Improved Quantum Circuits for Information Set Decoding with Application to Code-Based Cryptography

Hyunji Kim[1], Kyungbae Jang[2], and Hwajeong Seo[1†]

[1] Division of IT Convergence Engineering, Hansung University, Seoul, South Korea
[2] College of Computing and Data Science, Nanyang Technological University, Singapore
khj1594012@gmail.com, starj1023@gmail.com, hwajeong84@gmail.com

Abstract. Information set decoding (ISD) is the standard generic decoding attack considered for code-based cryptography. A concrete quantum-resource estimate for Grover-accelerated ISD requires an oracle whose dominant component is Gauss–Jordan elimination.

We improve the elimination circuit of Perriello et al. [25] and Jang et al. [15] by not updating the entries that no later pivot or the final weight predicate reads. The required result vector is recovered by a parallel back-substitution on the syndrome register. For the target schemes, our elimination circuit improves the qubit count by about 22% compared to [15]. The Toffoli count improves by about 20% compared to both [25] and [15]. The Toffoli depth improves by about 67% compared to [25] but degrades by about 0.3% compared to [15].

We report logical resource estimates for the quantum ISD attack on HQC and Classic McEliece. The product of the total gate count and the full depth exceeds the NIST post-quantum security thresholds, and the full depth exceeds the MAXDEPTH upper bound.

We also provide fault-tolerant estimates of the physical qubit count and the runtime under a surface-code model with magic-state distillation. As one example, HQC-128 requires about 2^{43} physical qubits and about 2^{53} years at a $1\,\mu$s code cycle.

Keywords: Grover's Algorithm · Code-based Cryptography · Quantum Cryptanalysis · Information Set Decoding

1 Introduction

Code-based cryptography is a leading family of post-quantum cryptography: NIST has recently selected HQC [1] for standardization as a key-encapsulation mechanism [23], while Classic McEliece [6] remains a long-studied, conservative alternative. Both schemes rely on the hardness of syndrome decoding, believed to resist both classical and quantum attack, though their code families differ: Classic McEliece uses binary Goppa codes and HQC uses quasi-cyclic codes.

†: Corresponding author.

Knowing how much quantum effort such an attack would actually take is therefore important for setting and trusting their parameters.

The best-known generic decoding attack on these schemes is *information set decoding* (ISD), a family of decoders that still runs in exponential time [26]. On a quantum computer, Grover's search lowers the cost of ISD by a square root [5], so its concrete cost is a central measure of post-quantum security against generic decoding. Turning this idea into a resource estimate requires an explicit quantum circuit for one ISD iteration, whose dominant component is a Gauss–Jordan elimination over $\mathbb{F}_2$.

Several circuit-level studies give concrete estimates: Perriello et al. presented the first complete quantum circuit for ISD [24] and later improved it [25], reporting resource estimates for HQC, BIKE, and Classic McEliece. Jang et al. focused on the dominant linear-algebra subroutine and proposed a depth-optimized parallel circuit for binary Gauss–Jordan elimination [15]. More recently, Chevignard et al. studied a low-qubit alternative based on Wiedemann sparse inversion, which optimizes a different point of the time–space trade-off [8].

In this work we improve the elimination oracle, lowering its qubit and gate counts with only a lower-order depth change, and evaluate the ISD attack on HQC and Classic McEliece at both the logical and physical level.

Contribution

Our contributions are the following.

- We propose a modified elimination procedure based on [15]. It keeps their parallel pivot fan-out, but avoids the redundant row updates and instead adds a parallel back-substitution that recovers $H_{\mathcal{S}}^{-1}s$ (Section 3.2). This lowers the qubit count by about 22% and the T-count by about 20%, while the T-depth increases only by the lower-order term $4r - 4$.
- We report logical estimates for the quantum ISD attack on HQC and Classic McEliece and evaluate them against the NIST criteria (Sections 4). For every target the gate count times the full depth and the full depth exceed the NIST security-level threshold and the MAXDEPTH upper bound, respectively.
- For a more realistic cost assessment, we provide fault-tolerant estimates under the surface code with magic-state distillation (Section 4.2). As one of our results, HQC-128 requires about 2^{43} physical qubits and 2^{53} years.

