Improved Classical and Quantum Algorithms for Subset-Sum

Xavier Bonnetain[1], Rémi Bricout[2,3], André Schrottenloher[3], and Yixin Shen[4]

[1] Institute for Quantum Computing, Department of Combinatorics and Optimization, University of Waterloo, Waterloo, ON, Canada
[2] Sorbonne Université, Collège Doctoral, F-75005 Paris, France
[3] Inria, Paris, France
[4] Université de Paris, IRIF, CNRS, F-75006 Paris, France

Abstract. We present new classical and quantum algorithms for solving random subset-sum instances. First, we improve over the Becker-Coron-Joux algorithm (EUROCRYPT 2011) from $\widetilde{\mathcal{O}}\left(2^{0.291n}\right)$ down to $\widetilde{\mathcal{O}}\left(2^{0.283n}\right)$, using more general representations with values in $\{-1, 0, 1, 2\}$.
Next, we improve the state of the art of quantum algorithms for this problem in several directions. By combining the Howgrave-Graham-Joux algorithm (EUROCRYPT 2010) and quantum search, we devise an algorithm with asymptotic running time $\widetilde{\mathcal{O}}\left(2^{0.236n}\right)$, lower than the cost of the quantum walk based on the same classical algorithm proposed by Bernstein, Jeffery, Lange and Meurer (PQCRYPTO 2013). This algorithm has the advantage of using *classical* memory with quantum random access, while the previously known algorithms used the quantum walk framework, and required *quantum* memory with quantum random access.
We also propose new quantum walks for subset-sum, performing better than the previous best time complexity of $\widetilde{\mathcal{O}}\left(2^{0.226n}\right)$ given by Helm and May (TQC 2018). We combine our new techniques to reach a time $\widetilde{\mathcal{O}}\left(2^{0.216n}\right)$. This time is dependent on a heuristic on quantum walk updates, formalized by Helm and May, that is also required by the previous algorithms. We show how to partially overcome this heuristic, and we obtain an algorithm with quantum time $\widetilde{\mathcal{O}}\left(2^{0.218n}\right)$ requiring only the standard classical subset-sum heuristics.

Keywords: subset-sum, representation technique, quantum search, quantum walk, list merging.

1 Introduction

We study the *subset-sum problem*, also known as *knapsack problem*: given n integers $\mathbf{a} = (a_1, \dots a_n)$, and a target integer S, find an n-bit vector $\mathbf{e} = (e_1, \dots e_n) \in \{0, 1\}^n$ such that $\mathbf{e} \cdot \mathbf{a} = \sum_i e_i a_i = S$. The *density* of the knapsack instance is defined as $d = n/(\log_2 \max_i a_i)$, and for a random instance $\mathbf{a}$, it is related to the number of solutions that one can expect.

The decision version of the knapsack problem is NP-complete [16]. Although certain densities admit efficient algorithms, related to lattice reduction [27,28], the best algorithms known for the knapsack problem when the density is close to 1 are exponential-time, which is why we name these instances "hard" knapsacks. This problem underlies some cryptographic schemes aiming at post-quantum security (see *e.g.* [29]), and is used as a building block in some quantum hidden shift algorithms [7], which have some applications in quantum cryptanalysis of isogeny-based [11] and symmetric cryptographic schemes [9].

In this paper, we focus on the case where $d = 1$, where expectedly a single solution exists. Instead of naively looking for the solution $\mathbf{e}$ via exhaustive search, in time 2^n, Horowitz and Sahni [20] proposed to use a meet-in-the-middle approach in $2^{n/2}$ time and memory. The idea is to find a collision between two lists of $2^{n/2}$ subknapsacks, *i.e.* to merge these two lists for a single solution. Schroeppel and Shamir [37] later improved this to a 4-list merge, in which the memory complexity can be reduced down to $2^{n/4}$.

The Representation Technique. At EUROCRYPT 2010, Howgrave-Graham and Joux [21] (HGJ) proposed a heuristic algorithm solving *random* subset-sum instances in time $\widetilde{\mathcal{O}}\left(2^{0.337n}\right)$, thereby breaking the $2^{n/2}$ bound. Their key idea was to represent the knapsack solution ambiguously as a sum of vectors in $\{0,1\}^n$. This *representation technique* increases the search space size, allowing to merge more lists, with new arbitrary constraints, thereby allowing for a more time-efficient algorithm. The time complexity exponent is obtained by numerical optimization of the list sizes and constraints, assuming that the individual elements obtained in the merging steps are well-distributed. This is the standard heuristic of classical and quantum subset-sum algorithms. Later, Becker, Coron and Joux [3] (BCJ) improved the asymptotic runtime down to $\widetilde{\mathcal{O}}\left(2^{0.291n}\right)$ by allowing even more representations, with vectors in $\{-1,0,1\}^n$.

The BCJ representation technique is not only a tool for subset-sums, as it has been used to speed up generic decoding algorithms, classically [31,4,32] and quantumly [22]. Therefore, the subset-sum problem serves as the simplest application of representations, and improving our understanding of the classical and quantum algorithms may have consequences on these other generic problems.

Quantum Algorithms for the Subset-Sum Problem. Cryptosystems based on hard subset-sums are natural candidates for post-quantum cryptography, but to understand precisely their security, we have to study the best generic algorithms for solving subset-sums. The first quantum time speedup for this problem was obtained in [6], with a quantum time $\widetilde{\mathcal{O}}\left(2^{0.241n}\right)$. The algorithm was based on the HGJ algorithm. Later on, [18] devised an algorithm based on BCJ, running in time $\widetilde{\mathcal{O}}\left(2^{0.226n}\right)$. Both algorithms use the corresponding classical merging structure, wrapped in a quantum walk on a Johnson graph, in the MNRS quantum walk framework [30]. However, they suffer from two limitations.

First, both use the model of *quantum memory with quantum random-access* (QRAQM), which is stronger than the standard quantum circuit model, as it allows unit-time lookups in superposition of all the qubits in the circuit. The

QRAQM model is used in most quantum walk algorithms to date, but its practical realizations are still unclear. With a more restrictive model, i.e. *classical memory with quantum random-access* (QRACM), no quantum time speedup over BCJ was previously known. This is not the case for some other hard problems in post-quantum cryptography, *e.g.* heuristic lattice sieving for the Shortest Vector Problem, where the best quantum algorithms to date require only QRACM [25].

Second, both use a conjecture (implicit in [6], made explicit in [18]) about quantum walk updates. In short, the quantum walk maintains a data structure, that contains a merging tree similar to HGJ (resp. BCJ), with lists of smaller size. A quantum walk step is made of updates that changes an element in the lowest-level lists, and requires to modify the upper levels accordingly, *i.e.* to track the partial collisions that must be removed or added. In order to be efficient, the update needs to run in polynomial time. Moreover, the resulting data structure shall be a function of the lowest-level list, and not depend on the path taken in the walk. The conjecture states that it should be possible to guarantee sound updates without impacting the time complexity exponent. However, it does not seem an easy task and the current literature on subset-sums lacks further justification or workarounds.

Contributions. In this paper, we improve classical and quantum subset-sum algorithms based on representations. We write these algorithms as sequences of "merge-and-filter" operations, where lists of subknapsacks are first merged with respect to an arbitrary constraint, then *filtered* to remove the subknapsacks that cannot be part of a solution.

First, we propose a more time-efficient classical subset-sum algorithm based on representations. We have two classical improvements: we revisit the previous algorithms and show that some of the constraints they enforced were not needed, and we use more general distributions by allowing "2"s in the representations. Overall, we obtain a better time complexity exponent of 0.283.

