

An Optimized Quantum Implementation of ISD on Scalable Quantum Resources

Andre Esser¹, Sergi Ramos-Calderer^{1,2}, Emanuele Bellini¹, José I. Latorre^{1,2,3},
and Marc Manzano^{4*}

¹ Technology Innovation Institute, United Arab Emirates

{andre.esser, sergi.ramos, emanuele.bellini, jose.ignacio.latorre}@tii.ae

² Departament de Física Quàntica i Astrofísica and Institut de Ciències del Cosmos,
Universitat de Barcelona, Spain

³ Centre for Quantum Technologies, National University of Singapore, Singapore

⁴ Sandbox@Alphabet, Mountain View, CA, United States
marc.manzano@google.com

Abstract. The security of code based constructions is usually assessed by Information Set Decoding (ISD) algorithms. In the quantum setting, amplitude amplification yields an asymptotic square root gain over the classical analogue. However, it is still unclear whether a real quantum circuit could yield actual improvements or suffer an enormous overhead due to its implementation. This leads to different considerations of these quantum attacks in the security analysis of code based proposals. In this work we clarify this doubt by giving the first quantum circuit design of the fully-fledged ISD procedure, an implementation in the quantum simulation library Qibo as well as precise estimates of its complexities. We show that against common belief, Prange’s ISD algorithm can be implemented rather efficiently on a quantum computer, namely with only a logarithmic overhead in circuit depth compared to a classical implementation.

As another major contribution, we leverage the idea of classical co-processors to design hybrid classical-quantum trade-offs, that allow to tailor the necessary qubits to any available amount, while still providing quantum speedups. Interestingly, when constraining the width of the circuit instead of its depth we are able to overcome previous optimality results on constraint quantum search.

Keywords: ISD, decoding, quantum circuit, classical-quantum trade-offs

1 Introduction

The growing threat to modern widespread cryptography posed by the advancing development of quantum computers has led to a focus on other hardness assumptions. One of the leading and most promising proposals for post quantum

* This work was conducted while the author was affiliated with Technology Innovation Institute.

cryptography is code based cryptography. It has a long history of withstanding classical as well as quantum attacks and is considered to rely on one of the most well understood hardness assumptions. The list of the four KEM finalists of the ongoing NIST standardization process for post quantum cryptography [1] includes one code based proposal (McEliece [14]) and two more can be found on the alternate candidate list (BIKE [2] and HQC [30]) .

At the heart of all these code based constructions lies the binary decoding or *syndrome decoding* problem. This problem asks to find a low Hamming weight solution $\mathbf{e} \in \mathbb{F}_2^n$ to the equation $H\mathbf{e} = \mathbf{s}$, where $H \in \mathbb{F}_2^{(n-k) \times n}$ is a random binary matrix and $\mathbf{s} \in \mathbb{F}_2^{n-k}$ a binary vector.

The best known strategy to solve this problem is based on Information Set Decoding (ISD) [32], a technique introduced by Prange in 1962. Since then, there has been a series of works improving on his original algorithm [5, 11, 15, 28, 29, 35], mostly by leveraging additional memory, exploiting some meet-in-the-middle strategies.

In the quantum setting Bernstein showed how to speed up Prange’s algorithm by an amplitude amplification routine [6], which results in an asymptotic square root gain over the classical running time. The translation of advanced ISD algorithm to the quantum setting [25, 26] yields so far only small asymptotic improvements. Further these algorithms rely on the existence of an exponential amount of quantum RAM, which is considered very unrealistic even for mid term quantum developments. Due to this fact, all code based NIST submissions exclude these algorithms when conducting their security analysis. Moreover, the McEliece submission states that ”Known quantum attacks multiply the security level of both ISD and AES by an asymptotic factor $0.5 + o(1)$, but a closer look shows that the application of Grover’s method to ISD suffers much more overhead in the inner loop” [14].

So far it was unclear if such a statement is well-founded and how much overhead a quantum implementation of the procedure by Prange would really cause. In this work, we carefully design every part of Prange’s algorithm as a quantum circuit, analyze its complexities and show how to incorporate the pieces in a fully-fledged quantum ISD procedure.

In our design we put a special focus on the necessary amount of qubits. Note that several prior works also focus on qubit reduction in the *few qubits* or *polynomial memory* setting [3, 7, 9, 18, 23], in which the quantum algorithm is limited to the use of a polynomial amount of qubits only. Prange’s algorithm falls into this regime by default, since asymptotically it *only* uses a polynomial amount of memory. Nevertheless, it is especially this need for memory which limits its applicability, as all code based constructions involve parity-check matrices consisting of millions of bits. Hence, we investigate different optimizations of our initial design with regards to the amount of qubits. Furthermore, we extend the *few qubits* setting by developing hybrid algorithms that enable us to reduce the already polynomial demand of qubits to any available amount while still providing a quantum speedup.

In this context we leverage the idea of classical co-processors resulting in hybrid trade-off algorithms between classical-time and quantum memory (and time). The idea of such co-processors has mostly been used to parallelize quantum circuits or instantiate circuits under depth constraints, e.g. when analyzing the quantum security of schemes under the MAXDEPTH constraint specified by NIST [2, 9, 10, 24]. Under depth constraints, Zalka [38] showed that the optimal way to perform a quantum search is by partitioning the search space in small enough sets such that the resulting circuit only targeting one set at a time does not exceed the maximum depth. Then the search has to be applied for every set of the partition. However, this optimality result only holds under depth constraints, when instead imposing constraints on the width of the circuit, our trade-offs yield more efficient strategies.

Our Contribution. As a first contribution we design and analyze the full circuit performing the quantum version of Prange’s algorithm. We give precise estimates for the circuit depth and width in the quantum circuit model. Our design shows that, against common belief, Prange’s algorithm can be implemented rather efficiently on a quantum computer, namely with only a logarithmic overhead in the depth. Through further optimizations, our width optimized circuit only needs $(n-k)k$ bits to store and operate on the input matrix $H \in \mathbb{F}_2^{(n-k) \times n}$ and roughly $n-k$ ancillas. Additionally, we provide functional implementations of our circuits in the quantum simulation library Qibo [16,17], which is accessible on github [33]. We also explore different optimizations regarding the circuit depth, including a quantum version of the Lee-Brickell improvement [27] and an adaptation of our circuits to benefit from quasi-cyclic structures in the BIKE / HQC case.

Our second major contribution is the design of hybrid quantum-classical trade-offs that address the practical limitation on the amount of qubits. In particular, these trade-offs enable quantum speedups for any available amount of qubits. We study the behavior of our trade-offs for various different choices of code parameters. Besides the coding-theoretic motivated settings of full and half distance decoding, this includes also the parameter choices made by the NIST PQC candidates McEliece, BIKE and HQC. Our trade-offs perform best on the BIKE and HQC schemes, which is a result of a combination of a very low error weight and a comparably low code rate used by these schemes.

Fig. 1 shows the behavior of our trade-off achieving the best results under limited width. Here, we measure the performance of the trade-offs in form of a qubit-reduction factor δ and a speedup $t(\delta)$. In comparison to an entirely quantum based computation, performed using a specific amount of qubits and taking time $T^{\frac{1}{2}}$, the trade-off reduces the amount of qubits by a factor of δ , while maintaining a time complexity of $T^{t(\delta)}$. For instance in the BIKE and HQC setting we can reduce the amount of qubits to only 1% ($\delta = 0.01$) of an entire quantum based computation and still achieve a speedup of roughly 0.87 compared to a classical computation.

The rest of this work is structured as follows. In Section 2 we set up the necessary notation, give a precise definition of the problem under consideration

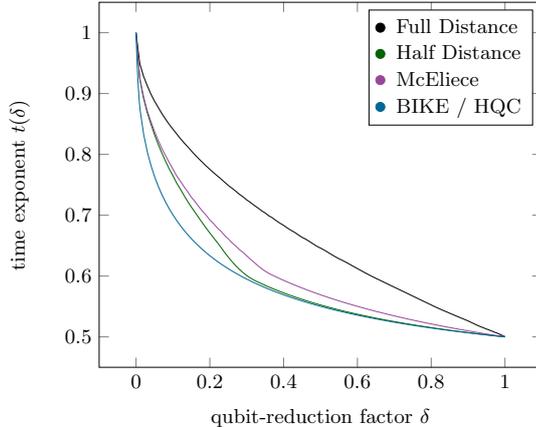


Fig. 1: Comparison of the achieved speedups of our trade-offs $t(\delta)$ (y-axis) plotted as a function of the qubit-reduction factor δ (x-axis).

and elaborate on the quantum model used for our analysis. In the subsequent Section 3 we present the original algorithm by Prange. In Section 4 we model every step of Prange’s algorithm as a quantum circuit and show how to make use of an amplitude amplification step. Finally, in Section 5 we give improvements for our initial design, including a quantization of the Lee-Brickell improvement as well as our hybrid classical-quantum trade-offs.

2 Preliminaries

For two integers $a, b \in \mathbb{N}$ with $a \leq b$ let $[a, b] := \{a, a+1, \dots, b\}$. Further we write conveniently $[b] := [1, b]$. Let H be an $m \times n$ matrix and $I \subseteq [n]$, we write H_I to denote the projection of H onto the columns indexed by I . We use the same notation for vectors, so for a vector \mathbf{v} of length n , then $\mathbf{v}_I := (v_{i_1}, v_{i_2}, \dots, v_{i_{|I|}})$, where $i_j \in I$. For a binary vector \mathbf{w} we define $\text{wt}(\mathbf{w}) := |\{i \in [n] \mid w_i = 1\}|$ as the Hamming weight of \mathbf{w} . For two reals $c, d \in \mathbb{R}$ we let $\llbracket c, d \rrbracket := \{x \in \mathbb{R} \mid c \leq x \leq d\}$ be the (including) interval of all reals between c and d .

We use standard Landau notation for complexity statements, where $\tilde{\mathcal{O}}$ -notation suppresses polylogarithmic factors, meaning $\tilde{\mathcal{O}}(f(x)) = \mathcal{O}(f(x) \log^i f(x))$ for any constant i . Besides standard binomial coefficients, for a set S , we let $\binom{S}{k}$ denote the set containing all size- k subsets of S . All logarithms are binary if not stated otherwise. We define $H(x) := -x \log(x) - (1-x) \log(1-x)$ to be the binary entropy function and make use of the well-known approximation

$$\binom{n}{k} = \tilde{\mathcal{O}}\left(2^{nH\left(\frac{k}{n}\right)}\right), \quad (1)$$

which can be derived from Stirling’s formula.

Decoding and linear codes. A binary linear $[n, k]^5$ or $[n, k, d]$ code \mathcal{C} is a k dimensional subspace of \mathbb{F}_2^n and minimum distance d , which is defined as the minimum Hamming weight of the elements of \mathcal{C} . We call n the code length and $R := \frac{k}{n}$ the code rate of \mathcal{C} . The code \mathcal{C} can be defined via the kernel of a matrix $H \in \mathbb{F}_2^{(n-k) \times n}$, so that $\mathcal{C} := \{\mathbf{c} \in \mathbb{F}_2^n \mid H\mathbf{c}^T = \mathbf{0}\}$, where H is called a *parity-check matrix*. Note that for ease of exposition, we treat all vectors as column vectors so that we can omit vector transpositions.

A given point $\mathbf{x} = \mathbf{c} + \mathbf{e}$ that differs from a codeword by an error \mathbf{e} can be uniquely decoded to \mathbf{c} as long as $\text{wt}(\mathbf{e}) \leq \lfloor \frac{d-1}{2} \rfloor$. This setting, in which the error weight is bounded by half of the minimum distance, is also known as *half distance* decoding. Another well-studied case upper bounds the error weight by the full minimum distance d and is hence known as *full distance* decoding. As the running time of decoding algorithms is solely increasing in the error weight, in those settings, we study the complexity for the case of equality to the respective upper bounds. Also in those cases we assume d to meet the Gilbert-Varshamov bound [19, 37] for random binary linear codes, which gives $d \approx H^{-1}(1 - R)n$.