2 Background

2.1 Information Set Decoding

Code-based cryptosystems rely on the hardness of the *Syndrome Decoding Problem* (SDP), whose bounded-weight decision version is NP-complete [4]. For a binary $[N, k]$ code with parity-check matrix $H \in \mathbb{F}_2^{(N-k)\times N}$ and redundancy $r = N - k$, the SDP seeks, given a syndrome $s = He^\top \in \mathbb{F}_2^r$ and a weight t, an error e with $\mathrm{wt}(e) = t$.

The two schemes we analyze instantiate the SDP over different code families: Classic McEliece [6] over binary Goppa codes, and HQC [1], recently selected by NIST, over quasi-cyclic codes. For these schemes, *information set decoding* (ISD) is regarded as the best known attack.

We adopt the parameter sets from the current specifications, Classic McEliece [9, Section 7] and HQC [14, Section 4.1]. For Classic McEliece, each set lists m, N, and t, from which $r = mt$ and $k = N - r$. For HQC, the key-recovery attack decodes the underlying $[2n, n]$ quasi-cyclic code with secret weight $t = 2w$, so that $N = 2n$ and $r = n$, where n and w are the ambient length and weight given in [14, Table 5]. The circuit cost depends only on N and r, while the iteration count also depends on t.

Given an SDP instance (H, s, t), Prange's information set decoding [26] selects r column positions $\mathcal{S}$ and forms the submatrix $H_{\mathcal{S}}$. If $H_{\mathcal{S}}$ is invertible, the vector e equal to $H_{\mathcal{S}}^{-1}s$ on $\mathcal{S}$ and zero elsewhere satisfies $He^{\top} = s$, and it solves the SDP when all t error positions lie in $\mathcal{S}$.

A random $\mathcal{S}$ captures all t errors with probability $p_{\text{sup}} = \binom{r}{t}/\binom{N}{t}$ and is invertible with probability about 0.288, giving the Prange success probability p. Grover's search turns the classical $1/p$ repetitions into about $1/\sqrt{p}$ iterations, and we take the concrete counts from the Prange analysis of [25], which already includes the invertibility factor. More advanced ISD and quantum-walk variants lower the asymptotic work factor [17,19], but no circuit translation considered in [25,10] gives a clear concrete advantage under our cost model, so we adopt Prange as the circuit-level baseline [5].

2.2 Grover's Algorithm in Quantum ISD

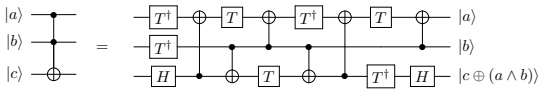
Grover's algorithm gives a quadratic speedup for unstructured search [13]. In this work, the search problem is to find a column set S for which one Prange iteration succeeds. The search register is therefore prepared in a uniform superposition over all r-column choices, not only over nonsingular choices. The oracle marks a choice S when $H_{\mathcal{S}}$ is nonsingular and the corresponding vector $H_{\mathcal{S}}^{-1}s$ has Hamming weight t.

For each S in superposition, the oracle gathers the selected columns, forms the submatrix $H_{\mathcal{S}}$, runs the elimination circuit, and checks the Hamming weight of the resulting syndrome vector. Singular branches are not aborted; they simply fail the nonsingularity condition and are not marked. If the nonsingularity and weight tests both succeed, a decision qubit, equivalently a phase oracle, flips the phase of that basis state. The work registers are then uncomputed, so that only the phase mark on the successful choices remains.

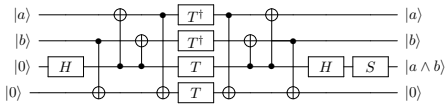
After the oracle, the diffusion operator amplifies the marked choices. Repeating the oracle and diffusion operator increases the amplitude of the successful column sets, and a measurement returns a marked set with high probability. If one classical Prange iteration succeeds with probability p, Grover search uses about $\pi/(4\sqrt{p})$ oracle calls instead of the classical $1/p$ trials. In this paper, p is the Prange success probability defined in Section 2.1, and the oracle whose cost we optimize is the Prange-oracle circuit constructed in Section 3.