Most of our contributions concern quantum algorithms. As a generic tool, we introduce *quantum filtering*, which speeds up the filtering of representations with a quantum search. We use this improvement in all our new quantum algorithms.

We give an improved quantum walk based on quantum filtering and our extended $\{-1, 0, 1, 2\}$ representations. Our best runtime exponent is 0.216, under the quantum walk update heuristic of [18]. Next, we show how to overcome this heuristic, by designing a new data structure for the vertices in the quantum walk, and a new update procedure with guaranteed time. We remove this heuristic from the previous algorithms [6,18] with no additional cost. However, we find that removing it from our quantum walk increases its cost to 0.218.

In a different direction, we devise a new quantum subset-sum algorithm based on HGJ, with time $\widetilde{\mathcal{O}}\left(2^{0.236n}\right)$. It is the first quantum time speedup on subset-sums that is *not* based on a quantum walk. The algorithm performs instead a depth-first traversal of the HGJ tree, using quantum search as its only building block. Hence, by construction, it does not require the additional heuristic of [18] *and* it only uses *classical memory with quantum random-access*, giving also the first quantum time speedup for subset-sum in this memory model.

A summary of our contributions is given in Table 1[5]. All these complexity exponents are obtained by numerical optimization. Our code is available at `https://github.com/xbonnetain/optimization-subset-sum`.

Table 1. Previous and **new** algorithms for subset-sum, classical and quantum, with time and memory exponents rounded upwards. We note that the removal of Heuristic 2 in [6,18] comes from our new analysis in Section 6.4. QW: Quantum Walk. QS: Quantum Search. CF: Constraint filtering (not studied in this paper). QF: Quantum filtering.

Time exp.	Memory exp.	Representations	Memory model	Techniques	Requires Heur. 2	Reference
			Classical			
0.3370	0.3113	$\{0,1\}$	RAM			[21]
0.2909	0.2909	$\{-1,0,1\}$	RAM			[3]
0.287		$\{-1,0,1\}$	RAM	CF		[36]
0.2830	**0.2830**	$\{-1,0,1,2\}$	RAM			Sec. 2.5
			Quantum			
0.241	0.241	$\{0,1\}$	QRAQM	QW	**No**	[6] + Sec. 6.4
0.226	0.226	$\{-1,0,1\}$	QRAQM	QW	**No**	[18] + Sec. 6.4
0.2356	**0.2356**	$\{0,1\}$	**QRACM**	QS + QF	No	Sec 4.3
0.2156	**0.2110**	$\{-1,0,1,2\}$	QRAQM	QW + QF	Yes	Sec. 5.3
0.2182	**0.2182**	$\{-1,0,1,2\}$	QRAQM	QW + QF	No	Sec. 6.4

Outline. In Section 2, we study classical algorithms. We review the representation technique, the HGJ algorithm and introduce our new $\{-1,0,1,2\}$ representations to improve over [3]. In Section 3, we move to the quantum setting, introduce some preliminaries and the previous quantum algorithms for subset-sum. In Section 4, we present and study our new quantum algorithm based on HGJ and quantum search. We give different optimizations and time-memory tradeoffs. In Section 5, we present our new quantum algorithm based on a quantum walk. Finally, in Section 6 we show how to overcome the quantum walk update conjecture, up to a potential increase in the update cost. We conclude, and give a summary of our new results in Section 7.

2 List Merging and Classical Subset-sum Algorithms

In this section, we remain in the classical realm. We introduce the standard subset-sum notations and heuristics and give a new presentation of the HGJ

[5] After this work, Alexander May has informed us that the thesis [14] contains unpublished results using more symbols, with the best exponent of 0.2871 obtained with the symbol set $\{-2,-1,0,1,2\}$.

algorithm, putting an emphasis on the *merge-and-filter* operation. We introduce our extended $\{-1,0,1,2\}$ representations and detail our improvements over BCJ.

2.1 Notations and Conventions

Hereafter and in the rest of the paper, all time and memory complexities, classical and quantum, are exponential in n. We use the soft-O notation $\widetilde{\mathcal{O}}$ which removes polynomial factors in n, and focus on the asymptotic exponent, relative to n. We use $\text{negl}(n)$ for any function that vanishes inverse-exponentially in n. We often replace asymptotic exponential time and memory complexities (*e.g.* $\widetilde{\mathcal{O}}(2^{\alpha n})$) by their exponents (*e.g.* α). We use capital letters (*e.g.* L) and corresponding letters (*e.g.* ℓ) to denote the same value, in $\log_2$ and relatively to n: $\ell = \log_2(L)/n$.

Definition 1 (Entropies and multinomial functions). *We define the following functions:*

Hamming Entropy: $h(x) = -x\log_2 x - (1-x)\log_2(1-x)$

Binomial: $\text{bin}(\omega,\alpha) = h(\alpha/\omega)\omega$

2-way Entropy: $g(x,y) = -x\log_2 x - y\log_2 y - (1-x-y)\log_2(1-x-y)$

Trinomial: $\text{trin}(\omega,\alpha,\beta) = g(\alpha/\omega,\beta/\omega)\omega$

3-way Entropy: $f(x,y,z) = -x\log_2 x - y\log_2 y - z\log_2 z - (1-x-y-z)\log_2(1-x-y-z)$

Quadrinomial: $\text{quadrin}(\omega,\alpha,\beta,\gamma) = f(\alpha/\omega,\beta/\omega,\gamma/\omega)\omega$

Property 1 (Standard approximations). We have the following approximations, asymptotically in n:

$$\text{bin}(\omega,\alpha) \simeq \frac{1}{n}\log_2\binom{\omega n}{\alpha n} \;\; ; \;\; \text{trin}(\omega,\alpha,\beta) \simeq \frac{1}{n}\log_2\binom{\omega n}{\alpha n,\beta n}$$
$$\text{quadrin}(\omega,\alpha,\beta,\gamma) \simeq \frac{1}{n}\log_2\binom{\omega n}{\alpha n,\beta n,\gamma n}$$

Definition 2 (Distributions of knapsacks). *A* knapsack *or* subknapsack *is a vector* $\mathbf{e} \in \{-1,0,1,2\}^n$*. The set of* $\mathbf{e}$ *with* αn *"-1",* $(\alpha+\beta-2\gamma)n$ *"1",* γn *"2" and* $(1-2\alpha-\beta+\gamma)n$ *"0" is denoted* $D^n[\alpha,\beta,\gamma]$*. If* $\gamma = 0$*, we may omit the third parameter. This coincides with the notation* $D^n[\alpha,\beta]$ *from [18].*

Note that we always add vectors *over the integers*, and thus, the sum of two vectors of $D^n[*,*,*]$ may contain unwanted symbols $-2, 3$ or 4.

Property 2 (Size of knapsack sets). We have:

$$\frac{1}{n}\log_2|D^n[0,\beta,0]| \simeq h(\beta) \;\; ; \;\; \frac{1}{n}\log_2|D^n[\alpha,\beta,0]| \simeq g(\alpha,\alpha+\beta)$$

$$\frac{1}{n}\log_2|D^n[\alpha,\beta,\gamma]| \simeq f(\alpha,\alpha+\beta-2\gamma,\gamma) \;\; .$$

Subset-sum. The problem we will solve is defined as follows:

Definition 3 (Random subset-sum instance of weight $n/2$). *Let* $\mathbf{a}$ *be chosen uniformly at random from* $(\mathbb{Z}_N)^n$*, where* $N \simeq 2^n$*. Let* $\mathbf{e}$ *be chosen uniformly at random from* $D^n[0,1/2,0]$*. Let* $t = \mathbf{a}\cdot\mathbf{e} \pmod N$*. Then* $(\mathbf{a},t)$ *is a* random subset-sum instance*. A* solution *is a vector* $\mathbf{e}'$ *such that* $\mathbf{a}\cdot\mathbf{e}' = t \pmod N$*.*

Sampling. Throughout this paper, we assume that we can classically sample uniformly at random from $D^n[\alpha, \beta, \gamma]$ in time poly(n). (Since αn, βn and γn will in general not be integer, we suppose to have them rounded to the nearest integer.) This comes from an efficient bijection between representations and integers (see Appendix A in the full version of the paper [8]). In addition, we can efficiently produce the uniform superposition of vectors of $D^n[\alpha, \beta, \gamma]$, using poly($n$) quantum gates, and we can perform a quantum search among representations.