Note that the definition of the code via its parity-check matrix allows to treat the decoding procedure independently of the specific codeword by considering the *syndrome* \mathbf{s} of a given faulty codeword \mathbf{x} , where $\mathbf{s} := H\mathbf{x} = H(\mathbf{c} + \mathbf{e}) = H\mathbf{e}$.

Now, if one is able to recover \mathbf{e} from H and \mathbf{s} , the codeword can be recovered from \mathbf{x} as $\mathbf{c} = \mathbf{x} + \mathbf{e}$. This leads to the definition of the *syndrome decoding problem*.

Definition 1 (Syndrome Decoding Problem). *Let \mathcal{C} be a linear $[n, k]$ code with parity-check matrix $H \in \mathbb{F}_2^{(n-k) \times n}$ and constant rate $\frac{k}{n}$. For $\mathbf{s} \in \mathbb{F}_2^{n-k}$ and $\omega \in [n]$, the syndrome decoding problem $\mathcal{SD}_{n,k,\omega}$ asks to find a vector $\mathbf{e} \in \mathbb{F}_2^n$ of weight $\text{wt}(\mathbf{e}) = \omega$ satisfying $H\mathbf{e} = \mathbf{s}$. We call any such \mathbf{e} a solution while we refer to (H, \mathbf{s}, ω) as an instance of the syndrome decoding problem.*

Quantum Circuits. Our algorithms are built in the quantum circuit model, where we assume a certain familiarity of the reader (for an introduction see [31]). The circuits are presented using general multi-qubit gates for simplicity, but we analyze their depth and complexity using their decomposition into basic implementable gates. Particularly, the decomposition of multi-controlled NOT gates into Toffoli gates is the main factor affecting depth and gate count.

A multi-controlled NOT gate with n controls can be decomposed using $\mathcal{O}(n)$ regular Toffoli gates [4]. If $n - 2$ ancillary qubits are available the procedure can be implemented in logarithmic depth. Note that we use the term circuit depth and time complexity interchangeably when analyzing our quantum circuits.

3 Prange’s Information Set Decoding

Let us introduce the original ISD algorithm by Prange [32]. Given an instance (H, \mathbf{s}, ω) of the $\mathcal{SD}_{n,k,\omega}$ Prange’s algorithm starts by choosing a random set

⁵ Note that we also use this notation to indicate the set of integers between n and k , but the concrete meaning will be clear from the context.

Algorithm 1 PRANGE

Require: parity-check matrix $H \in \mathbb{F}_2^{(n-k) \times n}$, syndrome $\mathbf{s} \in \mathbb{F}_2^{n-k}$, weight $\omega \in [n]$

Ensure: error vector \mathbf{e} with $\text{wt}(\mathbf{e}) = \omega$ satisfying $H\mathbf{e} = \mathbf{s}$

- 1: **repeat**
 - 2: choose random permutation matrix $P \in \mathbb{F}_2^{n \times n}$ and set $H_I \leftarrow (HP)_{[n-k]}$
 - 3: solve linear system $H_I \mathbf{e}_1 = \mathbf{s}$ for \mathbf{e}_1
 - 4: **until** $\text{wt}(\mathbf{e}_1) = \omega$
 - 5: **return** $P(\mathbf{e}_1, 0^k)$
-

$I \subseteq [n]$ of size $n - k$ and then solves the corresponding linear system

$$H_I \mathbf{e}_1 = \mathbf{s} \quad (2)$$

for \mathbf{e}_1 .⁶ Note that the solution \mathbf{e}_1 of the above linear system with weight $\omega' := \text{wt}(\mathbf{e}_1)$ can always be naively extended to a vector $\tilde{\mathbf{e}}$ of length n and weight ω' satisfying $H\tilde{\mathbf{e}} = \mathbf{s}$.

For the construction of $\tilde{\mathbf{e}}$ one simply sets the coordinates corresponding to the disregarded columns of H to zero. Hence, if $\omega' = \omega$ the vector $\tilde{\mathbf{e}}$ forms a solution to the syndrome decoding problem. The algorithm now chooses random subsets I until $\omega' = \omega$ holds.

Let us briefly analyze when Algorithm 1 succeeds in finding the solution. Assuming a unique solution \mathbf{e} , the algorithm is successful whenever \mathbf{e} projected to the coordinates given by I is a solution to the linear system in Eq. (2), hence if $\mathbf{e}_1 = \mathbf{e}_I$. This happens whenever \mathbf{e}_I covers the full weight of \mathbf{e} , in which case I or more precisely $[n] \setminus I$ is called an *information set*. Transferred to Algorithm 1 this applies whenever, for the permutation chosen in line 2, it holds that $P^{-1}\mathbf{e} = (\mathbf{e}_1, 0^k)$ for $\mathbf{e}_1 \in \mathbb{F}_2^{n-k}$. The probability that the permutation distributes the weight in such a way is

$$q := \Pr [P^{-1}\mathbf{e} = (\mathbf{e}_1, 0^k)] = \frac{\binom{n-k}{\omega}}{\binom{n}{\omega}}. \quad (3)$$

Hence, the expected number of tries until we draw a suitable permutation P becomes q^{-1} and the expected time complexity is $T = q^{-1} \cdot T_G$, where T_G describes the cost for solving the linear system and performing the weight check.

Remark 1. Note that in the case of S existent solutions the time complexity to retrieve a single solution with Prange's algorithm becomes $\frac{T}{S}$.

4 A first design of a quantum ISD circuit

In this section we give an initial design for the quantum version of Prange's ISD algorithm. Our design is composed of the following three main building blocks:

⁶ Note that in Algorithm 1 we model H_I as the first $n - k$ columns of HP , where P is a random permutation matrix.

- 1) The creation of the uniform superposition over all size- k subsets of $[n]$ (corresponding to the selection of information sets in line 2 of Algorithm 1).
- 2) The Gaussian elimination step to derive the error related to a given information set (line 3 of Algorithm 1).
- 3) A quantum search for an information set yielding an error of the desired weight (substituting the repeat loop in line 1 of Algorithm 1).

Next, we give independent descriptions of our circuit designs for the different steps after which we discuss how to incorporate them in a quantum search. We provide implementations of all described circuits in the quantum simulation library Qibo with the source code accompanying this work.

4.1 Superposition over size- k subsets

We represent a size- k subset $S \subset [n]$ via a binary vector \mathbf{b} of length n with exactly k bits set to one, where $i \in S$ iff $b_i = 1$. Let $\mathcal{B}_{n,k}$ denote the set of all such binary vectors. Our circuit builds the uniform superposition over $\mathcal{B}_{n,k}$ in a bit by bit fashion. Grover and Rudolph in [20] follow a similar approach of dividing the probability space into smaller parts to provide a feasibility argument for creating quantum states following a broad class of probability distributions. However, they leave it open how to create these quantum states for concrete distributions.

Our design relies on the observation that the fraction of vectors from $\mathcal{B}_{n,k}$ starting with a zero or a one respectively is known a priori and independent of subsequent bits. Therefore note that $\mathcal{B}_{n,k}$ splits into

$$|\mathcal{B}_{n,k}| = \binom{n}{k} = \underbrace{|\mathcal{B}_{n-1,k}|}_{\text{elements starting with zero}} + \underbrace{|\mathcal{B}_{n-1,k-1}|}_{\text{elements starting with one}}. \quad (4)$$

Hence, we start by rotating the first qubit such that we measure a zero with probability $a := \frac{|\mathcal{B}_{n-1,k}|}{|\mathcal{B}_{n,k}|} = \frac{\binom{n-1}{k}}{\binom{n}{k}}$ and respectively a one with probability $b := \frac{|\mathcal{B}_{n-1,k-1}|}{|\mathcal{B}_{n,k}|} = \frac{\binom{n-1}{k-1}}{\binom{n}{k}}$, i.e.,

$$|0\rangle \mapsto \sqrt{a}|0\rangle + \sqrt{b}|1\rangle.$$

For the second bit we proceed similar. In the case of the first qubit being zero the remaining combinations split analogously to Eq. (4) in $|\mathcal{B}_{n-1,k}| = |\mathcal{B}_{n-2,k}| + |\mathcal{B}_{n-2,k-1}|$, otherwise we find $|\mathcal{B}_{n-1,k-1}| = |\mathcal{B}_{n-2,k-1}| + |\mathcal{B}_{n-2,k-2}|$. This allows us to rotate the second bit accordingly (this time depending on the first qubit) such that

$$\begin{aligned} \sqrt{a}|00\rangle + \sqrt{b}|10\rangle &\mapsto \sqrt{a} \left(\sqrt{\frac{|\mathcal{B}_{n-2,k}|}{|\mathcal{B}_{n-1,k}|}} |00\rangle + \sqrt{\frac{|\mathcal{B}_{n-2,k-1}|}{|\mathcal{B}_{n-1,k}|}} |01\rangle \right) \\ &+ \sqrt{b} \left(\sqrt{\frac{|\mathcal{B}_{n-2,k-1}|}{|\mathcal{B}_{n-1,k-1}|}} |10\rangle + \sqrt{\frac{|\mathcal{B}_{n-2,k-2}|}{|\mathcal{B}_{n-1,k-1}|}} |11\rangle \right). \end{aligned}$$

Algorithm 2 Generate Uniform Superposition over $\mathcal{B}_{n,k}$

Require: $n, k \in \mathbb{N}$ with $k \leq n$, n qubits b_i and $\lceil \log(k+1) \rceil$ ancilla qubits (to store c)**Ensure:** Uniform superposition over $\mathcal{B}_{n,k}$, represented by the b_i 's

```

 $c \leftarrow k$ 
for  $i = 0$  to  $n - 1$  do
  for  $j = k$  down to  $1$  do
    if  $c = j$  then
      Rotate  $b_i$  such that  $1$  is measured with probability  $\frac{|\mathcal{B}_{n-i-1,j-1}|}{|\mathcal{B}_{n-i,j}|}$ 
    if  $b_i = 1$  then
       $c \leftarrow c - 1$ 
return  $\mathbf{b}$ 

```

We now proceed analogously for the remaining bits, where each bit depends on the state of its successors. A crucial observation is that for a general position i , the fraction of elements having the i -th bit equal to zero or one does not depend on the *exact* pattern of the previous $i - 1$ bits, but only on their weight. Therefore, consider a general state where the first $i - 1$ qubits have already been processed and their weight is equal to j . Then all combinations for the remaining $n - (i - 1)$ bits are given by $\mathcal{B}_{n-(i-1),k-j}$. Again the number of elements starting with a zero (or one respectively), can be derived, analogously to Eq. (4), as

$$|\mathcal{B}_{n-(i-1),k-j}| = \underbrace{|\mathcal{B}_{n-i,k-j}|}_{\text{elements starting with zero}} + \underbrace{|\mathcal{B}_{n-i,k-j-1}|}_{\text{elements starting with one}}.$$

Now, by keeping track of the weight via some auxilliary qubits we can perform the needed rotations for every bit controlled on these ancillas. Algorithm 2 gives a description in pseudocode on how to construct the circuit.

As in general there are k possibilities for the weight of the succeeding bits and we need to process a total of n bits, our circuit design achieves a depth of $\mathcal{O}(n \cdot k)$. To keep track of the weight c of processed bits we use $\lceil \log(k+1) \rceil$ ancillas to store the binary representation of c . The binary additions enlarge the circuit depth by a factor of $\mathcal{O}(\log k)$. The execution of the rotation gates controlled by the ancillary register c involves multi-controlled gates with $\log k$ controls, whose decomposition contributes with an additional $\mathcal{O}(\log \log k)$ factor, if we allow for further $\log k$ ancillary qubits. In total this yields a circuit depth of $D_{\text{sup}} = \mathcal{O}(nk \log k \log \log k)$ using $n + 2\lceil \log(k+1) \rceil$ qubits to create the superposition over $\mathcal{B}_{n,k}$.