2.3 Quantum Gates

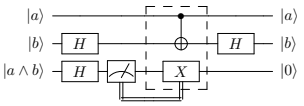
Figure 1 illustrates the quantum gates employed in our implementation. We adopt a Toffoli decomposition using 7 T-gates at a T-depth of 4 [3]. Where the target qubit is guaranteed clean, we instead use the AND gate of Jaques et al. [16], which costs 4 T-gates at a T-depth of 1 with one ancilla, and whose uncomputation $\mathrm{AND}^\dagger$ needs only Clifford gates and a measurement [12]. The input stage also uses the single-qubit rotation R_y, which lies outside Clifford+T and is synthesized into about 149 T-gates at precision 10^{-15} [20].



(a) Toffoli gate (decomposition from [3]).



(b) AND gate.



(c) $\mathrm{AND}^\dagger$ gate.

Figure 1: Non-Clifford primitives used in the resource estimates.

3 Quantum Circuit for Information Set Decoding

This section presents our quantum circuit for the Grover oracle. The oracle marks the column sets $\mathcal{S}$ in superposition for which $H_{\mathcal{S}}^{-1}s$ has weight t. The QISD oracle is composed of the following three arithmetic modules. The *input setting* first prepares the superposition and gathers the selected columns into the submatrix $H_{\mathcal{S}}$ (see Section 3.1). From this submatrix, the *Gauss–Jordan elimination* computes $H_{\mathcal{S}}^{-1}s$ (see Section 3.2), after which a *Hamming weight check* determines the decision qubit whenever the result has weight t (see Section 3.3). The elimination is by far the most expensive stage, so we focus on optimizing the Gauss–Jordan elimination circuit.

3.1 Input Setting

The input setting places the $\binom{N}{r}$ column choices in superposition and gathers the selected submatrix $H_{\mathcal{S}}$ into a contiguous block for the later stages. The choice is a *Dicke state* $\left|D_r^N\right\rangle$ on an N-qubit index register, prepared by the construction of [21], and the marked columns of H are then moved to the left by the reversible sorting network of [25, Section 3.1], to which we refer for both constructions. Perriello et al. [25, Section 3.1 and the footnote of Table 3] fix every gate to depth one. Table 1 therefore reports the required quantum resources, re-derived under the Toffoli decomposition we adopt (see Section 2.3).

Table 1: Quantum resources for the input setting, re-expressed in closed form.

Stage	#Qubit	#T	T-depth
Dicke state [21]	N	$149\,(4Nr - 4r^2 - 2N + 1)$	–
Sorting network [25]	$Nr + N + \frac{N\ell(1+\ell)}{4}$	$\frac{N\ell(1+\ell)}{4}\,(7r + 7)$	$2\ell(1+\ell) + 4r$

N: code length, $r = N - k$, and $\ell = \lceil \log_2 N \rceil$.

3.2 Gauss–Jordan Elimination

The elimination stage transforms $[\,H_{\mathcal{S}} \mid s\,]$ until the syndrome register holds $H_{\mathcal{S}}^{-1}s$, processing the r pivots one after another. Unlike a full identity-reduction of the selected matrix, our circuit only preserves the entries needed to recover $[\,H_{\mathcal{S}}^{-1} \mid s\,]$ for the final weight test. Each pivot is handled in two steps. The first sets the diagonal entry H_{jj} to 1: if it is already 1, nothing is done, and otherwise a lower row whose entry in column j is 1 is added (i.e., XOR) to the pivot row. The second uses the pivot row to set the entries below the pivot in column j to 0, by adding the pivot row to each lower row whose entry in column j is 1.

Gauss–Jordan elimination without redundant updates. We update only the matrix elements that later steps use. This lowers both the gate count and the qubit count. At each pivot, [15, Section 3.3] adds the pivot row to every lower row with a 1 in the pivot column and updates the whole row: the left elements of the pivot (L), its right elements (R), and the syndrome elements S, requiring ancilla copies of the pivot data for each update.

Here, our key observation is that later steps do not refer to L. A later pivot updates only the right columns R, while the back-substitution and the Hamming weight test refer only to the upper-triangular part and the syndrome.