2.2 Merging and Filtering

In all subset-sum algorithms studied in this paper, we repeatedly sample vectors with certain distributions $D^n[*, *, *]$, then combine them. Let $D_1 = D^n[\alpha_1, \beta_1, \gamma_1]$, $D_2 = D^n[\alpha_2, \beta_2, \gamma_2]$ be two input distributions and $D = D^n[\alpha, \beta, \gamma]$ be a target. Given two lists $L_1 \in D_1^{|L_1|}$ and $L_2 \in D_2^{|L_2|}$, we define:

- the *merged list* $L = L_1 \bowtie_c L_2$ containing all vectors $\mathbf{e} = \mathbf{e_1} + \mathbf{e_2}$ such that: $\mathbf{e_1} \in L_1, \mathbf{e_2} \in L_2$, $(\mathbf{e}_1 + \mathbf{e}_2) \cdot \mathbf{a} = s \mod M$, $s \leq M$ is an arbitrary integer and $M \approx 2^{cn}$ (we write $L_1 \bowtie_c L_2$ because s is an arbitrary value, whose choice is without incidence on the algorithm)
- the *filtered list* $L^f = (L \cap D) \subseteq L$, containing the vectors with the target distribution of $1, -1, 2$ (the target D will always be clear from context).

In general, L is exponentially bigger than L^f and does not need to be written down, as vectors can be filtered on the fly. The algorithms then repeat the merge-and-filter operation on multiple levels, moving towards the distribution $D^n[0, 1/2]$ while increasing the bit-length of the modular constraint, until we satisfy $\mathbf{e} \cdot \mathbf{a} = t \mod 2^n$ and obtain a solution. Note that this merging-and-filtering view that we adopt, where the merged list is repeatedly sampled before an element passes the filter, has some similarities with the ideas developed in the withdrawn article [15].

The standard subset-sum heuristic assumes that vectors in L^f are drawn independently, uniformly at random from D. It simplifies the complexity analysis of both classical and quantum algorithms studied in this paper. Note that this heuristic, which is backed by experiments, actually leads to provable probabilistic algorithms in the classical setting (see [3, Theorem 2]). We adopt the version of [18].

Heuristic 1 *If input vectors are uniformly distributed in $D_1 \times D_2$, then the filtered pairs are uniformly distributed in D (more precisely, among the subset of vectors in D satisfying the modular condition).*

Filtering Representations. We note $\ell = (1/n)\log_2|L|$, and so on for ℓ_1, ℓ_2, ℓ^f. By Heuristic 1, the average sizes of L_1, L_2, L and L^f are related by:

- $\ell = \ell_1 + \ell_2 - c$
- $\ell^f = \ell + \mathsf{pf}$, where pf is negative and $2^{\mathsf{pf} n}$ is the probability that a pair $(\mathbf{e_1}, \mathbf{e_2})$, drawn uniformly at random from $D_1 \times D_2$, has $(\mathbf{e_1} + \mathbf{e_2}) \in D$.

In particular, the occurrence of collisions in L^f is a negligible phenomenon, unless ℓ^f approaches $(\log_2 |D|/n) - c$, which is the maximum number of vectors in D with constraint c. For a given random knapsack problem, with high probability, the size of any list built by sampling, merging and filtering remains very close to its average (by a Chernoff bound and a union bound on all lists).

Here, pf depends only on D_1, D_2 and D. Working with this *filtering probability* is especially useful for writing down our algorithm in Section 4. We give its formula for $\{0,1\}$ representations below. Two similar results for $\{-1,0,1\}$ and $\{-1,0,1,2\}$ can be found in the full version of the paper [8].

Lemma 1 (Filtering HGJ-style representations). *Let* $\mathbf{e_1} \in D^n[0,\alpha]$ *and* $\mathbf{e_2} \in D^n[0,\beta]$ *be drawn uniformly at random. The probability that* $\mathbf{e_1} + \mathbf{e_2} \in D^n[0,\alpha+\beta]$ *is* 0 *if* $\alpha + \beta > 1$*, and* $2^{\mathsf{pf}_1(\alpha,\beta)n}$ *otherwise, with*

$$\mathsf{pf}_1(\alpha,\beta) = \text{bin}(1-\alpha,\beta) - h(\beta) = \text{bin}(1-\beta,\alpha) - h(\alpha) \ .$$

Proof. The probability that a $\mathbf{e}_1 + \mathbf{e}_2$ survives the filtering is:

$$\binom{n-\alpha n}{\beta n} / \binom{n}{\beta n} = \binom{n-\beta n}{\alpha n} / \binom{n}{\alpha n} \ .$$

Indeed, given a choice of αn bit positions among n, the other βn bit positions must be compatible, hence chosen among the $(1-\alpha)n$ remaining positions. By taking the $\log_2$, we obtain the formula for the filtering probability. □

Time Complexity of Merging. Classically, the time complexity of the merge-and-filter operation is related to the size of the *merged list*.

Lemma 2 (Classical merging with filtering). *Let* L_1 *and* L_2 *be two sorted lists stored in classical memory with random access. In* $\log_2$*, relatively to* n*, and discarding logarithmic factors, merging and filtering* L_1 *and* L_2 *costs a time* $\max(\min(\ell_1,\ell_2), \ell_1+\ell_2-c)$ *and memory* $\max(\ell_1,\ell_2,\ell^f)$*, assuming that we must store the filtered output list.*

Proof. Assuming sorted lists, there are two symmetric ways to produce a stream of elements of $L_1 \bowtie_c L_2$: we can go through the elements of L_1, and for each one, find the matching elements in L_2 by dichotomy search (time $\ell_1 + \max(0, \ell_2 - c)$) or we can exchange the role of L_1 and L_2. Although we do not need to store $L_1 \bowtie_c L_2$, we need to examine all its elements in order to filter them. □

2.3 Correctness of the Algorithms

While the operation of *merging and filtering* is the same as in previous works, our complexity analysis differs [21,3,6,18]. We enforce the constraint that the final list contains a single solution, hence if it is of size $2^{n\ell_0}$, we constrain $\ell_0 = 0$. Next, we limit the sizes of the lists so that they do not contain duplicate vectors: these are *saturation constraints*. A list of size $2^{n\ell}$, of vectors sampled from a distribution

D, with a constraint of cn bits, has the constraint: $\ell \leq \frac{1}{n}\log_2|D| - c$. This says that there are not more than $|D|/2^{cn}$ vectors $\mathbf{e}$ such that $\mathbf{e}\cdot\mathbf{a} = r \pmod{2^{cn}}$ for the (randomly chosen) arbitrary constraint r.