In Fig. 2 we give an example of our circuit for the case of $n = 5$ and $k = 2$.

4.2 Quantum Gaussian Elimination

Classical Gaussian elimination splits in the transformation to reduced row echelon form and finally the matrix diagonalization or solving step via back substitution. The first step performs a pivot search to enforce a non-zero diagonal

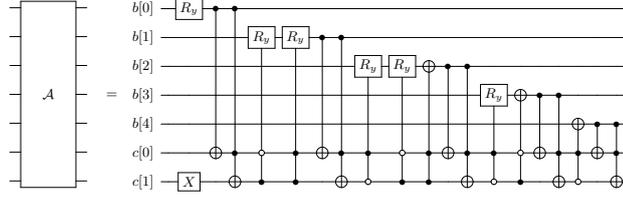


Fig. 2: Quantum circuit generating uniform superposition over $\mathcal{B}_{5,2}$.

Algorithm 3 Solve Linear System

Require: matrix $A \in \mathbb{F}_2^{n \times n}$, target vector $\mathbf{t} \in \mathbb{F}_2^n$ and $n(n+1)$ qubits

Ensure: $\mathbf{x} \in \mathbb{F}_2^n$ with $A\mathbf{x} = \mathbf{t}$

- 1: Initialize qubits with augmented matrix $A \leftarrow (A \mid \mathbf{t})$
 - 2: **for** $i = 1$ **to** $n - 1$ **do**
 - 3: **for** $j = i + 1$ **to** n **do** ▷ pivot search
 - 4: **if** $\forall_{i \leq k < j} a_{ki} = 0$ **then**
 - 5: add row j to row i starting from column $i + 1$
 - 6: **for** $j = i + 1$ **to** n **do** ▷ row reduce
 - 7: **if** $a_{ji} = 1$ **then**
 - 8: add row i to row j starting from column $i + 1$
 - 9: **for** $i = n$ **down to** 2 **do** ▷ back substitution
 - 10: **for** $j = i - 1$ **down to** 1 **do**
 - 11: **if** $a_{ji} = 1$ **then**
 - 12: add row i to row j only on the last column: $a_{j(n+1)} \leftarrow a_{j(n+1)} + a_{i(n+1)}$
 - 13: **return** $\mathbf{x} = (a_{1(n+1)}, \dots, a_{n(n+1)})$
-

entry and eliminates entries below the diagonal, while the latter eliminates entries above the diagonal. By performing all operations on the augmented matrix, containing the target vector as a last column, the solution to the system can be obtained from the last column of the final matrix.

Our Gaussian elimination circuit mostly resembles this classical procedure but modeled only with quantum gates. To save on depth and width we perform row additions during the transformation to reduced row echelon form only on columns succeeding the current column. Similarly, during the back substitution step row additions are only performed on the actual solution register. The pseudocode to generate our quantum circuit is given in Algorithm 3.

Our circuit needs no additional qubits besides the description of the linear system and achieves a depth analogous to the classical counterpart of $\mathcal{O}(n^3)$ but involving multi-controlled gates. Our decomposition strategy for these multi-controls introduced by the pivot search (compare to line 4 in Algorithm 3) costs an additional factor of $\mathcal{O}(\log n)$ in depth and an additional amount of $n - 2$ qubits, resulting in a total depth of $D_{\text{gauss}} = \mathcal{O}(n^3 \log n)$ and a width of $W_{\text{gauss}} = n^2 + n - 2$.

Algorithm 4 Combined Circuit

Require: matrix $H \in \mathbb{F}_2^{(n-k) \times n}$, syndrome $\mathbf{s} \in \mathbb{F}_2^{n-k}$, $n + (n+1)(n-k) + \lceil \log(n-k) \rceil$ qubits

Ensure: Uniform superposition over weight of all \mathbf{e}_1 with $H_I \mathbf{e}_1 = \mathbf{s}$ for $I \in \binom{[n]}{n-k}$

- 1: Initialize qubits with $(H \mid \mathbf{s})$
 - 2: Generate uniform superposition over $\mathcal{B}_{n,n-k}$ on qubits (b_1, \dots, b_n)
 - 3: **for** $i = 1$ **to** n **do** ▷ swap H_I to the front
 - 4: **if** $b_i = 1$ **then**
 - 5: **for** $j = i - 1$ **down to** 1 **do**
 - 6: swap column j and $j + 1$
 - 7: Apply Gaussian elimination circuit (Algorithm 3)
 - 8: **return** $c \leftarrow$ weight of last column
-

4.3 Designing a combined circuit

Next, we show how to combine both previously introduced building blocks to create a circuit which generates the uniform superposition over the solutions to the linear systems $H_I \mathbf{e}_1 = \mathbf{s}$, for $I \in \binom{[n]}{k}$.

Our combined circuit works in-place by first swapping the columns corresponding to the selected subset to the front of the matrix. Then we apply the Gaussian Elimination circuit to the first $n - k$ columns.

Circuit Description. First we generate the uniform superposition over $\mathcal{B}_{n,n-k}$, which determines the current selection of columns belonging to H_I . Next we swap all columns belonging to H_I to the front of the matrix (controlled on the chosen subset). Now, we implement the previously described Gaussian elimination circuit on the first $n - k$ columns of H . Since the final goal is to find a low weight solution it follows an accumulation of the weight of the solution in a separate register, later used by the amplitude amplification procedure. In Algorithm 4 we give a pseudocode description for the combined circuit generation. Fig. 3 illustrates the different operators needed in order to execute the combined circuit.

Circuit Complexity. In total the amount of qubits necessary for our initial design can be summarized as

$$\begin{aligned}
 W_{\text{combined}} &= \underbrace{n}_{\text{Permutation}} + \underbrace{(n-k) \cdot (n+1)}_{\text{Matrix}} + \underbrace{n-k-2}_{\text{Auxiliary}} \\
 &= (n-k+1)(n+2) - 4.
 \end{aligned}$$

Note that all subroutines use the same auxiliary qubits as we ensure they return to the zero state after each procedure.

The depth of our circuit is dominated by the circuit that swaps the columns belonging to H_I to the beginning of H as well as the Gaussian elimination and

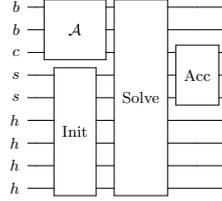


Fig. 3: Different parts detailed in Algorithm 4 as quantum operators. \mathcal{A} represents the creation of superposition, **Init** initializes the registers with $(H \mid s)$, **Solve** solves the linear system, **Acc** accumulates the weight of the solution.

summarizes as

$$\begin{aligned}
 D_{\text{combined}} = \mathcal{O} & \left(\underbrace{n(n-k)^2 \log(n-k)}_{\text{Gaussian}} + \underbrace{n(n-k) \log(n-k) \log \log(n-k)}_{\text{Permutation}} \right. \\
 & \left. + \underbrace{n^2(n-k)}_{\text{Swaps}} \right) = \mathcal{O}(n^3 \log n) \quad (5)
 \end{aligned}$$

4.4 Amplitude Amplification

Amplitude amplification was introduced as a generalization of Grover's algorithm [21] in [12, 22] and analogously allows to obtain a square-root advantage over a classical search. More precisely, given a quantum operation \mathcal{A} , that creates a quantum state with non-zero overlap with a target state of amplitude a , one can amplify the probability of measuring the desired state to $\Theta(1)$ using $\mathcal{O}(1/a)$ iterations of operator \mathcal{A} , whereas classical sampling would require $\mathcal{O}(1/a^2)$. The amplitude amplification operator is defined as

$$\mathcal{Q} = -\mathcal{A}S_0\mathcal{A}^\dagger S_t, \quad (6)$$

where S_0 and S_t are operators that flip the sign of the initial state and target state respectively. The \mathcal{Q} operator, when applied to the quantum state $\mathcal{A}|0\rangle$ amplifies the probability of measuring the quantum state targeted by S_t . Applying the operator $\lceil \pi/(4a) \rceil$ times then results in measuring one of the target states with probability close to 1.

In our case \mathcal{A} is the operator that creates the uniform superposition over all size- $(n-k)$ subsets. The oracle S_t is comprised of the Gaussian elimination circuit and an ancilla initialized in the $|-\rangle$ state to flip the sign of the states with target Hamming weight. Eventually, the inverse of the Gaussian elimination circuit is applied in order to clean up the ancillary register. The shape of S_t is detailed in Fig. 4.

Summarizing, we can find a solution to the syndrome decoding problem after $\mathcal{O}(\sqrt{q^{-1}})$ applications of operator \mathcal{Q} , where q is the proportion of suitable subsets yielding a solution among all size- $(n-k)$ subsets defined in Eq. (3).

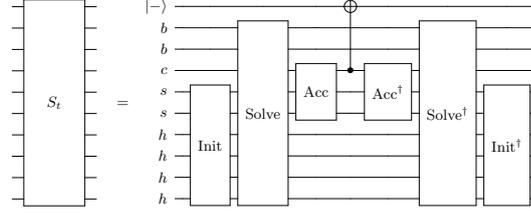


Fig. 4: Oracle designed to change the sign of all states with target Hamming weight after solving the linear system.

The expected time complexity, becomes $\mathcal{O}(\sqrt{q^{-1}}T_{\mathcal{Q}})$, where $T_{\mathcal{Q}}$ is the cost of performing the quantum operator \mathcal{Q} . Note that $T_{\mathcal{Q}} = \mathcal{O}(D_{\text{combined}})$ is dominated by the cost of the combined circuit given in Eq. (5). Hence, the depth of the full quantum ISD procedure can be summarized as

$$D_{\text{full}} = \mathcal{O}(\sqrt{q^{-1}}D_{\text{combined}}). \quad (7)$$

The amplitude amplification procedure only requires a single additional qubit for the sign flip, hence the number of qubits is equal to $W_{\text{full}} = W_{\text{combined}} + 1$.

Remark 2. Note that similar to Remark 1, in the case of S existent solutions the proportion of subsets yielding a solution increases to $S \cdot q$. Hence $\mathcal{O}(\sqrt{(Sq)^{-1}})$ applications of the operator Q suffice to find any one of these solutions.

5 Optimizing the quantum circuit

In this section we introduce further optimizations of our initial design. First we show how to extend our circuit by the ISD improvement made by Lee-Brickell [27] and how to exploit the cyclicity in case of BIKE and HQC. Then, we show how, in the case of Prange, a preprocessing of the matrix to *systematic form* and a clever adaptation of our combined circuit allows us to save $(n-k)^2$ input qubits without any increase in depth.

5.1 Quantum Lee-Brickell

Lee and Brickell [27] observed that allowing for a small weight p outside of the selected subset of columns can yield a polynomial runtime improvement. Therefore they aim for a permutation P , that distributes the weight on $P^{-1}\mathbf{e} = (\mathbf{e}_1, \mathbf{e}_2) \in \mathbb{F}_2^{n-k} \times \mathbb{F}_2^k$ such that

$$\text{wt}(\mathbf{e}_1) = \omega - p \text{ and } \text{wt}(\mathbf{e}_2) = p. \quad (8)$$

Again by Gaussian elimination (here modelled via the multiplication by the matrix Q) one transforms the identity $HP(\mathbf{e}_1, \mathbf{e}_2) = \mathbf{s}$ into

$$QHP(\mathbf{e}_1, \mathbf{e}_2) = (I_{n-k}H_2)(\mathbf{e}_1, \mathbf{e}_2) = Q\mathbf{s}.$$

The algorithm then enumerates in each iteration all possible candidates for \mathbf{e}_2 of weight p and checks for every such candidate \mathbf{x} if

$$\text{wt}(\mathbf{e}_1) = \text{wt}(Q\mathbf{s} + H_2\mathbf{x}) = \omega - p,$$

and if so outputs the solution $\mathbf{e} = P(\mathbf{e}_1, \mathbf{x})$. Note that the probability for the permutation satisfying Eq. (8) improves to

$$q_{\text{LB}} = \frac{\binom{n-k}{\omega-p} \binom{k}{p}}{\binom{n}{\omega}}. \quad (9)$$

On the downside in each iteration now the Gaussian elimination as well as the enumeration of all candidates must be performed which results in a total classical running time of

$$T = (q_{\text{LB}})^{-1} \cdot \left(T_{\text{G}} + \binom{k}{p} \right),$$

which is optimal for a p satisfying $T_{\text{G}} \approx \binom{k}{p}$, which implies constant p and hence a polynomial speedup.