We therefore update only R and S without L (see Figure 2). The back-substitution recovers $H_{\mathcal{S}}^{-1}s$ from the retained part, so the result is identical. Removing the L-updates lowers the leading Toffoli and qubit coefficients, while the added back-substitution contributes only lower-order terms.

	col0	col1	col2	col3	col4	col5	col6	col7	s
row0	.	.	.	.	.	.	.	.	.
row1	.	.	.	.	.	.	.	.	.
row2	.	.	P	.	.	.	.	.	.
row3	L	L	.	R	R	R	R	R	S
row4	L	L	.	R	R	R	R	R	S
row5	L	L	.	R	R	R	R	R	S
row6	L	L	.	R	R	R	R	R	S
row7	L	L	.	R	R	R	R	R	S

	col0	col1	col2	col3	col4	col5	col6	col7	s
row0	.	.	.	.	.	.	.	.	.
row1	.	.	.	.	.	.	.	.	.
row2	.	.	P	.	.	.	.	.	.
row3	.	.	.	R	R	R	R	R	S
row4	.	.	.	R	R	R	R	R	S
row5	.	.	.	R	R	R	R	R	S
row6	.	.	.	R	R	R	R	R	S
row7	.	.	.	R	R	R	R	R	S

P: pivot L: lower-left updates R: lower-right updates S: s updates .: not updated at this pivot

Figure 2: Update targets at one pivot P (left: baseline [15]; right: ours).

Algorithm 1 states the full procedure, writing H for $H_{\mathcal{S}}$. Its forward loop updates only R and the syndrome. In general, some column choices give a singular $H_{\mathcal{S}}$ where not every pivot can be set, so PivotFix records the success of each pivot. The phase flip is applied only when all pivots succeed and the weight is t. On a nonsingular branch the forward elimination applies a product E of elementary row operations, leaving the unit upper-triangular matrix $U = EH_{\mathcal{S}}$ in the matrix register and the eliminated syndrome $\tilde{s} = Es$ in the syndrome register. The back-substitution solves $Ue = \tilde{s}$, returning $e = H_{\mathcal{S}}^{-1}s$ because $U^{-1}\tilde{s} = (EH_{\mathcal{S}})^{-1}Es = H_{\mathcal{S}}^{-1}s$. With $U_{i,j} = H_{ij}$ for the retained coefficients ($i < j$), it computes $e_i = \tilde{s}_i + \sum_{j=i+1}^{r-1} U_{i,j}\, e_j$ for $i = r-1$ down to 0, in place on the syndrome register.

We parallelize the back-substitution so that its effect on the depth stays small. Each solved e_j is fanned out to copies, so the updates within a row run in parallel. The fan-out for the next row further overlaps the uncomputation of the current one. Together these keep the back-substitution at lower-order T-depth. Figure 3 shows the circuit for $r = 3$.

Algorithm 1: Gauss–Jordan elimination with syndrome back-substitution.

Require: $[\,H \mid s\,]$ with $H \in \mathbb{F}_2^{r\times r}$ (H_{ik} is row i, column k)
Ensure: $s_{\text{out}} = H^{-1}s_{\text{in}}$ on the nonsingular branches

for $j = 0$ **to** $r-1$ **do** ▷ forward elimination
  $\text{ok}_j \leftarrow \textsc{PivotFix}(H, j)$ ▷ set the pivot H_{jj} to 1; record success
  for $i = j+1$ **to** $r-1$ **do**
    for $k = j+1$ **to** $r-1$ **do** ▷ right block R only
      $H_{ik} \leftarrow H_{ik} \oplus H_{ij}\, H_{jk}$
    end for
    $s_i \leftarrow s_i \oplus H_{ij}\, s_j$ ▷ syndrome
  end for
end for
for $j = r-1$ **downto** 1 **do** ▷ back-substitution
  for $i = 0$ **to** $j-1$ **do**
    $s_i \leftarrow s_i \oplus H_{ij}\, s_j$
  end for
end for

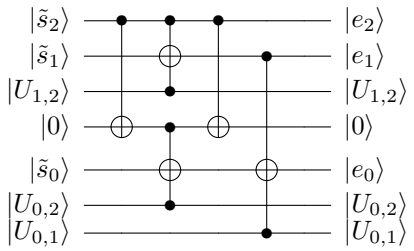


Figure 3: Closing back-substitution at the gate level for $r = 3$.