Previous works focus on the solution vector $\mathbf{e}$ and compute the number of *representations* of $\mathbf{e}$, that is, the number of ways it can be decomposed as a sum: $\mathbf{e} = \mathbf{e}_1 + \ldots + \mathbf{e}_t$ of vectors satisfying the constraints on the distributions. Then, they compare this with the probability that a given representation passes the arbitrary constraints imposed by the algorithm. As their lists contains all the subknapsacks that fulfill the constraint, this really reflects the number of duplicates, and it suffices to enforce that the number of representations is equal to the inverse probability that a representation fulfills the constraint. If the two lists we merge are not of maximal size, the size of the merged list is the number of elements that fulfill the corresponding distribution times the probability that such an element is effectively the sum of two elements in the initial lists.

The two approaches are strictly equivalent, as the probability that the sum of two subknapsacks is valid is exactly the number of representations of the sum, divided by the number of pairs of subknapsacks.

2.4 The HGJ Algorithm

We start our study of classical subset-sum by recalling the algorithm of Howgrave-Graham and Joux [21], with the corrected time complexity of [3]. The algorithm builds a merging tree of lists of subknapsacks, with four levels, numbered 3 down to 0. Level j contains 2^j lists. In total, 8 lists are merged together into one.

Level 3. We build 8 lists denoted $L_0^3 \ldots L_7^3$. They contain *all* subknapsacks of weight $\frac{n}{16}$ on $\frac{n}{2}$ bits, either left or right:

$$\begin{cases} L_{2i}^3 &= D^{n/2}[0,1/8] \times \{0^{n/2}\} \\ L_{2i+1}^3 &= \{0^{n/2}\} \times D^{n/2}[0,1/8] \end{cases}$$

From Property 2, these level-3 lists have size $\ell_3 = h(1/8)/2$. As the positions set to 1 cannot interfere, these is no filtering when merging L_{2i}^3 and L_{2i+1}^3.

Level 2. We merge the lists pairwise with a (random) constraint on $c_2 n$ bits, and obtain 4 filtered lists. The size of the filtered lists plays a role in the memory complexity of the algorithm, but the time complexity depends on the size of the unfiltered lists.

In practice, when we say "with a constraint on $c_j n$ bits", we assume that given the subset-sum objective t modulo 2^n, random values r_i^j such that $\sum_i r_i^j = t \mod 2^{c_j n}$ are selected at level j, and the r_i^j have $c_j n$ bits only. Hence, at this step, we have selected 4 integers on $c_2 n$ bits $r_0^1, r_1^1, r_2^1, r_3^1$ such that $r_0^1 + r_1^1 + r_2^1 + r_3^1 = t \mod 2^{c_2 n}$. The 4 level-2 lists $L_0^2, L_1^2, L_2^2, L_3^2$ have size $\ell_2 = (h(1/8) - c_2)$, they contain subknapsacks of weight $\frac{n}{8}$ on n bits.

Remark 1. The precise values of these r_i are irrelevant, since they cancel out each other in the end. They are selected at random during a run of the algorithm, and although there could be "bad" values of them that affect significantly the computation, this is not expected to happen.

Level 1. We merge the lists pairwise with $(c_1 - c_2)n$ new bits of constraint, ensuring that the constraint is compatible with the previous ones. We obtain two filtered lists L_0^1, L_1^1, containing subknapsacks of weight $n/4$. They have size:

$$\ell_1 = 2\ell_2 - (c_1 - c_2) + \mathsf{pf}_1\left(1/8, 1/8\right)$$

where $\mathsf{pf}_1\left(1/8, 1/8\right)$ is given by Lemma 1.

Level 0. We find a solution to the subset-sum problem with the complete constraint on n bits. This means that the list L^0 must have expected length $\ell_0 = 0$. Note that there remains $(1 - c_1)n$ bits of constraint to satisfy, and the filtering term is similar as before, so:

$$\ell_0 = 2\ell_1 - (1 - c_1) + \mathsf{pf}_1\left(1/4, 1/4\right) \quad .$$

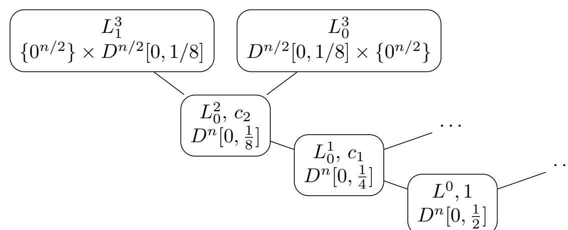

Fig. 1. The HGJ algorithm (duplicate lists are omitted)

By Lemma 2, the time complexity of this algorithm is determined by the sizes of the unfiltered lists: $\max\left(\ell_3, 2\ell_3 - c_2, 2\ell_2 - (c_1 - c_2), 2\ell_1 - (1 - c_1)\right)$. The memory complexity depends of the sizes of the filtered lists: $\max\left(\ell_3, \ell_2, \ell_1\right)$. By a numerical optimization, one obtains a time exponent of $0.337n$.

2.5 The BCJ Algorithm and our improvements

The HGJ algorithm uses representations to increase artificially the search space. The algorithm of Becker, Coron and Joux [3] improves the runtime exponent down to 0.291 by allowing even more freedom in the representations, which can now contain "-1"s. The "-1"s have to cancel out progressively, to ensure the validity of the final knapsack solution.

We improve over this algorithm in two different ways. First, we relax the constraints $\ell_j + c_j = g(\alpha_j, 1/2^{j+1})$ enforced in [3], as only the inequalities $\ell_j + c_j \leqslant g(\alpha_j, 1/2^{j+1})$ are necessary: they make sure the lists are not larger than the number of distinct elements they can contain. This idea was also implicitly used in [13], in the context of syndrome decoding. When optimizing the parameters under these new constraints, we bring the asymptotic time exponent down to $0.289n$.

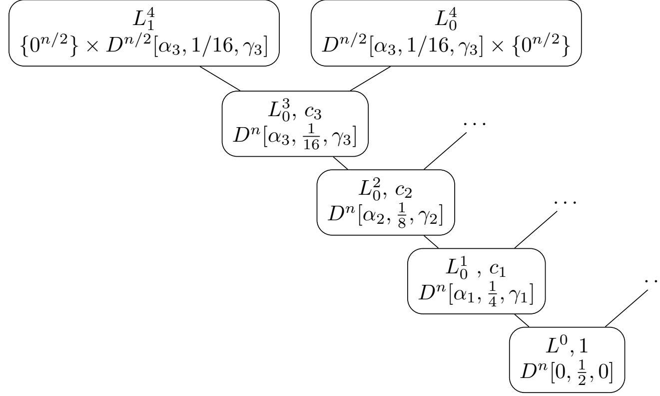

Fig. 2. Our improved algorithm (duplicate lists are omitted).

$\{-1, 0, 1, 2\}$ *representations.* Next, we allow the value "2" in the subknapsacks. This allows us to have more representations for the final solution from the same initial distributions. Indeed, in BCJ, if on a bit the solution is the sum of a "-1" and two "1"s, then it can only pass the merging steps if we first have the "-1" that cancels a "1", and then the addition of the second "1". When allowing "2"s, we can have the sum of the two "1's and then at a later step the addition of a "-1". The algorithm builds a merging tree with five levels, numbered 4 down to 0. Level j contains 2^j lists. In total, 16 lists are merged together into one.

Level 4. We build 16 lists $L_0^4 \dots L_{15}^4$. They contain complete distributions on $\frac{n}{2}$ bits, either left or right, with $\frac{n}{32} + \frac{\alpha_3 n}{2} - \gamma_3 n$ "1", $\frac{\alpha_3 n}{2}$ "-1" and $\frac{\gamma_3 n}{2}$ "2":

$$\begin{cases} L_{2i}^4 &= D^{n/2}[\alpha_3, 1/16, \gamma_3] \times \{0^{n/2}\} \\ L_{2i+1}^4 &= \{0^{n/2}\} \times D^{n/2}[\alpha_3, 1/16, \gamma_3] \end{cases}$$

As before, this avoids filtering at the first level. These lists have size: $\ell_4 = f(\alpha_3, 1/16 + \alpha_3 - 2\gamma_3, \gamma_3)/2$.