Circuit Adaptation. Note that the Lee-Brickell improvement requires knowledge of H_2 , which corresponds to the last k columns of H after the Gaussian elimination. Hence, we need to extend all row additions of Algorithm 3 to be also applied to these columns.⁷ Note that this does not affect the running time in \mathcal{O} -notation.

Now, after the Gaussian elimination (compare to Algorithm 4) we add a circuit that for each selection of p columns of H_2 , adds those to the last column. After the addition we check if the weight of the last column is equal to $\omega - p$ and if so set the control bit of the amplitude amplification procedure to one. After that, we reverse the addition by adding those columns again. The total depth of this enumeration circuit is $\mathcal{O}(p \binom{k}{p})$.

The circuit depth of the Lee-Brickell quantum algorithm can be summarized similar to before in the case of Prange as

$$D_{\text{LB}} = \mathcal{O} \left(\sqrt{(q_{\text{LB}})^{-1}} \left(n^3 \log n + p \binom{k}{p} \right) \right).$$

In terms of width, the Lee-Brickell circuit has the same performance as our initial Prange design, namely

$$W_{\text{LB}} = W_{\text{full}} = (n - k + 1)(n + 2) - 3.$$

The source code accompanying this work also provides an implementation of this improvement in *Qibo* [33].

⁷ Precisely, the row additions of the back substitution step (line 13 of Algorithm 3) now need to be applied on the last $k+1$ columns rather than only on the last column.

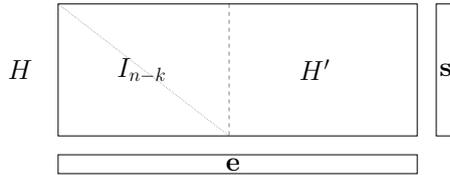


Fig. 5: Problem shape for input matrix in systematic form.

5.2 The case of BIKE and HQC – exploiting the cyclicity

BIKE and HQC use double-circulant codes of rate $\frac{1}{2}$, i.e., $n = 2k$. For those codes, a given syndrome decoding instance (H, \mathbf{s}, ω) allows to obtain k different instances $(H, \mathbf{s}_i, \omega)$ ⁸, where the solution to any one of these instances leads directly to a solution to the original instance. Sendrier has shown [34] that in the case enumeration based ISD, this setting allows for a classical-speedup of \sqrt{k} , known as *Decoding-One-Out-of-Many*. However, we observe that in the case of Prange these k instances also allow for a *quantum*-speedup of $\mathcal{O}(\sqrt{k})$ (corresponding to a classical speedup of order k). Therefore instead of performing the Gaussian elimination on the matrix $(H | \mathbf{s})$ it is performed on $(H | \mathbf{s}_1 | \dots | \mathbf{s}_k)$. Now, whenever one of the last k columns after the Gaussian elimination admits weight ω we set the sign-flip bit of the amplitude amplification procedure. Note that this change does not effect the depth of the circuit in \mathcal{O} -notation. It increases the cost for a row addition from $\mathcal{O}(n)$ to $\mathcal{O}(n + k) = \mathcal{O}(n)$ and the weight-accumulation has to be performed k times instead of once, which is still surpassed by the cost of the Gaussian elimination. On the upside there exist a total of k solutions to these k instances, which according to Remark 2 results in a speedup of \sqrt{k} of the quantum search. In terms of width we need $(n - k)(k - 1)$ more qubits to represent all \mathbf{s}_i .

This strategy is also compatible with the Lee-Brickell improvement from the previous Section 5.1. Therefore we need to extend the enumeration part to all \mathbf{s}_i increasing the cost of that step by a factor of k to $\mathcal{O}(kp(n - k)\binom{k}{p})$. We summarize the width and depth of this combination in Table 2 at the end of Section 6.

5.3 Saving quadratically many qubits for free

In the following we apply a Gaussian elimination to the first $n - k$ columns of H resulting in a problem shape as shown in Fig. 5, also known as *systematic form*.⁹ We now describe how to adapt our circuit to only require the matrix H' as well as the corresponding syndrome as an input, reducing the amount of qubits by $(n - k)^2$.

In a classical setting, preprocessing the matrix to systematic form allows to reduce the number of operations needed for subsequent Gaussian eliminations [8].

⁸ All defined on the same parity-check matrix H and with the same error-weight ω .

⁹ If the first $n - k$ columns do not form a matrix of full rank we permute the columns accordingly.

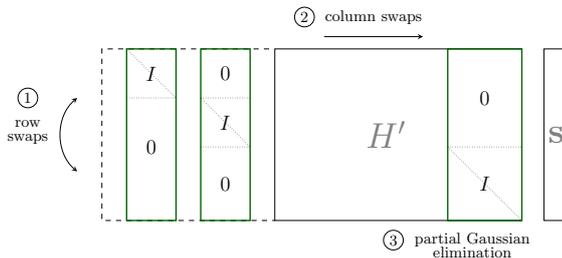


Fig. 6: Procedure to perform quantum version of Prange without first $n - k$ columns as input. Colored framed parts indicate columns belonging to the current selected subset.

This holds as long as there is a non-empty intersection between the new subset of columns and the first $n - k$ columns, which already form the identity matrix.

Inspired by the classical time complexity improvement, we implement for columns from the identity part that belong to the selected subset only a corresponding row-swap. After the necessary row-swaps are performed, all columns of H' belonging to the corresponding subset are swapped to the *back*. Subsequently we perform the Gaussian elimination only on the last columns of H' that belong to the current selection. This procedure is depicted in Fig. 6, which shows the state of the matrix after all three operations have been performed for the chosen subset. Note that the first $n - k$ columns only serve an illustrative purpose and are not part of the input.

Circuit Adaptation . In the following we assume the input matrix to be in *systematic form*. Now, whenever the selected subset includes any column \mathbf{h}_j for $j \leq n - k$ of H , a single row swap is sufficient to obtain the desired unit vector for this column. This implies that all operations resulting from column \mathbf{h}_j of H , where $j \leq n - k$, being part of the selected subset as column i are fully determined by j and i . The necessary operations can, hence, be embedded into the circuit directly without the need of the first $n - k$ columns as an input. However, this design requires to implement the procedure for every combination of j and i , which are $\mathcal{O}(n^2)$ possibilities. Furthermore, each swap comes at a cost of $\mathcal{O}(n)$.

Now, it follows the part of the circuit swapping all columns of H' belonging to the respective subset to the back. During these swaps we craft an ancillary state $x = |0^{k-r}1^r\rangle$, where r is the number of selected columns from H' . This state allows us to perform the Gaussian elimination only on the last r columns of H' (by controlling all operations that depend on column j of H' on bit x_j). Since we do not know a priori how many columns that might be, we start the Gaussian elimination, contrary to the description before, from the back, where it always starts with the last column.

In total, our modified circuit then needs an amount of qubits equal to

$$\begin{aligned}
 W_{\text{improved}} &= \underbrace{n}_{\text{Permutation}} + \underbrace{(n-k) \cdot (k+1)}_{\text{Matrix}} + \underbrace{k+n-k-2}_{\text{Auxiliary}} + \underbrace{1}_{\text{Sign-flip}} \\
 &= (n-k+2)(k+3) - 7,
 \end{aligned} \tag{10}$$

since we save the identity part, but need an additional k ancillas for representing x . The depth is still dominated by the Gaussian elimination on the last (possibly) $n-k$ columns and, hence, stays as in the initial design in Eq. (7). We again provide an implementation of this design with the source code belonging to this work. A pseudocode description of the adapted combined circuit can be found in Appendix I.

6 Classical-time quantum-memory trade-offs

Despite our optimization, the quadratic amount of qubits required for the representation of the input matrix is still the limiting factor with respect to real quantum implementations.

We overcome this issue by introducing hybrid trade-offs between classical-time and quantum-memory for our ISD circuit, allowing for an adaptive scaling of the algorithm to the available amount of qubits. Our trade-offs divide in a classical and quantum computation part, where a decrease of the amount of qubits comes at the cost of an increased classical running time. Since this increase in running time is of exponential nature we neglect the polynomial factors of the implementation by switching to \tilde{O} -notation. Our trade-offs allow for a smooth interpolation between purely classical computations at a running time of

$$T_C := \tilde{O} \left(\frac{\binom{n}{\omega}}{\binom{n-k}{\omega}} \right), \tag{11}$$

(compare to the analysis in Section 3) and a purely quantum based computation taking time $\sqrt{T_C}$, as given in Eq. (7). We interpolate between both complexities using a qubit reduction factor δ , where a fully classical computation corresponds to $\delta = 0$ and an entirely quantum based execution implies $\delta = 1$. For each trade-off we then state the running time for a given reduction factor δ as $t(\delta) \in \llbracket 0.5, 1 \rrbracket$, meaning that a reduction of the amount of qubits by a factor of δ implies a total running time of $(T_C)^{t(\delta)}$.

We start at first with a quite straightforward trade-off, which already achieves a better than linear dependence between δ and $t(\delta)$. This first trade-off also achieves good results for concrete medium sized parameters. After that, we present a second trade-off which asymptotically outperforms the first one. However, for concrete parameters in medium scale both trade-offs remain superior to each other for certain values of δ . Finally, we show how to combine both trade-offs to obtain an improved version. For large reduction factors, meaning close to one, this combination obtains the minimum of both previous trade-offs while for

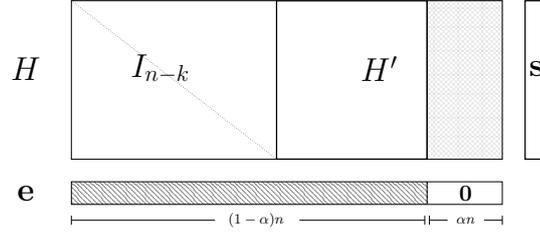


Fig. 7: Parity-check matrix where αn zero positions of \mathbf{e} are guessed. Striped region of \mathbf{e} indicates parts containing weight, crosshatched columns of H' do not affect \mathbf{s} . Framed parts are used as input to the quantum algorithm.

small reduction factors, which are most important when aiming at near future realizations, an improved running time is achieved in most settings.

We also provide an implementation of our classical co-processor in Sage that invoke the quantum circuit (implemented in the quantum simulation library Qibo) on the respective instances.

6.1 A hybrid version of Prange

Our first trade-off is a hybrid version of Prange’s original algorithm. In his original algorithm k zero positions of \mathbf{e} are guessed and then the linear system corresponding to the non-zero positions is solved in polynomial time. In our hybrid version the classical part consists in guessing $\alpha n \leq k$ zero coordinates of \mathbf{e} , which allows to shorten the code and, hence, reduce the problem to an $[(1 - \alpha)n, k - \alpha n]$ code, while the error weight remains unchanged (compare to Fig. 7). This reduced instance is then solved with our previously constructed quantum circuit. Should the quantum computation not result in an actual solution, the initial guess of zero coordinates was incorrect and we proceed with a new guess.