Table 2 compares the elimination costs with those of Perriello et al. [25] and Jang et al. [15][3]. The result of [15] includes concrete counts only for matrix sizes up to 48, and that of [25] uses a different Toffoli decomposition. We express all three in closed form[4]. We compare at the Toffoli level to assess the effect of the algorithmic optimization on its own, independent of the choice of Toffoli decomposition. For the parameters of target schemes, our circuit improves the qubit count by about 22% and the Toffoli count by about 20% over [15], while the Toffoli-depth increases only by the lower-order term $r-1$, which is negligible relative to the leading $\frac{1}{2}r^2$. Against [25] the Toffoli-depth is improved by about 67%: to leading order it is $\frac{1}{2}r^2$ for our circuit and for that of [15], against $\frac{3}{2}r^2$.

Table 2: Quantum resources for the Gauss-Jordan elimination compared with prior works [25,15].

Source	#Qubit*	#T*	T-depth*	#Toffoli*	Toffoli-depth*
P$^+$ [25]	–	–	–	$\frac{r(r-1)(5r+8)}{6}$	$\frac{3\,r(r-1)}{2}$
J$^+$ [15]	$\frac{9r^2-11r+6}{2}$	$\frac{7r(5r^2+3r-8)}{6}$	$2r^2+2r-4$	$\frac{r(5r^2+3r-8)}{6}$	$\frac{r^2+r-2}{2}$
This work✿	$\frac{7r^2-5r+4}{2}$ (22%)	$\frac{7r(2r^2+3r-5)}{3}$ (20%)	$2r^2+6r-8$ (−0.3%)	$\frac{r(2r^2+3r-5)}{3}$ (20%/20%)	$\frac{r^2+3r-4}{2}$ (67%/−0.3%)

✻ and ✽: Improvement percentages relative to J$^+$ [15] and P$^+$ [25]/J$^+$ [15], respectively (a negative value denotes a degradation).
✿: Our work includes the added syndrome back-substitution.

3.3 Hamming Weight Check

The last stage of the oracle turns $H_{\mathcal{S}}^{-1}s$ in the syndrome register into a marking decision: it computes the Hamming weight of this r-bit vector and flips the sign

[3] For [25], we do not use their reported full-depth numbers directly, because their estimates assign unit depth to high-level gates. Instead, we re-derived a network-topology Toffoli-layer depth from their QGJE schedule. Thus, the Perriello depth in Table 2 should be read as a normalized Toffoli-layer comparison, not as the full depth reported in [25].

[4] The full depth is not given in closed form, as it depends on how the Clifford and Toffoli gates are scheduled, and is reported only for explicitly generated circuits.

of the state when the weight equals t. Prior work computes this weight with the adder tree of [25, Section 2.6]. We instead accumulate all r bits in parallel with the tree-based quantum carry-save adder [18], which lowers the depth. This requires more ancillas than the adder tree, but we reuse the clean ancilla qubits from the elimination phase, achieving a lower T-depth without additional qubits.

4 Quantum Resource Estimates for ISD on HQC and Classic McEliece

In this section, we present quantum resource estimates for the ISD attack on the two schemes, obtained by combining the oracle of Section 3 with the Grover iteration count of [25]. We report logical costs, then derive physical qubits and runtime from them under a surface-code model, and finally assess both against the NIST security levels and the MAXDEPTH.

Table 3[5] reports the resulting logical quantum resources, which include the three oracle stages (i.e., sorting, elimination, and weight check). For HQC, the iteration count includes the decoding-one-out-of-many speedup that the quasi-cyclic structure permits [25, Section 3.1], lowering it by about $\sqrt{r}$. Classic McEliece, whose Goppa codes lack this structure, uses the plain single-syndrome count. The two schemes exhibit different cost trade-offs: HQC has larger N and r, which decide the matrix size, so its oracle uses more qubits, whereas its iteration count is smaller.

Another line of work [8] replaces Gauss–Jordan elimination with Wiedemann's sparse matrix inversion, minimizing the number of qubits instead of circuit depth and gates. As one illustrative point in this trade-off, we roughly compare with the space-optimized variant of Chevignard et al. [8, Table 1] (excluding the Grover iteration count). For instance, on Classic McEliece-348864 our circuit requires about 2^8 fewer Toffoli gates ($2^{30.26}$[6] against $2^{37.8}$ of their space-optimized version), while theirs requires about 2^8 fewer qubits ($2^{14.1}$ against our 2^{22}). Calculating the other parameter sets in the same way, we can observe a trade-off: the Gauss-Jordan line we improve lowers gates and depth, while the matrix inversion approach lowers qubits.