Level 3. We merge into 8 lists $L_0^3 \dots L_7^3$, with a constraint on c_3 bits. As there is no filtering, these lists have size: $\ell_3 = f(\alpha_3, 1/16 + \alpha_3 - 2\gamma_3, \gamma_3) - c_3$.

Level 2. We now merge and filter. We force a target distribution $D^n[\alpha_2, 1/8, \gamma_2]$, with α_2 and γ_2 to be optimized later. There is a first filtering probability p_2. We have $\ell_2 = 2\ell_3 - (c_2 - c_3) + p_2$.

Level 1. Similarly, we have: $\ell_1 = 2\ell_2 - (c_1 - c_2) + p_1$.

Level 0. We have $\ell_0 = 2\ell_1 - (1 - c_1) + p_0 = 0$, since the goal is to obtain one solution in the list L_0.

With these constraints, we find a time $\widetilde{\mathcal{O}}\left(2^{0.2830n}\right)$ (rounded upwards) with the following parameters:

$$\alpha_1 = 0.0340, \alpha_2 = 0.0311, \alpha_3 = 0.0202, \gamma_1 = 0.0041, \gamma_2 = 0.0006, \gamma_3 = 0.0001$$
$$c_1 = 0.8067, c_2 = 0.5509, c_3 = 0.2680, p_0 = -0.2829, p_1 = -0.0447, p_2 = -0.0135$$
$$\ell_1 = 0.2382, \ell_2 = 0.2694, \ell_3 = 0.2829, \ell_4 = 0.2755$$

Remark 2 (On numeric optimizations). All algorithms since HGJ, including quantum ones, rely on (nonlinear) numeric optimizations. Their correctness is easy to check, since the obtained parameters satisfy the constraints, but there is no formal proof that the parameters are indeed optimal for a given constraint set. The same goes for all algorithms studied in this paper. In order to gain confidence in our results, we tried many different starting points and several equivalent rewriting of the constraints.

Remark 3 (Adding more symbols). In general, adding more symbols ("-2"s, "3"s, *etc.*) can only increase the parameter search space and improve the optimal time complexity. However, we expect that the improvements from adding more symbols will become smaller and smaller, while the obtained constraints will become more difficult to write down and the parameters harder to optimize. Note that adding "-1"s decreases the time complexity exponent by 0.048, while adding "2"s decreases it only by 0.006.

Remark 4 (On the number of levels). Algorithms based on merging-and-filtering, classical and quantum, have a number of levels (say, 4 or 5) which must be selected before writing down the constraints. The time complexity is a decreasing function of the number of levels, which quickly reaches a minimum. In all algorithms studied in this paper, adding one more level does not change the cost of the upper levels, which will remain the most expensive.

3 Quantum Preliminaries and Previous Work

In this section, we recall some preliminaries of quantum computation (quantum search and quantum walks) that will be useful throughout the rest of this paper. We also recall previous quantum algorithms for subset-sum. As we consider all our algorithms from the point of view of asymptotic complexities, and neglect polynomial factors in n, a high-level overview is often enough, and we will use quantum building blocks as black boxes. The interested reader may find more details in [35].

3.1 Quantum Preliminaries

All the quantum algorithms considered in this paper run in the quantum circuit model, with quantum random-access memory, often denoted as qRAM. "Baseline" quantum circuits are simply built using a universal gate set. Many quantum algorithms use qRAM access, and require the circuit model to be augmented with the so-called "qRAM gate". This includes subset-sum, lattice sieving and generic decoding algorithms that obtain time speedups with respect to their classical counterparts. Given an input register $1 \leq i \leq r$, which represents the index of a memory cell, and many quantum registers $|x_1, \ldots x_r\rangle$, which represent stored data, the qRAM gate fetches the data from register x_i:

$$|i\rangle |x_1, \ldots x_r\rangle |y\rangle \mapsto |i\rangle |x_1, \ldots x_r\rangle |y \oplus x_i\rangle \quad .$$

We will use the terminology of [24] for the qRAM gate:

- If the input i is classical, then this is the plain quantum circuit model (with classical RAM);
- If the x_j are classical, we have *quantum-accessible classical memory* (QRACM)
- In general, we have *quantum-accessible quantum memory* (QRAQM)

All known quantum algorithms for subset-sum with a quantum time speedup over the best classical one require QRAQM. For comparison, speedups on heuristic lattice sieving algorithms exist in the QRACM model [26,23], including the best one to date [25]. While no physical architecture for quantum random access has been proposed that would indeed produce a constant or negligible overhead in time, some authors [24] consider the separation meaningful. If we assign a cost $\mathcal{O}(N)$ to a QRACM query of N cells, then we can replace it by *classical* memory. Subset-sum algorithms were studied in this setting by Helm and May [19].

Quantum Search. One of the most well-known quantum algorithms is Grover's unstructured search algorithm [17]. We present here its generalization, amplitude amplification [12].

Lemma 3 (Amplitude amplification, from [12]). *Let $\mathcal{A}$ be a reversible quantum circuit, f a computable boolean function over the output of $\mathcal{A}$, O_f its implementation as a quantum circuit, and a be the initial success probability of $\mathcal{A}$, that is, the probability that $O_f\mathcal{A}|0\rangle$ outputs "true". There exists a quantum reversible algorithm that calls $\mathcal{O}\left(\sqrt{1/a}\right)$ times $\mathcal{A}$, $\mathcal{A}^\dagger$ and O_f, uses as many qubits as $\mathcal{A}$ and O_f, and produces an output that passes the test f with probability greater than $\max(a, 1-a)$.*

This is known to be optimal when the functions are black-box oracles [5].

As we will use quantum search as a subprocedure, we make some remarks similar to [33, Appendix A.2] and [10, Section 5.2] to justify that, up to additional polynomial factors in time, we can consider it runs with no errors and allows to return all the solutions efficiently.

Remark 5 (Error in a sequence of quantum searches). Throughout this paper, we will assume that a quantum search in a search space of size S with T solutions runs in exact time $\sqrt{S/T}$. In practice, there is a constant overhead, but since S and T are always exponential in n, the difference is negligible. Furthermore, this is a probabilistic procedure, and it will return a wrong result with a probability of the order $\sqrt{T/S}$. As we can test if an error occurs, we can make it negligible by redoing the quantum search polynomially many times.

Remark 6 (Finding all solutions). Quantum search returns a solution among the T possibilities, selected uniformly at random. Finding all solutions is then an instance of the coupon collector problem with T coupons [34]; all coupons are collected after on average $\mathcal{O}(T\log(T))$ trials. However, *in the QRACM model*, which is assumed in this paper, this logarithmic factor disappears. We can run the search of Lemma 3 with a new test function that returns 0 if the output of $\mathcal{A}$ is incorrect, *or* if it is correct but has already been found. The change to the runtime is negligible, and thus, we collect all solutions with only $\mathcal{O}(T)$ searches.

Quantum Walks. Quantum walks can be seen as a generalization of quantum search. They allow to obtain polynomial speedups on many unstructured problems, with sometimes optimal results (*e.g.* Ambainis' algorithm for element distinctness [1]). In this paper, we consider walks in the MNRS framework [30].

Let $G = (V, E)$ be an undirected, connected, regular graph, such that some vertices of G are "marked". Let ϵ be the fraction of marked vertices, that is, a random vertex has a probability ϵ of being marked. Let δ be the spectral gap of G, which is defined as the difference between its two largest eigenvalues.