Algorithm 5 gives a pseudocode description of our hybrid Prange variant.

Algorithm 5 HYBRID-PRANGE

Require: parity-check matrix $H \in \mathbb{F}_2^{(n-k) \times n}$, syndrome $\mathbf{s} \in \mathbb{F}_2^{n-k}$, weight $\omega \in [n]$, qubit reduction factor $\delta \in \llbracket 0, 1 \rrbracket$

Ensure: error vector \mathbf{e} with $\text{wt}(\mathbf{e}) = \omega$ satisfying $H\mathbf{e} = \mathbf{s}$

- 1: $\alpha := (1 - \delta) \frac{k}{n}$
 - 2: **repeat**
 - 3: choose random permutation matrix $P \in \mathbb{F}_2^{n \times n}$ and set $\tilde{H} \leftarrow HP$
 - 4: solve instance $(\tilde{H}_{[(1-\alpha)n]}, \mathbf{s}, \omega)$ via quantum algorithm returning $\mathbf{e}_1 \in \mathbb{F}_2^{(1-\alpha)n}$
 - 5: $\mathbf{e} \leftarrow P(\mathbf{e}_1, 0^{\alpha n})$
 - 6: **until** $H\mathbf{e} = \mathbf{s}$
 - 7: **return** \mathbf{e}
-

Theorem 1 (Hybrid Prange). *Let $n \in \mathbb{N}$, $\omega = \tau n$ and $k = Rn$ for $\tau, R \in \llbracket 0, 1 \rrbracket$ and let T_C be as defined in Eq. (11). Then for any qubit reduction factor $\delta \in \llbracket 0, 1 \rrbracket$ Algorithm 5 solves the $\mathcal{SD}_{n,k,\omega}$ problem in time $(T_C)^{t(\delta)}$ using $\delta(1-R)Rn^2$ qubits for the matrix representation, where*

$$t(\delta) = 1 - \frac{\frac{1}{2} \left((1-\alpha)H\left(\frac{\tau}{1-\alpha}\right) - (1-R)H\left(\frac{\tau}{1-R}\right) \right)}{H(\tau) - (1-R)H\left(\frac{\tau}{1-R}\right)},$$

for $\alpha = (1-\delta)R$.

Proof. Assume that the permutation P distributes the error such that

$$P^{-1}\mathbf{e} = (\mathbf{e}_1, 0^{\alpha n}), \quad (12)$$

for α as defined in Algorithm 5. Then it follows, that \mathbf{e}_1 is a solution to syndrome decoding instance $((HP)_{\llbracket (1-\alpha)n \rrbracket}, \mathbf{s}, \omega)$. By the correctness of our quantum circuit the solution \mathbf{e}_1 is returned in line 4 and finally $\mathbf{e} = P(\mathbf{e}_1, 0^{\alpha n})$ is recovered.

Next let us analyze the running time of the algorithm. The probability of a random permutation distributing the error weight as given in Eq. (12) is

$$q_C := \Pr [P^{-1}\mathbf{e} = (\mathbf{e}_1, 0^{\alpha n})] = \frac{\binom{(1-\alpha)n}{\omega}}{\binom{n}{\omega}}.$$

Hence, we expect that after q_C^{-1} random permutations one of them induces the desired weight-distribution. The asymptotic time complexity for the execution of the quantum circuit to solve the corresponding $\mathcal{SD}_{(1-\alpha)n, (R-\alpha)n, \omega}$ problem can be derived from Eq. (7) as

$$T_Q = \tilde{O} \left(\sqrt{\frac{\binom{(1-\alpha)n}{\omega}}{\binom{(1-R)n}{\omega}}} \right).$$

Since for each classically chosen permutation we need to execute our quantum circuit the total running time becomes

$$T = q_C^{-1} \cdot T_Q = \tilde{O} \left(\frac{\binom{n}{\omega}}{\sqrt{\binom{(1-\alpha)n}{\omega} \binom{(1-R)n}{\omega}}} \right).$$

Now let us determine $t(\delta) := \frac{\log T}{\log T_C}$. First observe that $T = \frac{T_C}{T_Q}$, which can be rewritten as

$$\begin{aligned} \log T_C - \log T_Q &= \log T \\ \Leftrightarrow 1 - \frac{\log T_Q}{\log T_C} &= \frac{\log T}{\log T_C} =: t(\delta). \end{aligned}$$

An approximation of T_Q and T_C via the approximation for binomial coefficients given in Eq. (1) together with $\omega := \tau n$ and $k := Rn$ then yields

$$t(\delta) = 1 - \frac{\frac{1}{2} \left((1-\alpha)H\left(\frac{\tau}{1-\alpha}\right) - (1-R)H\left(\frac{\tau}{1-R}\right) \right)}{H(\tau) - (1-R)H\left(\frac{\tau}{1-R}\right)},$$

as claimed. Note that the input matrix of an $[(1-\alpha)n, (R-\alpha)n]$ -code requires $(1-R)(R-\alpha)n^2$ qubits for the matrix representation (compare to Eq. (10)). Hence, by setting $\alpha = (1-\delta)R$ we obtain a qubit reduction by

$$\frac{(1-R)(R-\alpha)n^2}{(1-R)Rn^2} = \frac{R - (1-\delta)R}{R} = \delta. \quad \square$$

Note that in the case of a sublinear error-weight, which is e.g. the case for the McEliece, BIKE and HQC crypto systems, T_C can be expressed as

$$T_C = \tilde{O} \left(\frac{\binom{n}{\omega}}{\binom{(1-R)n}{\omega}} \right) = \tilde{O} \left((1-R)^{-\omega} \right), \quad (13)$$

as shown in [36].

This allows us to simplify the statement of Theorem 1 in the following corollary.

Corollary 1 (Hybrid Prange for sublinear error weight). *Let all parameters be as in Theorem 1. For $\tau = o(1)$, we have*

$$t(\delta) = \frac{1}{2} \cdot \left(1 + \frac{\log(1 - (1-\delta)R)}{\log(1-R)} \right).$$

Proof. First we approximate T_Q similar to T_C in Eq. (13) as

$$T_Q = \tilde{O} \left(\sqrt{\frac{\binom{(1-\alpha)n}{\omega}}{\binom{(1-R)n}{\omega}}} \right) = \tilde{O} \left(\left(\frac{1-\alpha}{1-R} \right)^{\frac{\omega}{2}} \right).$$

Now we can derive the statement of the corollary as

$$\begin{aligned} t(\delta) &= 1 - \frac{\log T_Q}{\log T_C} = 1 - \frac{\frac{\omega}{2} (\log(1-\alpha) - \log(1-R))}{-\omega \log(1-R)} \\ &= \frac{1}{2} \cdot \left(1 + \frac{\log(1 - (1-\delta)R)}{\log(1-R)} \right). \quad \square \end{aligned}$$

Fig. 8 visualizes the relation between the qubit-reduction factor and the speedup for different choices of the code- and error-rate. We compare the full distance decoding setting with worst-case rate $R = 0.5$ and, hence, $\tau = H^{-1}(R) \approx 0.11$ and the half distance case with $\tau = \frac{H^{-1}(R)}{2}$ to the code parameters of the

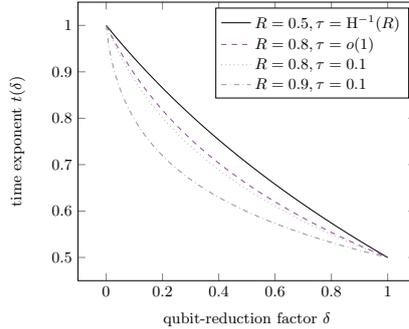


Fig. 8: Time exponent (y-axis) achieved by Theorem 1 plotted as a function of the qubit-reduction factor δ (x-axis).

McEliece scheme, which are $R = 0.8$ and $\tau = o(1)$, and the parameters of the BIKE and HQC schemes, which are specified as $R = 0.5$ and $\tau = o(1)$. Additionally, we give comparisons to higher code- and error-rates. It can be observed that the best results are obtained for high rates, where the code-rate is the more significant factor, which lies in favour to mounting an attack against codes using McEliece parameters. Note that especially for a rate close to 0.5 the trade-off is very insensitive to changes in the error-rate, such that the behaviour for the settings of full and half distance as well as BIKE and HQC are almost identical, hence, we only included the full distance case for the sake of clarity.

To give a concrete example, our HYBRID-PRANGE algorithm allows for a reduction of the necessary qubits by 80% (corresponding to $\delta = 0.2$), while still achieving a speedup of $t(\delta) \approx 0.82$ in the McEliece setting.

6.2 Puncturing the code

While our HYBRID-PRANGE decreases the amount of necessary qubits by shortening the code, our second trade-off instead aims at puncturing the code. In a nutshell we consider only $(1 - \beta)n - k$ parity-check equations, rather than all $n - k$, which is equivalent to omitting βn rows of the parity-check matrix. The subsequently applied quantum circuit, hence, needs fewer qubits to represent matrix and syndrome. The advantage over HYBRID-PRANGE partly comes from the fact that each row saves n instead of only $n - k$ bits. Also the generated classical overhead is significantly smaller. This variant has similarities with the Canteaut-Chabaud improvement [13]. Here only a certain amount of columns (originally only one) of the identity part are exchanged in each iteration rather than drawing a completely new permutation. In our case we fix βn columns of the permutation classically and then search for the remaining $n - k - \beta n$ quantumly. In addition we introduce a different weight distribution on the fixed columns, which does not yield improvements in a purely classical setting.

We again start with a parity-check matrix H in systematic form. Now consider the projection of H onto its first $n - k - \beta n$ rows, we call the resulting

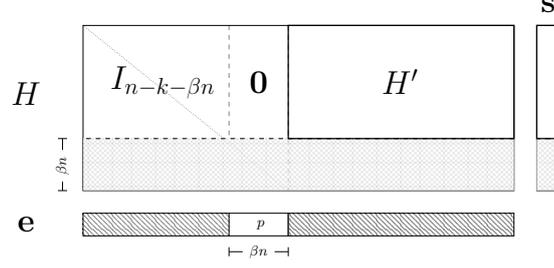


Fig. 9: Parity-check matrix where βn rows are omitted and \mathbf{e} contains weight p on βn coordinates. Framed parts are used as input to the quantum algorithm.

matrix \tilde{H} . Clearly, a solution \mathbf{e} to the instance (H, \mathbf{s}, ω) is still a solution to the instance $(\tilde{H}, \mathbf{s}_{[n-k-\beta n]}, \omega)$. Moreover, the matrix \tilde{H} includes βn zero columns, which can safely be removed (compare to Fig. 9). This results in a matrix $\tilde{H}' = (I_{n-k-\beta n} \mid H') \in \mathbb{F}_2^{(n-k-\beta n) \times (1-\beta)n}$ corresponding to an $[(1-\beta)n, k]$ code. Still, by removing the corresponding coordinates from \mathbf{e} we obtain a solution \mathbf{e}' to the instance $(\tilde{H}', \mathbf{s}_{[n-k-\beta n]}, \omega - p)$, where $p := \text{wt}(\mathbf{e}_{[n-k-\beta n+1, n-k]})$ is the weight of coordinates removed from \mathbf{e} . Eventually, once \mathbf{e}' is recovered we can obtain \mathbf{e} in polynomial time by solving the respective linear system.

A crucial observation is that disregarding βn parity-check equations could lead to the existence of multiple solutions to the reduced instance, i.e. multiple \mathbf{e}' satisfying $\tilde{H}'\mathbf{e}' = \mathbf{s}_{[n-k-\beta n]}$ but yielding an \mathbf{e} with $\text{wt}(\mathbf{e}) > \omega$. We can control this amount of solutions by shifting more weight onto the removed coordinates. Also our algorithm compensates for multiple solutions by recovering all solutions to the reduced instance by repeated executions of the quantum circuit. A pseudocode description of this trade-off is given in Algorithm 6.