[5] Here, G denotes the total gate count after decomposition, including both Clifford and T.

[6] This is a per-iteration Toffoli count, not the total gate count G in Table 3. From the total T-count 2^{105} for QISD, we divide out the 2^{72} iterations and the seven T gates per Toffoli in our decomposition: $2^{105} \div 2^{72} \div 7 \approx 2^{30.2}$.

Table 3: Quantum resources for the ISD attack in base-2 logarithm.

Scheme	Iterations	#Qubit	#Total gates (G)	T-depth (T_d)	Full-depth (FD)	$G \times FD^*$
HQC-128	59	30	105	89	101	206
HQC-192	93	32	142	125	138	280
HQC-256	123	34	175	156	170	345
Classic McEliece-348864	72	22	106	93	100	206
Classic McEliece-460896	93	23	129	116	123	252
Classic McEliece-6688128	131	24	168	155	163	331
Classic McEliece-8192128	150	24	187	174	182	369

$*$: The gate-depth cost metric used to evaluate the NIST post-quantum security levels.

4.1 NIST Post-Quantum Security

NIST specifies security levels, each defined by the resources required to break a reference primitive (e.g., AES) [22]. This cost is measured as the product of the gate count and the full circuit depth (i.e., $G \times FD$), which NIST requires to be at least 2^{157}, 2^{221}, and 2^{285} at levels 1, 3, and 5 [22]. Table 3 shows that $G \times FD$ exceeds the threshold of the corresponding NIST security level for the listed instances, so each scheme satisfies its claimed security. The full depth exceeds the MAXDEPTH upper bound.

4.2 Physical Resources

We estimate the physical-qubit count and the runtime under the surface code [11], following the methodology of Amy et al. [2]. On the surface code each T-gate is supplied by a magic state, which we produce by Reed–Muller 15-to-1 distillation [7]. For the physical parameters, we assume a physical error rate $p_{phys} = 10^{-3}$, a per-gate error rate $p_g = 10^{-4}$, and code-cycle times of $1\,\mu$s, $10\,\mu$s, $100\,\mu$s, and 1 ms.

As an illustrative example, we derive the physical cost for HQC-128. Following [2, Algorithm 4] with p_{phys} and p_g, we obtain three distillation layers at code distances $\{35, 14, 7\}$ such that the output error rate satisfies $1/T_c$, where the total T-count (i.e., T_c) is 2^{104}. A logical qubit at code distance d is encoded into $2.5 \times 1.25\, d^2$ physical qubits. The bottom layer is at distance 7 and consists of $16 \times 15^2 = 3600$ input states. Each state is a logical qubit, so the bottom layer occupies $3600 \times (2.5 \times 1.25 \times 7^2) = 551250$ physical qubits and runs in $10 \times 7 = 70$ surface-code cycles. Computed in the same way, the second layer occupies 147000 physical qubits and runs in 140 cycles, while the top layer is completed in 350 cycles. A single distillation takes $70 + 140 + 350 = 560$ surface-code cycles, and one factory supplies $\lfloor 551250/147000 \rfloor = 3$ magic states. We need to supply $32768\,(= \lceil T_c/T_d \rceil)$ magic states within 560 cycles, so we parallelize across multiple factories. The required number of factories is $\lceil 32768/3 \rceil = 10923$, occupying $10923 \times 551250 \approx 2^{32.5}$ physical qubits.

We also account for the physical qubits of the Clifford gates. For the 2^{30} logical qubits, we take the smallest code distance d_c with $(p_{phys}/0.0125)^{(d_c+1)/2} < 1/C_g$ for $C_g = 2^{105}$ Clifford gates, giving $d_c = 57$. The Clifford part then requires $2^{30} \times 2.5 \times 1.25 \times 57^2 \approx 2^{43.3}$ physical qubits. Here, the Clifford gates are executed within the magic-state distillation time, so they do not affect the runtime. For HQC-128, with 10923 factories supplying the magic states, the runtime is $2^{89} \times 560 \times 1\,\mu\text{s} \approx 2^{53}$ years with about 2^{43} physical qubits[7]. We compute the other parameter sets in the same manner and report all results in Table 4.