In a *classical* random walk on G, we can start from any vertex and reach the stationary distribution in approximately $\frac{1}{\delta}$ random walk steps. Then, such a random vertex is marked with probability ϵ. Assume that we have a procedure `Setup` that samples a random vertex to start with in time S, `Check` that verifies if a vertex is marked or not in time C and `Update` that performs a walk step in time U, then we will have found a marked vertex in expected time: $\mathsf{S} + \frac{1}{\epsilon}\left(\frac{1}{\delta}\mathsf{U} + \mathsf{C}\right)$.

Quantum walks reproduce the same process, except that their internal state is not a vertex of G, but a superposition of vertices. The walk starts in the uniform superposition $\sum_{v \in V} |v\rangle$, which must be generated by the `Setup` procedure. It repeats $\sqrt{1/\epsilon}$ iterations that, similarly to amplitude amplification, move the amplitude towards the marked vertices. An update produces, from a vertex, the superposition of its neighbors. Each iteration does not need to repeat $\frac{1}{\delta}$ vertex updates and, instead, takes a time equivalent to $\sqrt{1/\delta}$ updates to achieve a good mixing. Thanks to the following theorem , we will only need to specify the setup, checking and update unitaries.

Theorem 1 (Quantum walk on a graph (adapted from [30])). *Let $G = (V, E)$ be a regular graph with spectral gap $\delta > 0$. Let $\epsilon > 0$ be a lower bound on the probability that a vertex chosen randomly of G is marked. For a random walk on G, let $\mathsf{S}, \mathsf{U}, \mathsf{C}$ be the setup, update and checking cost. Then there exists a quantum algorithm that with high probability finds a marked vertex in time*

$$\mathcal{O}\left(\mathsf{S} + \frac{1}{\sqrt{\epsilon}}\left(\frac{1}{\sqrt{\delta}}\mathsf{U} + \mathsf{C}\right)\right).$$

3.2 Solving Subset-sum with Quantum Walks

In 2013, Bernstein, Jeffery, Lange and Meurer [6] constructed quantum subset sum algorithms inspired by Schroeppel-Shamir [37] and HGJ [21]. We briefly explain the idea of their quantum walk for HGJ. The graph G that they consider is a product Johnson graph. We recall formal definitions from [22].

Definition 4 (Johnson graph). *A Johnson graph $J(N, R)$ is an undirected graph whose vertices are the subsets of R elements among a set of size N, and there is an edge between two vertices S and S' iff $|S \cap S'| = R - 1$, in other words, if S' can be obtained from S by replacing an element. Its spectral gap is given by $\delta = \frac{N}{R(N-R)}$.*

Theorem 2 (Cartesian product of Johnson graphs [22]). *Let $J^m(N,R)$ be defined as the cartesian product of m Johnson graphs $J(N,R)$,* i.e.*, a vertex in $J^m(N,R)$ is a tuple of m subsets $S_1, \dots S_m$ and there is an edge between $S_1, \dots S_m$ and $S'_1, \dots S'_m$ iff all subsets are equal at all indices except one index i, which satisfies $|S_i \cap S'_i| = R-1$. Then it has $\binom{N}{R}^m$ vertices and its spectral gap is greater than $\frac{1}{m}\frac{N}{R(N-R)}$.*

In [6], a vertex contains a product of 8 sublists $L'^3_0 \subset L^3_0, \dots, L'^3_7 \subset L^3_7$ of a smaller size than the classical lists: $\ell < \ell_3$. There is an edge between two vertices if we can transform one into the other by replacing only one element in one of the sublists. The spectral gap of such a graph is (in $\log_2$, relative to n) $-\ell$.

In addition, each vertex has an internal data structure which reproduces the HGJ merging tree, from level 3 to level 0. Since the initial lists are smaller, the list L^0 is now of expected size $8(\ell - \ell_3)$ (in $\log_2$, relative to n), *i.e.* the walk needs to run for $4(\ell_3 - \ell)$ steps. Each step requires $\ell/2$ updates.

In the `Setup` procedure, we simply start from all choices for the sublists and build the tree by merging and filtering. Assuming that the merged lists have decreasing sizes, the setup time is ℓ. The vertex is marked if it contains a solution at level 0. Hence, checking if a vertex is marked takes time $\mathsf{C} = 1$, but the update procedure needs to ensure the consistency of the data structure. Indeed, when updating, we remove an element $\mathbf{e}$ from one of the lists L'^3_i and replace it by a $\mathbf{e}'$ from L^3_i. We then have to track all subknapsacks in the upper levels where $\mathbf{e}$ intervened, to remove them, and to add the new collisions where $\mathbf{e}'$ intervenes.

Assuming that the update can run in $\text{poly}(n)$, an optimization with the new parameter ℓ yields an exponent 0.241. In [6], the parameters are such that on average, a subknapsack intervenes only in a single sum at the next level. The authors propose to simply limit the number of elements to be updated at each level, in order to guarantee a constant update time.

Quantum Walk Based on BCJ. In [18], Helm and May quantize, in the same way, the BCJ algorithm. They add "-1" symbols and a new level in the merging tree data structure, reaching a time exponent of 0.226. But they remark that this result depends on a conjecture, or a heuristic, that was implicit in [6].

Heuristic 2 (Helm-May) *In these quantum walk subset-sum algorithms, an update with expected constant time U can be replaced by an update with exact time U without affecting the runtime of the algorithm, up to a polynomial factor.*

Indeed, it is easy to construct "bad" vertices and edges for which an exact update, *i.e.* the complete reconstruction of the merging tree, will take exponential time: by adding a single new subknapsack $\mathbf{e}$, we find an exponential number of pairs $\mathbf{e} + \mathbf{e}'$ to include at the next level. So we would like to update only a few elements among them. But in the MNRS framework, the data structure of a vertex must depend solely on the vertex itself (*i.e.* on the lowest-level lists in the merging tree). And if we do as proposed in [6], we add a dependency on the path that lead to the vertex, and lose the consistency of the walk.

In a related context, the problem of "quantum search with variable times" was studied by Ambainis [2]. In a quantum search for some x such that $f(x) = 1$, in a set of size N, if the time to evaluate f on x is always 1, then the search requires time $\mathcal{O}\left(\sqrt{N}\right)$. Ambainis showed that if the elements have different evaluation times $t_1, \dots t_N$, then the search now requires $\widetilde{\mathcal{O}}(\sqrt{t_1^2 + \dots + t_N^2})$, the geometric mean of $t_1, \dots t_N$. As quantum search can be seen as a particular type of quantum walk, this shows that Heuristic 2 is wrong in general, as we can artificially create a gap between the geometric mean and expectation of the update time U; but also, that it may be difficult to actually overcome. In this paper, we will obtain different heuristic and non-heuristic times.

4 Quantum Asymmetric HGJ

In this section, we give the first quantum algorithm for the subset-sum problem, *in the QRACM model*, with an asymptotic complexity smaller than BCJ.

4.1 Quantum Match-and-Filter

We open this section with some technical lemmas that replace the classical merge-and-filter Lemma 2. In this section, we will consider a merging tree as in the HGJ algorithm, but this tree will be built using quantum search. The following lemmas bound the expected time of merge-and-filter *and* match-and-filter operations performed quantumly, in the QRACM model. This will have consequences both in this section and in the next one.

First, we remark that we can use a much more simple data structure than the ones in [1,6]. In this data structure, we store pairs $\mathbf{e}, \mathbf{e} \cdot \mathbf{a}$ indexed by $\mathbf{e} \cdot \mathbf{a} \mod M$ for some $M \simeq 2^m$.