In the following theorem we first state the time complexity of Algorithm 6 in dependence on the qubit reduction factor δ . After this we derive the speedup $t(\delta)$ in a separate corollary.

Theorem 2 (Punctured Hybrid). *Let $n \in \mathbb{N}$, $\omega \in [n]$ and $k = Rn$ for $R \in [0, 1]$. Then for any qubit reduction factor $\delta \in [0, 1]$ Algorithm 6 solves the $\mathcal{SD}_{n,k,\omega}$ problem in expected time T_{PH} using $\delta(1-R)Rn^2$ qubits for the matrix representation, where*

$$T_{\text{PH}} = \tilde{O} \left(\frac{\binom{n}{\omega}}{\sqrt{\binom{(1-\beta)n}{\omega-p} \binom{(1-\beta-R)n}{\omega-p} \binom{\beta n}{p}}} \cdot \max \left(1, \sqrt{\binom{(1-\beta)n}{\omega-p} \cdot 2^{-(1-\beta-R)n}} \right) \right)$$

with $\beta = (1-\delta)(1-R)$ and $p \in [\min(\omega, \beta n)]$.

Proof. Assume that the permutation distributes the error weight, such that for $P^{-1}\mathbf{e} = (\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3) \in \mathbb{F}_2^{(1-\beta-R)n} \times \mathbb{F}_2^{\beta n} \times \mathbb{F}_2^{Rn}$ it holds $\text{wt}(\mathbf{e}_2) = p$. Now consider

Algorithm 6 Punctured Hybrid

Require: parity-check matrix $H \in \mathbb{F}_2^{(n-k) \times n}$, syndrome $\mathbf{s} \in \mathbb{F}_2^{n-k}$, weight $\omega \in [n]$,
qubit reduction factor $\delta \in \llbracket 0, 1 \rrbracket$

Ensure: error vector \mathbf{e} with $\text{wt}(\mathbf{e}) = \omega$ satisfying $H\mathbf{e} = \mathbf{s}$

- 1: choose p accordingly
- 2: $\beta := (1 - \delta)(1 - \frac{k}{n})$, $S := \frac{\binom{(1-\beta)n}{\omega-p}}{2^{(1-\beta)n-k}}$
- 3: **repeat**
- 4: choose random permutation matrix $P \in \mathbb{F}_2^{n \times n}$ and set $\tilde{H} \leftarrow HP$
- 5: transform \tilde{H} to systematic form, $\tilde{H} = \begin{pmatrix} I_{n-k-\beta n} & \mathbf{0} & H'_1 \\ \mathbf{0} & I_{\beta n} & H'_2 \end{pmatrix}$ with syndrome $\tilde{\mathbf{s}}$
- 6: $\tilde{H}' \leftarrow (I_{n-k-\beta n} \mid H'_1)$, $\mathbf{s}' \leftarrow \tilde{\mathbf{s}}_{[(1-\beta)n-k]}$
- 7: **for** $i = 1$ **to** $\text{poly}(n) \cdot S$ **do**
- 8: solve instance $(\tilde{H}', \mathbf{s}', \omega - p)$ via quantum algorithm returning $\mathbf{e}' \in \mathbb{F}_2^{(1-\beta)n}$
- 9: $\mathbf{e}'' \leftarrow H'_2 \mathbf{e}'_{[n-k-\beta n+1, (1-\beta)n]} + \tilde{\mathbf{s}}_{[n-k-\beta n+1, n-k]}$
- 10: **if** $\text{wt}(\mathbf{e}'') \leq p$ **then**
- 11: $\mathbf{e} \leftarrow P(\mathbf{e}'_{[n-k-\beta n]}, \mathbf{e}'', \mathbf{e}'_{[n-k-\beta n+1, (1-\beta)n]})$
- 12: **break**
- 13: **until** $H\mathbf{e} = \mathbf{s}$
- 14: **return** \mathbf{e}

the permuted parity-check matrix in systematic form \tilde{H} as given in line 5 of Algorithm 6 with corresponding syndrome $\tilde{\mathbf{s}}$. We obtain

$$\tilde{H}P^{-1}\mathbf{e} = (\mathbf{e}_1 + H'_1\mathbf{e}_3, \mathbf{e}_2 + H'_2\mathbf{e}_3) = \tilde{\mathbf{s}}.$$

This implies that $(\mathbf{e}_1, \mathbf{e}_3)$ is a solution to the syndrome decoding instance $(\tilde{H}', \mathbf{s}', \omega - p)$ with $\tilde{H}' = (I_{(1-\beta-R)n} \mid H'_1)$ and $\mathbf{s}' = \tilde{\mathbf{s}}_{[(1-\beta-R)n]}$. The solution is then recovered by the application of our quantum circuit in line 8. Note that in expectation there exist

$$S := \binom{(1-\beta)n}{\omega-p} \cdot 2^{-(1-\beta-R)n}$$

solutions to our reduced instance. Since we apply our quantum circuit $\text{poly}(n) \cdot S$ times and in each execution a random solution is returned, a standard coupon collector argument yields that we recover all S solutions with high probability. Now, when $\mathbf{e}' = (\mathbf{e}_1, \mathbf{e}_3)$ is returned by the quantum circuit, we recover $\mathbf{e}_2 = \tilde{\mathbf{s}}_{[(1-\beta-R)n+1, (1-R)n]} + H'_2\mathbf{e}_3$ and eventually return $\mathbf{e} = P(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)$.

Next let us consider the time complexity of the algorithm. Observe that the probability, that $\text{wt}(\mathbf{e}_2) = p$ for a random permutation holds is

$$q_C := \Pr[\text{wt}(\mathbf{e}_2) = p] = \frac{\binom{(1-\beta)n}{\omega-p} \binom{\beta n}{p}}{\binom{n}{\omega}}.$$

Hence, after q_C^{-1} iterations we expect that there is at least one iteration where $\text{wt}(\mathbf{e}_2) = p$. In each iteration we apply our quantum circuit $\tilde{\mathcal{O}}(S)$ times to solve

the reduced instance $(\tilde{H}', \mathbf{s}', \omega - p)$, corresponding to an $[(1 - \beta)n, (1 - R)n]$ -code. Since there exist S solutions the expected time to retrieve one of them at random is

$$T_Q = \tilde{\mathcal{O}} \left(\sqrt{\frac{\binom{(1-\beta)n}{\omega-p}}{\max(1, S) \cdot \binom{(1-\beta-R)n}{\omega-p}}} \right),$$

according to Remark 2. The maximum follows since we know that there exists at least one solution. In summary the running time becomes $T_{\text{PH}} = q_C^{-1} \cdot T_Q \cdot \max(1, S)$, as stated in the theorem.

The required amount of qubits of the quantum circuit for solving the syndrome decoding problem related to an $[(1 - \beta)n, (1 - R)n]$ -code are roughly $R(1 - \beta - R)n^2$ (compare to Eq. (10)). Thus, for $\beta := (1 - \delta)(1 - R)$ this corresponds to a qubit reduction of

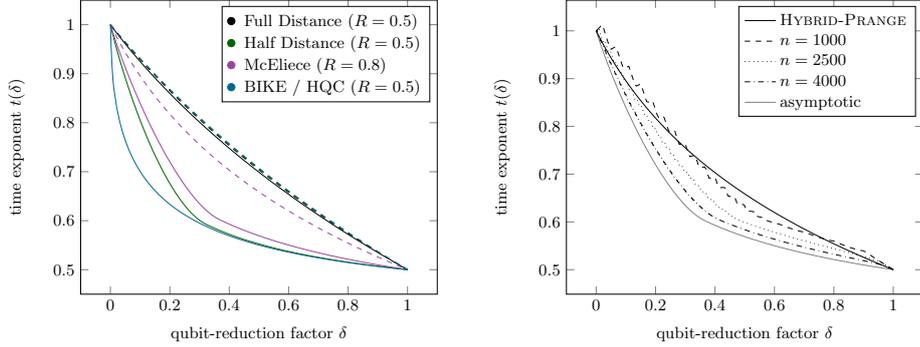
$$\frac{R(1 - \beta - R)}{R(1 - R)} = \frac{1 - R - (1 - \delta)(1 - R)}{1 - R} = \delta. \quad \square$$

Theorem 2 allows to easily determine the corresponding speedup, whose exact formula we give in Corollary 2 in Appendix II.

In Fig. 10a we compare the behavior of our new trade-off to our previously obtained HYBRID-PRANGE. Recall that the performance of HYBRID-PRANGE is not very sensitive to changes in the weight. Thus, for settings with a rate of $R = 0.5$ the dashed lines are almost on top of each other. The value p of our new trade-off (Theorem 2) were optimized numerically. It can be observed, that our second trade-off outperforms the first one for all parameters. We observe the best behaviour for low coderates and small error weights, which correspond to the case, where the solution is very unique. In these cases our PUNCTURED-HYBRID algorithm can disregard parity-check equations without introducing multiple solutions to the reduced instance. Hence, still a single execution of the quantum circuit suffices to recover the solution. Note that in the McEliece, BIKE and HQC setting the error weight is sublinear, which is in favour of our new trade-off. BIKE and HQC furthermore use a very small error weight of only $\mathcal{O}(\sqrt{n})$ and specify a rate of $R = 0.5$, which results in a very unique solution. Consequently, in Fig. 10a it can be observed, that asymptotically for these settings the second trade-off improves drastically on HYBRID-PRANGE.

Note that our formulation of the speedup for PUNCTURED-HYBRID in contrast to HYBRID-PRANGE (see Corollary 1) still depends on the error-rate, not exactly allowing for $\omega = o(n)$. Thus, to obtain the asymptotic plot we compared the result of Corollary 1 to Theorem 2 for McEliece [6688, 5024, 128], BIKE [81946, 40973, 264] and HQC [115274, 57637, 262], which are the suggested parameters for 256-bit security from the corresponding NIST submission documentation [2, 14, 30].

To quantify the result of our new trade-off take e.g. the case of McEliece and a qubit reduction by 80% ($\delta = 0.2$), as before. Here we improve to a speedup of $t(\delta) \approx 0.74$, compared to 0.82 for HYBRID-PRANGE.



(a) Asymptotic time exponents. New Theorem 2 depicted as solid line, Corollary 1 as dashed line.

(b) Time exponents for concrete parameter sets. McEliece parameter sets satisfy $k = 0.8n$ and $\omega = \lfloor \frac{n}{5 \log n} \rfloor$.

Fig. 10: Comparison of time exponents of HYBRID-PRANGE and PUNCTURED-HYBRID (y-axis) plotted as a function of the qubit-reduction factor δ (x-axis).

However, for concrete medium sized parameters this asymptotic behaviour is not necessarily obtained. In Fig. 10b we therefore show a comparison of both trade-offs for concrete McEliece parameter sets. Here we furthermore used the more accurate time complexity formula involving binomial coefficients rather than its asymptotic approximation to compute the speedup $t(\delta)$. Note that the discontinuity for our new trade-off in these cases is due to the limitation to discrete choices of p . We find that for parameters up to $n \approx 2500$ both trade-offs remain superior to each other for certain reduction factors δ . For larger values of n the PUNCTURED-HYBRID algorithm becomes favourable for all δ . In the BIKE and HQC settings the PUNCTURED-HYBRID algorithm is favourable already for small parameters corresponding to $n = 1000$.

6.3 Combining both trade-offs

Next we show how to combine both previous trade-offs to achieve an improved version. Therefore we first reduce the code length and dimension, again by guessing αn zero coordinates of \mathbf{e} and removing the corresponding columns from H . The remaining instance is then solved using our PUNCTURED-HYBRID algorithm (compare also to Fig. 11). If the initial guess was wrong, this procedure will not finish. Thus, we introduce an abort of the execution after the expected amount of iterations of PUNCTURED-HYBRID on a correct guess.