Table 4: Physical resource estimates for QISD using surface codes (in base-2 logarithm).

Scheme	#Physical qubits	Runtime for code cycle time			
		1 μs	10 μs	100 μs	1 ms
HQC-128	43	53	57	60	63
HQC-192	46	90	93	96	100
HQC-256	49	121	124	128	131
Classic McEliece-348864	35	57	61	64	67
Classic McEliece-460896	37	81	84	87	91
Classic McEliece-6688128	39	120	123	127	130
Classic McEliece-8192128	39	139	142	146	149

5 Conclusion

Post-quantum cryptography requires concrete resource estimates to assess its security against quantum attacks. This paper presents an improved quantum circuit for the Gauss–Jordan elimination that dominates an information-set decoding oracle, and estimates the cost of Grover-based ISD attacks on HQC and Classic McEliece.

At the elimination stage, our circuit updates only the matrix entries used by the remaining pivots and the syndrome entries needed for the final weight test. It then recovers the result vector by a parallel back-substitution on the syndrome register. Compared with [15], this reduces the qubit count by about 22%. The Toffoli count is reduced by up to about 20% compared with both [15] and [25]. The Toffoli depth remains close to that of [15], while it is reduced by about 67% compared with [25] under the normalized Toffoli-layer comparison used in Table 2.

Combining our oracle with the Grover iteration counts from [25], we report logical costs for the evaluated HQC and Classic McEliece instances in Table 3. Under the $G \times FD$ metric, these instances exceed the thresholds for the corre-

[7] The physical qubits for the Clifford part dominate the total.

sponding NIST post-quantum security levels. Separately, the reported full depths exceed the MAXDEPTH bounds.

For fault-tolerant resources, we estimate the physical qubit count and runtime under a surface-code model with magic-state distillation (Table 4). As an example, HQC-128 requires about 2^{43} physical qubits and about 2^{53} years at a $1\,\mu$s code cycle under a physical error rate $p_{phys} = 10^{-3}$.

Considering more advanced ISD variants and reassessing the physical cost under lower-overhead fault-tolerant architectures remain possible directions for future work.

References

1. Aguilar Melchor, C., Aragon, N., Bettaieb, S., Bidoux, L., Blazy, O., Deneuville, J.C., Gaborit, P., Persichetti, E., Zémor, G., Bourges, I.: Hamming quasi-cyclic (HQC). NIST Post-Quantum Cryptography Standardization submission (2018), https://pqc-hqc.org 1, 3
2. Amy, M., Di Matteo, O., Gheorghiu, V., Mosca, M., Parent, A., Schanck, J.: Estimating the cost of generic quantum pre-image attacks on SHA-2 and SHA-3. In: Avanzi, R., Heys, H. (eds.) Selected Areas in Cryptography – SAC 2016. pp. 317–337. Springer International Publishing, Cham (2017) 9
3. Amy, M., Maslov, D., Mosca, M., Roetteler, M.: A meet-in-the-middle algorithm for fast synthesis of depth-optimal quantum circuits. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems **32**(6), 818–830 (Jun 2013). https://doi.org/10.1109/tcad.2013.2244643, http://dx.doi.org/10.1109/TCAD.2013.2244643 4
4. Berlekamp, E.R., McEliece, R.J., van Tilborg, H.C.A.: On the inherent intractability of certain coding problems. IEEE Transactions on Information Theory **24**(3), 384–386 (1978) 2
5. Bernstein, D.J.: Grover vs. mceliece. In: Post-Quantum Cryptography: Third International Workshop, PQCrypto 2010, Darmstadt, Germany, May 25-28, 2010. Proceedings 3. pp. 73–80. Springer (2010) 2, 3
6. Bernstein, D.J., Chou, T., Lange, T., von Maurich, I., Misoczki, R., Niederhagen, R., Persichetti, E., Peters, C., Schwabe, P., Sendrier, N., Szefer, J., Wang, W.: Classic McEliece: conservative code-based cryptography. NIST Post-Quantum Cryptography Standardization submission (2017), https://classic.mceliece.org 1, 3
7. Bravyi, S., Kitaev, A.: Universal quantum computation with ideal clifford gates and noisy ancillas. Physical Review A—Atomic, Molecular, and Optical Physics **71**(2), 022316 (2005) 9
8. Chevignard, C., Fouque, P.A., Schrottenloher, A.: Reducing the number of qubits in quantum information set decoding. In: Advances in Cryptology – ASIACRYPT 2024. Lecture Notes in Computer Science, vol. 15491, pp. 299–329. Springer (2025) 2, 8
9. Classic McEliece team: Classic McEliece: round-4 specification. https://classic.mceliece.org/mceliece-spec-20221023.pdf (2022) 3
10. Esser, A., Ramos-Calderer, S., Bellini, E., Latorre, J.I., Manzano, M.: An optimized quantum implementation of ISD on scalable quantum resources. Cryptology ePrint Archive, Paper 2021/1608 (2021), https://eprint.iacr.org/2021/1608 3