Definition 5 (Unique modulus list). *A* unique modulus list *is a qRAM data structure $\mathcal{L}(M)$ that stores at most M entries $(\mathbf{e}, \mathbf{e} \cdot \mathbf{a})$, indexed by $\mathbf{e} \cdot \mathbf{a} \mod M$, and supports the following operations:*

- *Insertion: inserts the entry $(\mathbf{e}, \mathbf{e} \cdot \mathbf{a})$ if the modulus is not already occupied;*
- *Deletion: deletes $(\mathbf{e}, \mathbf{e} \cdot \mathbf{a})$ (not necessary in this section)*
- *Query in superposition: returns the superposition of all entries $(\mathbf{e}, \mathbf{e} \cdot \mathbf{a})$ with some modular condition on $\mathbf{e} \cdot \mathbf{a}$,* e.g. *$\mathbf{e} \cdot \mathbf{a} = t \mod M'$ for some t and some modulus M'.*

Note that all of these operations, *including* the query in superposition of all the entries with a given modulus, cost $\mathcal{O}(1)$ qRAM gates only. For the latter, we need only some Hadamard gates to prepare the adequate superposition of indices. Furthermore, the list remains sorted by design.

Next, we write a lemma for quantum *matching* with filtering, in which one of the lists is not written down. We start from a unitary that produces the uniform superposition of the elements of a list L_1, and we wrap it into an amplitude amplification, in order to obtain a unitary that produces the uniform superposition of the elements of the merged-and-filtered list.

Lemma 4 (Quantum matching with filtering). *Let L_2 be a list stored in QRACM (with the* unique modulus list *data structure of Definition 5). Assume given a unitary U that produces in time t_{L_1} the uniform superposition of $L_1 = x_0, \dots x_{2^m-1}$ where $x_i = (\mathbf{e_i}, \mathbf{e_i} \cdot \mathbf{a})$. We merge L_1 and L_2 with a modular condition of cn bits and a filtering probability p. Let L be the merged list and L^f the filtered list. Assume $|L^f| \geq 1$. Then there exists a unitary U' producing the uniform superposition of L^f in time: $\mathcal{O}\left(\frac{t_{L_1}}{\sqrt{p}} \max(\sqrt{2^{cn}/|L_2|}, 1)\right)$.*

Notice that this is also the time complexity to produce a single random element of L^f. If we want to produce and store the whole list L^f, it suffices to multiply this complexity by the number of elements in L^f (*i.e.* $p|L_1||L_2|/2^{cn}$). We would obtain: $\mathcal{O}\left(t_{L_1}\sqrt{p} \max\left(|L_1|\sqrt{\frac{|L_2|}{2^{cn}}}, \frac{|L_1||L_2|}{2^{cn}}\right)\right)$.

Proof. Since L_2 is stored in a unique modulus list, all its elements have distinct moduli. Note that the *expected* sizes of L and L^f follow from Heuristic 1. Although the number of iterations of quantum search should depend on the *real* sizes of these lists, the concentration around the average is so high (given by Chernoff bounds) that the error remains negligible if we run the search with the expected number of iterations. We separate three cases.

- If $|L_2| < 2^{cn}$, then we have no choice but to make a quantum search on elements of L_1 that match the modular constraint and pass the filtering step, in time: $\mathcal{O}\left(t_{L_1}\sqrt{\frac{2^{cn}}{L_2 p}}\right)$.
- If $|L_2| > 2^{cn}$ but $|L_2| < 2^{cn}/p$, an element of L_1 will always pass the modular constraint, with more than one candidate, but in general all these candidates will be filtered out. Given an element of L_1, producing the superposition of these candidates is done in time 1, so finding the one that passes the filter, if there is one, takes time $\sqrt{|L_2|/2^{cn}}$. Next, we wrap this in a quantum search to find the "good" elements of L_1 (passing the two conditions), with $\mathcal{O}\left(\sqrt{2^{cn}/pL_2}\right)$ iterations. The total time is:

$$\mathcal{O}\left(\sqrt{\frac{2^{cn}}{L_2 p}} \times \left(\sqrt{|L_2|/2^{cn}} \times t_{L_1}\right) = \frac{t_{L_1}}{\sqrt{p}}\right) .$$

- If $|L_2| > 2^{cn}/p$, an element of L_1 yields on average more than one filtered candidate. Producing the superposition of the modular candidates is done in time $\mathcal{O}(1)$ thanks to the data structure, then finding the superposition of filtered candidates requires $1/\sqrt{p}$ iterations. The total time is: $\mathcal{O}\left(t_{L_1}/\sqrt{p}\right)$.

The total time in all cases is: $\mathcal{O}\left(\frac{t_{L_1}}{\sqrt{p}} \max(\sqrt{2^{cn}/|L_2|}, 1)\right)$. Note that classically, the coupon collector problem would have added a polynomial factor, but this is not the case here thanks to QRACM (Remark 6). □

In the QRACM model, we have the following corollary for merging and filtering two lists of equal size. This result will be helpful in Section 4.3 and 5.

Corollary 1. *Consider two lists* L_1, L_2 *of size* $|L_1| = |L_2| = |L|$ *exponential in* n*. We merge* L_1 *and* L_2 *with a modular condition of* cn *bits, and filter with a probability* p*. Assume that* $2^{cn} < |L|$*. Then* L^f *can be written down in quantum time:* $\mathcal{O}\left(\sqrt{p}\frac{|L|^2}{2^{cn}}\right)$.

Proof. We do a quantum search to find each element of L^f. We have $t_{L_1} = \mathcal{O}(1)$ since it is a mere QRACM query, and we use Lemma 4. □

4.2 Revisiting HGJ

We now introduce our new algorithm for subset-sum in the QRACM model.

Our starting point is the HGJ algorithm. Similarly to [33], we use a merging tree in which the lists at a given level may have different sizes. Classically, this does not improve the time complexity. However, quantumly, we will use quantum filtering. Since our algorithm does not require to write data in superposition, only to read from classical registers with quantum random access, we require only QRACM instead of QRAQM.

In the following, we consider that all lists, except L_0^3, L_0^2, L_0^1, L^0, are built with classical merges. The final list L^0, containing (expectedly) a single element, and a branch leading to it, are part of a nested quantum search. Each list L_0^3, L_0^2, L_0^1, L^0 corresponds either to a search space, the solutions of a search, or both. We represent this situation on Fig. 3. Our procedure runs as follows:

1. (Classical step): build the *intermediate lists* L_1^3, L_1^2, L_1^1 and store them using a *unique modulus list* data structure (Definition 5).
2. (Quantum step): do a quantum search on L_0^3. To test a vector $\mathbf{e} \in L_0^3$:
   - Find $\mathbf{e}_3 \in L_1^3$ such that $\mathbf{e} + \mathbf{e}_3$ passes the $c_0^2 n$-bit modular constraint (assume that there is at most one such solution). There is no filtering here.
   - Find $\mathbf{e}_2 \in L_1^2$ such that $(\mathbf{e} + \mathbf{e}_3) + \mathbf{e}_2$ passes the additional $(c^1 - c_0^2)n$-bit constraint.
   - If it also passes the filtering step, find $\mathbf{e}_1 \in L_1^1$ such that $(\mathbf{e} + \mathbf{e}_3 + \mathbf{e}_2) + \mathbf{e}_1$ is a solution to the knapsack problem (and passes the filter).

Structural constraints are imposed on the tree, in order to guarantee that there exists a knapsack solution. The only difference between the quantum and classical settings is in the optimization goal: the final time complexity.