The pseudocode of the procedure is given in Algorithm 7. Note that here we use β and p as input parameters to PUNCTURED-HYBRID, rather than to the choice made in Algorithm 6 (PUNCTURED-HYBRID).

Theorem 3 (Combined Hybrid). *Let $n \in \mathbb{N}$, $\omega \in [n]$ and $k = Rn$ for $R \in \llbracket 0, 1 \rrbracket$. Then for any qubit reduction factor $\delta \in \llbracket 0, 1 \rrbracket$ the $\mathcal{SD}_{n,k,\omega}$ problem*

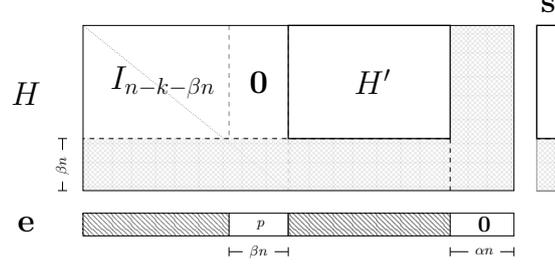


Fig. 11: Input matrix in systematic form where βn parity-check equations are omitted and αn zeros of \mathbf{e} are known. The vector \mathbf{e} is assumed to contain weight p on βn coordinates. Framed parts are used as input to the quantum algorithm.

Algorithm 7 COMBINED-HYBRID

Require: parity-check matrix $H \in \mathbb{F}_2^{(n-k) \times n}$, syndrome $\mathbf{s} \in \mathbb{F}_2^{n-k}$, weight $\omega \in [n]$, qubit reduction factor $\delta \in \llbracket 0, 1 \rrbracket$

Ensure: error vector \mathbf{e} with $\text{wt}(\mathbf{e}) = \omega$ satisfying $H\mathbf{e} = \mathbf{s}$

- 1: choose α and p accordingly
 - 2: $\beta := (1 - \frac{k}{n}) \left(\frac{\delta \frac{k}{n}}{\frac{k}{n} - \alpha} \right)$, $E := \frac{\binom{(1-\alpha)n}{\omega}}{\binom{(1-\alpha-\beta)n}{\omega-p} \binom{\beta n}{p}}$
 - 3: **repeat**
 - 4: choose random permutation matrix $P \in \mathbb{F}_2^{n \times n}$ and set $\tilde{H} \leftarrow HP$
 - 5: $\mathbf{e}' \leftarrow \text{PUNCTURED-HYBRID}(\tilde{H}_{[(1-\alpha)n]}, \mathbf{s}, \omega, \delta, \frac{\beta}{1-\alpha}, p)$ \triangleright abort after E iterations of the outer loop
 - 6: $\mathbf{e} \leftarrow P(\mathbf{e}', 0^{\alpha n})$
 - 7: **until** $H\mathbf{e} = \mathbf{s}$
 - 8: **return** \mathbf{e}
-

can be solved in expected time T_{CH} using $\delta(1-R)Rn^2$ qubits for the matrix representation, where

$$T_{\text{CH}} = \tilde{O} \left(\frac{\binom{n}{\omega}}{\sqrt{\binom{(1-\alpha-\beta)n}{\omega-p} \binom{(1-\beta-R)n}{\omega-p} \binom{\beta n}{p}}} \cdot \max \left(1, \sqrt{\binom{(1-\alpha-\beta)n}{\omega-p}} \cdot 2^{-(1-\beta-R)n} \right) \right)$$

with $\alpha \in \llbracket 0, R \rrbracket$, $\beta = (1-R) \left(1 - \frac{\delta R}{R-\alpha} \right)$ and $p \in [\min(\omega, \beta n)]$.

Proof. The correctness follows from the correctness of Algorithm 5 and Algorithm 6. Therefore observe that for a correct guess of αn zero positions of \mathbf{e} , the expected amount of permutations needed by PUNCTURED-HYBRID to find the solution is

$$E := \frac{\binom{(1-\alpha)n}{\omega}}{\binom{(1-\alpha-\beta)n}{\omega-p} \binom{\beta n}{p}}.$$

Also note that PUNCTURED-HYBRID is called on a code of length $n' = (1-\alpha)n$. Hence setting $\beta' = \frac{\beta}{1-\alpha}$ guarantees that $\beta'n' = \beta n$ parity equations are omitted.

For the time complexity we have again with probability

$$q_C := \Pr [P^{-1} \mathbf{e} = (\mathbf{e}_1, 0^{\alpha n})] = \frac{\binom{(1-\alpha)n}{\omega}}{\binom{n}{\omega}},$$

a correct guess for αn zero positions (compare to the proof of Theorem 1). In each iteration of our combined algorithm we call the PUNCTURED-HYBRID algorithm. Inside this subroutine E iterations of the outer loop are executed, each performing

$$S = \tilde{\Theta} \left(\max \left(1, \frac{\binom{(1-\beta-\alpha)n}{\omega-p}}{2^{-(1-R-\beta)n}} \right) \right)$$

calls to the quantum circuit. This quantum circuit is applied to solve the syndrome decoding problem defined on an $[(1-\alpha-\beta)n, (R-\alpha)n]$ -code with error-weight $\omega-p$ (compare to Fig. 11), which takes time

$$T_Q = \tilde{\mathcal{O}} \left(\sqrt{\frac{\binom{(1-\alpha-\beta)n}{\omega-p}}{S \cdot \binom{(1-\beta-R)n}{\omega-p}}} \right).$$

Thus, eventually, the time complexity of the whole algorithm summarizes as $T_{\text{CH}} = q_C^{-1} \cdot E \cdot T_Q \cdot S$, as claimed. Finally, note that for given $\beta = (1-R) \left(1 - \frac{\delta R}{R-\alpha}\right)$ we obtain a qubit reduction by

$$\frac{(R-\alpha)(1-R-\beta)}{R(1-R)} = \frac{(R-\alpha)(1-R)(1 - (1 - \frac{\delta R}{R-\alpha}))}{R(1-R)} = \delta. \quad \square$$

Next we give a comparison of the trade-off behavior in different settings. On the left in Fig. 12 we illustrate the asymptotic behaviors of the trade-offs, where p and α for the combined trade-off were numerically optimized. It shows that the combination of both trade-offs (dashed lines) for most parameters improves on PUNCTURED-HYBRID (solid line). Especially in the full distance decoding setting an improvement for nearly all δ is achieved. This is due to the fact, that the guessing of zero coordinates is an additional possibility to control the amount of solutions to the reduced instance and therefore to optimize the complexity of the PUNCTURED-HYBRID subroutine. This is also the reason why we achieve no (asymptotic) improvement in the BIKE and HQC settings, here the solution is already so unique that the trade-off can not benefit from the new degree of freedom.

But also in the McEliece setting we achieve notable improvements. If we again consider a reduction-factor of $\delta = 0.2$ the combination improves the speedup to $t(\delta) \approx 0.69$ from 0.74 achieved by PUNCTURED-HYBRID. Furthermore, when focusing on near future realizations, i.e., the regime of small reduction factors, it is for example possible with just one percent of the qubits ($\delta = 0.01$) to achieve a speedup of $t(\delta) \approx 0.92$.

On the right in Fig. 12 we show the relation between qubit reduction and speedup for concrete McEliece parameters. Here we restrict ourselves to the

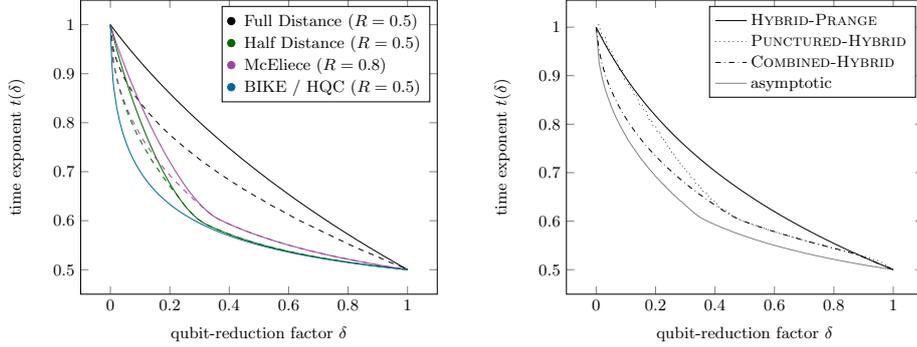


Fig. 12: Comparison of time exponents (y-axis) plotted as a function of the qubit-reduction factor δ (x-axis) for asymptotic (left, combined trade-off illustrated dashed, Theorem 2 solid) and concrete McEliece parameters (right).

instance $n = 2500$ for the sake of clarity. But note that for all parameter sets at least the minimum of both trade-offs is obtained with improvements especially for low reduction factors. In the BIKE and HQC setting for small parameter sets with $n \leq 3000$ we achieve (small) improvements in the regime of $\delta \leq 0.05$.

6.4 Overview and Discussion

For convenience we state in Table 1 the parameters of the reduced instances solved by the quantum circuits within each of our trade-offs. Table 2 then states the necessary amount of qubits and the resulting depth of the circuits to solve a respective instance with parameters (n, k, ω) .

	Quantum Instance Parameters		
	n'	k'	ω'
HYBRID-PRANGE	$(1 - \alpha)n$	$(R - \alpha)n$	ω
PUNCTURED-HYBRID	$(1 - \beta)n$	Rn	$\omega - p$
COMBINED-HYBRID	$(1 - \alpha - \beta)n$	$(R - \alpha)n$	$\omega - p$

Table 1: Parameters of the reduced subinstance solved by the quantum circuit called by the respective classical co-processor.

By plugging in the values from Table 1 into the formulas given in Table 2 one receives the quantum complexities of the respective classical co-processor. Here we differentiate between optimization regarding the amount of qubits and the circuit depth. The essential difference lies in the use of the Lee-Brickell improvement in case of an optimization of the depth, while the width optimized variant uses the qubit reduction technique from Section 5.3.

The depth of our circuits is mainly dominated by the application of the Gaussian elimination, where the additional $\log n$ factor results from the decomposition of multi-controlled gates. Remember that q is the proportion of subsets yielding a solution among all size- $(n - k)$ subsets in the case of Prange, given in Eq. (3), while q_{LB} is the proportion of good subsets in the case of the Lee-Brickell algorithm stated in Eq. (9).

	W -Optimized	D -Optimized	D -Optimized (cyclic)
Qubit	$(n - k + 2)(k + 3) - 7$	$(n - k + 1) \cdot (n + 2) - 3$	$(n - k) \cdot (n + k + 2) - 1$
Depth	$\mathcal{O}\left(\frac{n^3 \log n}{\sqrt{q}}\right)$	$\mathcal{O}\left(\frac{n^3 \log n + p \binom{k}{p}}{\sqrt{q_{\text{LB}}}}\right)$	$\mathcal{O}\left(\frac{n(n^2 \log n + p \binom{k}{p})}{\sqrt{k \cdot q_{\text{LB}}}}\right)$

Table 2: Depth and required qubits of the quantum circuit to solve the $\mathcal{SD}_{n,k,\omega}$.

Overall, we presented concrete depth and width optimized quantum circuits for the fully-fledged ISD procedure. Our tradeoffs put a special focus on the reduction of necessary qubits, targeting *near-term* realizations. Following this thought, we also provide the necessary implementations of our circuits in the simulation library Qibo making a transition to a real quantum computer as easy as possible.