11. Fowler, A.G., Mariantoni, M., Martinis, J.M., Cleland, A.N.: Surface codes: Towards practical large-scale quantum computation. Physical Review A—Atomic, Molecular, and Optical Physics **86**(3), 032324 (2012) 9
12. Gidney, C.: Halving the cost of quantum addition. Quantum **2**, 74 (2018) 4
13. 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) 3
14. HQC team: HQC: Hamming quasi-cyclic — updated specification. https://pqc-hqc.org/doc/hqc_specifications_2025_08_22.pdf (2025) 3
15. Jang, K., Oh, Y., Seo, H.: Depth-optimized quantum circuit of Gauss–Jordan elimination. Applied Sciences **14**(19), 8579 (2024) 1, 2, 5, 7, 10
16. Jaques, S., Naehrig, M., Roetteler, M., Virdia, F.: Implementing grover oracles for quantum key search on aes and lowmc. In: Advances in Cryptology–EUROCRYPT 2020: 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10–14, 2020, Proceedings, Part II 30. pp. 280–310. Springer (2020) 4
17. Kachigar, G., Tillich, J.P.: Quantum information set decoding algorithms. In: Post-Quantum Cryptography (PQCrypto 2017). LNCS, vol. 10346, pp. 69–89. Springer (2017) 3
18. Kim, H., Lim, S., Jang, K., Wang, S., Baksi, A., Chattopadhyay, A., Seo, H.: Tree-based quantum carry-save adder. Cryptology ePrint Archive, Paper 2025/626 (2025), https://eprint.iacr.org/2025/626 8
19. Kirshanova, E.: Improved quantum information set decoding. In: Post-Quantum Cryptography (PQCrypto 2018). LNCS, vol. 10786, pp. 507–527. Springer (2018) 3
20. Kliuchnikov, V., Maslov, D., Mosca, M.: Practical approximation of single-qubit unitaries by single-qubit quantum Clifford and T circuits. IEEE Transactions on Computers **65**(1), 161–172 (2016) 4
21. Mukherjee, C.S., Maitra, S., Gaurav, V., Roy, D.: Preparing Dicke states on a quantum computer. IEEE Transactions on Quantum Engineering **1**, 1–17 (2020) 4, 5
22. National Institute of Standards and Technology: Submission requirements and evaluation criteria for the post-quantum cryptography standardization process. https://csrc.nist.gov/CSRC/media/Projects/Post-Quantum-Cryptography/documents/call-for-proposals-final-dec-2016.pdf (2016) 9
23. National Institute of Standards and Technology: HQC announced as a fourth round selection for standardization. https://csrc.nist.gov/news/2025/hqc-announced-as-a-4th-round-selection (2025) 1
24. Perriello, S., Barenghi, A., Pelosi, G.: A complete quantum circuit to solve the information set decoding problem. In: 2021 IEEE International Conference on Quantum Computing and Engineering (QCE). pp. 366–377. IEEE (2021) 2
25. Perriello, S., Barenghi, A., Pelosi, G.: Improving the efficiency of quantum circuits for information set decoding. ACM Transactions on Quantum Computing **4**(4), 1–40 (2023) 1, 2, 3, 4, 5, 7, 8, 10
26. Prange, E.: The use of information sets in decoding cyclic codes. IRE Transactions on Information Theory **8**(5), 5–9 (1962) 2, 3