2^{86.87}/1969/1$	$2^{86.53}/1969/1$	$2^{86.53}/988/1$
192	$2^{119.23}/2225/1$	$2^{118.89}/2225/1$	$2^{118.89}/1115/1$
256	$2^{151.96}/4009/1$	$2^{151.03}/4009/1$	$2^{151.03}/1340/1$

Table 2: Comparison of quantum resource estimates for Grover vs the modified *STO* algorithm applied to cryptanalysis of single-target AES

As we can see from Table 2, even choosing an extremely large value of r to uniquely specify the user's key, we have that this method is essentially the same cost as if we used the bare minimum ($r = 2/3$) to ensure that the keys are uniquely identified. Full details and parameters will be in the final paper.

Conclusions We have extended the results of Kimmel et al. [14] and introduced a modified solution to the STO problem that solves an open problem. We additionally have found that it offers favourable quantum resource estimates for well-studied quantum resource estimates for cryptographic problems.

Caution must evidently be applied in choosing parameters for quantum-resistant cryptosystems in relation to quantum resource estimates, particularly if these estimates are based upon current best-known attacks using Grover's algorithm, as it is simply a special-case of amplitude amplification. We note that these optimisations do not impact upon query-complexity, hence basing parameters upon query-complexity is still a safe choice with respect to the optimisations presented in this paper. We have additionally demonstrated that low-qubit implementations in particular can exploit the STO method to lower the circuit-size for these problems and that sometimes there is little advantage with regards to quantum circuit-size in using large numbers of qubits other than to improve the performance of the implementation of the cheaper quantum oracle $\mathcal{O}_S$.

Explicit details of all proofs and theorems will be in the final paper and the code used in computations will be made available.

Acknowledgements Benjamin Pring was funded during the development of this research by the EPRSC grant EP/M50645X/1.

References

[1] Ahlgren, J.: The probability distribution for draws until first success without replacement. arXiv preprint arXiv:1404.1161 (2014)
[2] Amy, M., Maslov, D., Mosca, M., Roetteler, M.: A meet-in-the-middle algorithm for fast synthesis of depth-optimal quantum circuits. IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems 32(6), 818–830 (2013)
[3] Arunachalam, S., de Wolf, R.: Optimizing the number of gates in quantum search. Quantum Information & Computation 17(3&4), 251–261 (2017)
[4] Bouillaguet, C., Chen, H.C., Cheng, C.M., Chou, T., Niederhagen, R., Shamir, A., Yang, B.Y.: Fast exhaustive search for polynomial systems in $\mathbb{F}_2$. In: Int. Workshop on Cryptographic Hardware and Embedded Systems. pp. 203–218. Springer (2010)
[5] Bouillaguet, C., Cheng, C.M., Chou, T., Niederhagen, R., Yang, B.Y.: Fast exhaustive search for quadratic systems in $\mathbb{F}_2$ on FPGAs. In: International Conference on Selected Areas in Cryptography. pp. 205–222. Springer (2013)
[6] Boyer, M., Brassard, G., Høyer, P., Tapp, A.: Tight bounds on quantum searching. arXiv quant-ph/9605034 (1996)
[7] Brassard, G., Hoyer, P., Mosca, M., Tapp, A.: Quantum amplitude amplification and estimation. Contemporary Mathematics 305, 53–74 (2002)

[8] Casanova, A., Faugère, J.C., Macario-Rat, G., Patarin, J. Perret, L., Ryckeghem, J.: GeMSS–submission to the NIST post-quantum cryptography project. (2017), https://www-polsys.lip6.fr/Links/NIST/GeMSS_specification.pdf
[9] Chen, M.S., Hülsing, A., Rijneveld, J., Samardjiska, S., Schwabe, P.: MQDSS—submission to the NIST post-quantum cryptography project. (2017), http://mqdss.org/files/mqdss.pdf
[10] Ding, J., Chen, M.S., Petzoldt, A., Schmidt, D., Yang, B.Y.: Gui—submission to the nist post-quantum cryptography project. specification (2017), https://csrc.nist.gov/Projects/Post-Quantum-Cryptography/Round-1-Submissions
[11] Fusco, G., Bach, E.: Phase transition of multivariate polynomial systems. Mathematical Structures in Computer Science 19(1), 9–23 (2009)
[12] Grassl, M., Langenberg, B., Roetteler, M., Steinwandt, R.: Applying Grover's algorithm to AES: quantum resource estimates. In: International Workshop on Post-Quantum Cryptography. pp. 29–43. Springer (2016)
[13] Grover, L.K.: A fast quantum mechanical algorithm for database search. In: Proc. of the 28^{th} annual ACM symp. on Theory of computing. pp. 212–219. ACM (1996)
[14] Kimmel, S., Yen-Yu Lin, C., Han-Hsuan, L.: Oracles with costs. 10th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2015, May 20-22, 2015, Brussels, Belgium 44 (2015)
[15] Korepin, V.E., Grover, L.K.: Simple algorithm for partial quantum search. Quantum Information Processing 5(1), 5–10 (2006)
[16] Maslov, D.: Advantages of using relative-phase Toffoli gates with an application to multiple control Toffoli optimization. Physical Review A 93(2), 022311 (2016)
[17] Nielsen, M.A., Chuang, I.L.: Quantum Computation and Quantum Information. Cambridge University Press (2010)
[18] Petzoldt, A., Chen, M.S., Ding, J., Yang, B.Y.: HMFEV- an efficient multivariate signature scheme. In: International workshop on post-quantum cryptography. pp. 205–223. Springer (2017)
[19] Petzoldt, A., Chen, M.S., Ding, J., Yang, B.Y.: HMFEV- an efficient multivariate signature scheme (slides). International workshop on post-quantum cryptography (2017), https://2017.pqcrypto.org/conference/slides/mqI/HMFEv.pdf
[20] Petzoldt, A., Chen, M.S., Yang, B.Y., Tao, C., Ding, J.: Design principles for HFEv-based multivariate signature schemes. In: Int. Conference on the Theory and Application of Cryptology and Information Security. pp. 311–334. Springer (2015)
[21] Pring, B.: Exploiting preprocessing for quantum search to break parameters for mq cryptosystems. In: Arithmetic of Finite Fields-7th International Workshop, WAIFI 2018, Revised Selected Papers. WAIFI (2018)
[22] Pub, N.F.: 197: Advanced encryption standard (aes). Federal information processing standards publication 197(441), 0311 (2001)
[23] Schwabe, P., Westerbaan, B.: Solving binary $\mathcal{MQ}$ with Grover's algorithm. In: SPACE 2016. pp. 303–322. Springer (2016)
[24] Selinger, P.: Quantum circuits of T-depth one. Phys. Rev. A 87(4), 042302 (2013)
[25] of Standards, N.I., Technology.: Nist project for post-quantum cryptography standardization. https://csrc.nist.gov/Projects/Post-Quantum-Cryptography (2016), accessed: 07/10/2018
[26] of Standards, N.I., Technology.: Submission requirements and evaluation criteria for the post-quantum cryptography standardization process. (2016)