Although we placed a strong focus on the circuit width, we have shown that ISD can also be implemented efficiently on a quantum computer from a depth perspective. Thus, we cannot confirm the mentioned statement of the McEliece submission regarding a higher overhead when applying a Grover search to ISD rather than AES. However, we admit that a single application of AES has a lower complexity than one iteration of an ISD algorithm, which lies in favor of the quantum security of code-based schemes. Since NIST imposes a depth-limitation on the used quantum circuits, the more depth is needed for the implementation of one iteration, the less Grover iterations can be performed. We made a first step in the direction of overcoming this issue by giving a quantized version of the Lee-Brickell improvement and by exploiting the cyclicity in the BIKE / HQC cases. Both approaches tackle the problem by reducing the number of necessary Grover iterations. The second possibility is targeting a depth reduction of a single iteration, which is dominated by performing the Gaussian elimination. We leave it as an open problem to further study the concrete quantum circuit design of advanced Gaussian elimination procedures, such as M4RI or Strassen.

References

1. Alagic, G., Alperin-Sheriff, J., Apon, D., Cooper, D., Dang, Q., Kelsey, J., Liu, Y.K., Miller, C., Moody, D., Peralta, R., et al.: Status report on the second round of the NIST post-quantum cryptography standardization process. US Department of Commerce, NIST (2020)

2. Aragon, N., Barreto, P., Bettaieb, S., Bidoux, L., Blazy, O., Deneuville, J.C., Gaborit, P., Gueron, S., Guneyasu, T., Melchor, C.A., et al.: BIKE: bit flipping key encapsulation (2020)
3. Banegas, G., Bernstein, D.J., van Hoof, I., Lange, T.: Concrete quantum cryptanalysis of binary elliptic curves. *IACR Transactions on Cryptographic Hardware and Embedded Systems* pp. 451–472 (2021)
4. 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. *Physical review A* **52**(5), 3457 (1995)
5. Becker, A., Joux, A., May, A., Meurer, A.: Decoding random binary linear codes in $2^{n/20}$: How $1+1=0$ improves information set decoding. In: Annual international conference on the theory and applications of cryptographic techniques. pp. 520–536. Springer (2012)
6. Bernstein, D.J.: Grover vs. McEliece. In: International Workshop on Post-Quantum Cryptography. pp. 73–80. Springer (2010)
7. Bernstein, D.J., Biasse, J.F., Mosca, M.: A low-resource quantum factoring algorithm. In: International Workshop on Post-Quantum Cryptography. pp. 330–346. Springer (2017)
8. Bernstein, D.J., Lange, T., Peters, C.: Attacking and defending the McEliece cryptosystem. In: International Workshop on Post-Quantum Cryptography. pp. 31–46. Springer (2008)
9. Biasse, J.F., Bonnetain, X., Pring, B., Schrottenloher, A., Youmans, W.: A trade-off between classical and quantum circuit size for an attack against CSIDH. *Journal of Mathematical Cryptology* **15**(1), 4–17 (2020)
10. Biasse, J.F., Pring, B.: A framework for reducing the overhead of the quantum oracle for use with Grover’s algorithm with applications to cryptanalysis of SIKE. *Journal of Mathematical Cryptology* **15**(1), 143–156 (2020)
11. Both, L., May, A.: Decoding linear codes with high error rate and its impact for LPN security. In: International Conference on Post-Quantum Cryptography. pp. 25–46. Springer (2018)
12. Brassard, G., Hoyer, P.: An exact quantum polynomial-time algorithm for Simon’s problem. In: Proceedings of the Fifth Israeli Symposium on Theory of Computing and Systems. pp. 12–23. IEEE (1997)
13. Canteaut, A., Chabaud, F.: A new algorithm for finding minimum-weight words in a linear code: application to McEliece’s cryptosystem and to narrow-sense BCH codes of length 511. *IEEE Transactions on Information Theory* **44**(1), 367–378 (1998)
14. Chou, T., Cid, C., UiB, S., Gilcher, J., Lange, T., Maram, V., Misoczki, R., Niederhagen, R., Paterson, K.G., Persichetti, E., et al.: Classic mceliece: conservative code-based cryptography 10 october 2020 (2020)
15. Dumer, I.: On minimum distance decoding of linear codes. In: Proc. 5th Joint Soviet-Swedish Int. Workshop Inform. Theory. pp. 50–52 (1991)
16. Efthymiou, S., Ramos-Calderer, S., Bravo-Prieto, C., Pérez-Salinas, A., García-Martín, D., Garcia-Saez, A., Latorre, J.I., Carrazza, S.: Qibo: a framework for quantum simulation with hardware acceleration. arXiv preprint arXiv:2009.01845 (2020)
17. Efthymiou, S., Ramos-Calderer, S., Bravo-Prieto, C., Pérez-Salinas, A., García-Martín, D., Garcia-Saez, A., Latorre, J.I., Carrazza, S.: Quantum-tii/qibo: Qibo (Aug 2020). <https://doi.org/10.5281/zenodo.3997195>, <https://doi.org/10.5281/zenodo.3997195>

18. Ekerå, M., Håstad, J.: Quantum algorithms for computing short discrete logarithms and factoring RSA integers. In: International Workshop on Post-Quantum Cryptography. pp. 347–363. Springer (2017)
19. Gilbert, E.N.: A comparison of signalling alphabets. The Bell system technical journal **31**(3), 504–522 (1952)
20. Grover, L., Rudolph, T.: Creating superpositions that correspond to efficiently integrable probability distributions. arXiv preprint quant-ph/0208112 (2002)
21. Grover, L.K.: A fast quantum mechanical algorithm for database search. In: Proceedings of the twenty-eighth annual ACM symposium on Theory of computing. pp. 212–219 (1996)
22. Grover, L.K.: Quantum computers can search rapidly by using almost any transformation. Physical Review Letters **80**(19), 4329 (1998)
23. Helm, A., May, A.: The power of few qubits and collisions–subset sum below Grover’s bound. In: International Conference on Post-Quantum Cryptography. pp. 445–460. Springer (2020)
24. Jaques, S., Naehrig, M., Roetteler, M., Virdia, F.: Implementing grover oracles for quantum key search on AES and LowMC. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 280–310. Springer (2020)
25. Kachigar, G., Tillich, J.P.: Quantum information set decoding algorithms. In: International Workshop on Post-Quantum Cryptography. pp. 69–89. Springer (2017)
26. Kirshanova, E.: Improved quantum information set decoding. In: International Conference on Post-Quantum Cryptography. pp. 507–527. Springer (2018)
27. Lee, P.J., Brickell, E.F.: An observation on the security of McEliece’s public-key cryptosystem. In: Workshop on the Theory and Application of Cryptographic Techniques. pp. 275–280. Springer (1988)
28. May, A., Meurer, A., Thomae, E.: Decoding random linear codes in $\tilde{\gamma}(2^{0.054n})$. In: International Conference on the Theory and Application of Cryptology and Information Security. pp. 107–124. Springer (2011)
29. May, A., Ozerov, I.: On computing nearest neighbors with applications to decoding of binary linear codes. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 203–228. Springer (2015)
30. Melchor, C.A., Aragon, N., Bettaieb, S., Bidoux, L., Blazy, O., Deneuville, J.C., Gaborit, P., Persichetti, E., Zémor, G., Bourges, I.: Hamming quasi-cyclic (HQC) (2020)
31. Nielsen, M.A., Chuang, I.L.: Quantum information and quantum computation. Cambridge: Cambridge University Press **2**(8), 23 (2000)
32. Prange, E.: The use of information sets in decoding cyclic codes. IRE Transactions on Information Theory **8**(5), 5–9 (1962)
33. Ramos-Calderer, S., Esser, A.: qiboteam/qisd (Dec 2021), <https://github.com/qiboteam/qISD>
34. Sendrier, N.: Decoding one out of many. In: International Workshop on Post-Quantum Cryptography. pp. 51–67. Springer (2011)
35. Stern, J.: A method for finding codewords of small weight. In: International Colloquium on Coding Theory and Applications. pp. 106–113. Springer (1988)
36. Torres, R.C., Sendrier, N.: Analysis of information set decoding for a sub-linear error weight. In: Post-Quantum Cryptography. pp. 144–161. Springer (2016)
37. Varshamov, R.R.: Estimate of the number of signals in error correcting codes. Doklady Akad. Nauk, SSSR **117**, 739–741 (1957)
38. Zalka, C.: Grover’s quantum searching algorithm is optimal. Physical Review A **60**(4), 2746 (1999)

Appendix

I Width reduced circuit

Algorithm 8 shows the pseudocode for generating our width optimized circuit, not requiring the identity part of H as an input, described in Section 5.3.

Algorithm 8 Width-reduced Combined Circuit

Require: matrix $H' \in \mathbb{F}_2^{(n-k) \times k}$, syndrome $\mathbf{s} \in \mathbb{F}_2^{n-k}$, $(k+1)(n-k) + n + \lceil \log(n-k) \rceil$ qubits

Ensure: Uniform superposition over weight of all \mathbf{x} with $H_I \mathbf{x} = \mathbf{s}$ for $I \in \binom{[n]}{n-k}$ where $H = (I_{n-k} \mid H')$

- 1: Initialize qubits with $(H' \mid \mathbf{s})$
 - 2: Generate uniform superposition over $\mathcal{B}_{n,n-k}$ on qubits (b_1, \dots, b_n)
 - 3: $c \leftarrow n - k$
 - 4: **for** $i = 1$ **to** $n - k$ **do** ▷ row swaps depending on first $n - k$ columns
 - 5: **for** $j = 1$ **to** $i - 1$ **do**
 - 6: **if** $b_i = 1$ **and** $c = n - k - j + 1$ **then** swap row i and row j
 - 7: **if** $b_i = 1$ **then** $c \leftarrow c - 1$
 - 8: **for** $i = n$ **down to** $n - k + 1$ **do** ▷ swap selected columns of H' to the back
 - 9: **if** $b_i = 1$ **then**
 - 10: $c \leftarrow c - 1$
 - 11: $x_i \leftarrow 1$
 - 12: **for** $j = i$ **to** $n - 1$ **do**
 - 13: swap column j and $j + 1$
 - 14: swap x_j and x_{j+1}
 - 15: Apply Gaussian elimination circuit starting with the last column, where each operation depending on column j is controlled by x_j (Algorithm 3)
 - 16: **return** $c \leftarrow$ weight of last column
-

II Punctured Hybrid

In the following corollary we state the exact form of the speedup $t(\delta)$ for our PUNCTURED-HYBRID (Theorem 2).

Corollary 2 (Punctured Hybrid Speedup). *Let $n \in \mathbb{N}$, $\omega = \tau n$ and $k = Rn$, $p = \rho n$ for $\tau, R, \rho \in \llbracket 0, 1 \rrbracket$ and let T_C be as defined in Eq. (11). Then for any qubit reduction factor $\delta \in \llbracket 0, 1 \rrbracket$ Algorithm 6 solves the $\mathcal{SD}_{n,k,\omega}$ problem in time*

$(T_C)^{t(\delta)}$ using $\delta(1-R)Rn^2$ qubits for the matrix representation, where

$$t(\delta) = \frac{\mathbb{H}(\tau) - \beta\mathbb{H}\left(\frac{\rho}{\beta}\right) - \frac{1-\beta}{2} \cdot \mathbb{H}\left(\frac{\tau-\rho}{1-\beta}\right) - \frac{(1-\beta-R)}{2} \cdot \mathbb{H}\left(\frac{\tau-\rho}{1-\beta-R}\right) + \max(0, \sigma)}{\mathbb{H}(\tau) - (1-R)\mathbb{H}\left(\frac{\tau}{1-R}\right)}$$

for $\beta = (1-\delta)(1-R)$ and $\sigma = (1-\beta)\mathbb{H}\left(\frac{\tau-\rho}{1-\beta}\right) - (1-\beta-R)$.

Proof. Recall that $t(\delta) = \frac{\log T_{\text{PH}}}{\log T_C}$, where T_{PH} is the running time of Algorithm 6, given in Theorem 2. Now the statement of the corollary follows immediately by approximating the binomial coefficients in T_{PH} and T_C via Stirling's formula (see Eq. (1)). \square