Structural Constraints. We now introduce the variables and the structural constraints that determine the shape of the tree in Fig. 3. The asymmetry happens both in the weights at level 0 and at the constraints at level 1 and 2. We write $\ell_i^j = (\log_2 |L_i^j|)/n$. With the lists built classically, we expect a symmetry to be respected, so we have: $\ell_2^3 = \ell_3^3$, $\ell_4^3 = \ell_5^3 = \ell_6^3 = \ell_7^3$, $\ell_2^2 = \ell_3^2$. We also tweak the left-right split at level 0: lists from L_2^3 to L_7^3 have a standard balanced left-right split; however, we introduce a parameter r that determines the proportion of

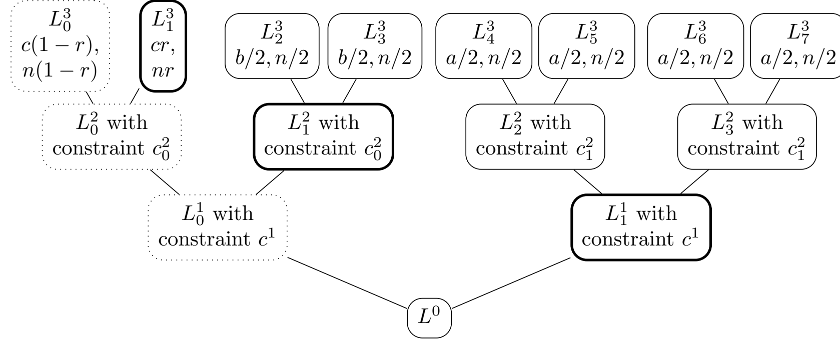

Fig. 3. Quantum HGJ algorithm. Dotted lists are search spaces (they are not stored). Bold lists are stored in QRACM. In Section 4.3, L_2^2 and L_3^2 are also stored in QRACM.

positions set to zero in list L_0^3: in L_0^3, the vectors weigh $cn(1-r)$ on a support of size $n(1-r)$, instead of $cn/2$ on a support of size $n/2$. In total we have $c+b+2a=\frac{1}{2}$, as the weight of the solution is supposed to be exactly $n/2$.

Then we note that:

- The lists at level 3 have a maximal size depending on the corresponding weight of their vectors:

$$\ell_0^3 \leq h(c)(1-r), \quad \ell_1^3 \leq h(c)r, \quad \ell_2^3 = \ell_3^3 \leq h(b)/2, \quad \ell_4^3 \leq h(a)/2$$

- The lists at level 2 cannot contain more representations than the filtered list of all subknapsacks of corresponding weight:

$$\ell_0^2 \leq h(c) - c_0^2, \quad \ell_1^2 \leq h(b) - c_0^2, \quad \ell_2^2 = \ell_3^2 \leq h(a) - c_1^2$$

- Same at levels 1 and 0: $\ell_0^1 \leq h(c+b) - c^1, \quad \ell_1^1 \leq h(2a) - c^1$
- The merging at level 2 is exact (there is no filtering):

$$\ell_0^2 = \ell_0^3 + \ell_1^3 - c_0^2, \quad \ell_1^2 = \ell_2^3 + \ell_3^3 - c_0^2, \quad \ell_2^2 = \ell_4^3 + \ell_5^3 - c_1^2, \quad \ell_3^2 = \ell_6^3 + \ell_7^3 - c_1^2,$$

- At level 1, with a constraint $c^1 \geq c_0^2, c_1^2$ that subsumes the previous ones:

$$\ell_0^1 = \ell_0^2 + \ell_1^2 - c^1 + c_0^2 + \mathsf{pf}_1(b, c) \ , \qquad \ell_1^1 = \ell_2^2 + \ell_3^2 - c^1 + c_1^2 + \mathsf{pf}_1(a, a)$$

- And finally at level 0: $\ell^0 = 0 = \ell_0^1 + \ell_1^1 - (1-c^1) + \mathsf{pf}_1(b+c, 2a)$

Classical Optimization. All the previous constraints depend on the problem, not on the computation model. Now we can get to the time complexity in the classical setting, that we want to minimize:

$$\max\big(\ell_4^3, \ell_2^3, \ell_1^3, \ell_2^3 + \ell_3^3 - c_0^2, \ell_4^3 + \ell_5^3 - c_1^2, \ell_2^2 + \ell_3^2 - c^1 + c_1^2,$$
$$\ell_0^3 + \max(\ell_1^3 - c_0^2, 0) + \max(\ell_1^2 - c^1 + c_0^2, 0) + \max(\mathsf{pf}_1(b,c) + \ell_1^1 - (1-c^1), 0)\big) \ .$$

The last term corresponds to the exhaustive search on ℓ_0^3. In order to keep the same freedom as before, it is possible that an element of L_0^3 matches against *several* elements of L_1^3, all of which yield a potential solution that has to be matched against L_1^2, *etc.* Hence for each element of L_0^3, we find the expected $\max(\ell_1^3 - c_0^2, 0)$ candidates matching the constraint c_0^1. For each of these candidates, we find the expected $\max(\ell_1^2 - c^1 + c_0^2, 0)$ candidates matching the constraint c^1. For each of these candidates, if it passes the filter, we search for a collision in L_1^1; this explains the $\max(\mathsf{pf}_1(b,c) + \ell_1^1 - (1 - c^1), 0)$ term. In the end, we check if the final candidates pass the filter on the last level.

We verified that optimizing the classical time under our constraints gives the time complexity of HGJ.

Quantum Optimization. The time complexity for producing the intermediate lists is unchanged. The only difference is the way we find the element in L_0^3 that will lead to a solution, which is a nested sequence of quantum searches.

- We can produce the superposition of all elements in L_0^2 in time

$$t_2 = \frac{1}{2}\max(c_0^2 - \ell_1^3, 0)$$

- By Lemma 4, we can produce the superposition of all elements in L_0^1 in time

$$t_2 - \frac{1}{2}\mathsf{pf}_1(b,c) + \frac{1}{2}\max\left(c^1 - c_0^2 - \ell_1^2, 0\right)$$

- Finally, we expect that there are $\left(\ell_0^2 + \ell_1^2 - c^1 + c_0^2 + \mathsf{pf}_1(b,c)\right)$ elements in L_0^1, which gives the number of iterations of the quantum search.

The time of this search is:

$$\frac{1}{2}\left(\ell_0^1 + \max(c_0^2 - \ell_1^3, 0) - \mathsf{pf}_1(b,c) + \max\left(c^1 - c_0^2 - \ell_1^2, 0\right)\right)$$

and the total time complexity is:

$$\max\left(\ell_4^3, \ell_2^3, \ell_1^3, \ell_2^3 + \ell_3^3 - c_0^2, \ell_4^3 + \ell_5^3 - c_1^2, \ell_2^2 + \ell_3^2 - c^1 + c_1^2,\right.$$
$$\left.\frac{1}{2}\left(\ell_0^1 + \max(c_0^2 - \ell_1^3, 0) - \mathsf{pf}_1(b,c) + \max\left(c^1 - c_0^2 - \ell_1^2, 0\right)\right)\right)$$

We obtain a quantum time complexity exponent of 0.2374 with this method (the detailed parameters are given in Table 2).

4.3 Improvement via Quantum Filtering

Let us keep the tree structure of Figure 3 and its structural constraints. The final quantum search step is already made efficient with respect to the filtering of representations, as we only pay half of the filtering term $\mathsf{pf}_1(b,c)$. However,

we can look towards the *intermediate lists* in the tree, *i.e.* L_1^3, L_1^2, L_1^1. The merging at the first level is exact: due to the left-right split, there is no filtering of representations, hence the complexity is determined by the size of the output list. However, the construction of L_1^1 contains a filtering step. Thus, we can use Corollary 1 to produce the elements of L_1^1 faster and reduce the time complexity from: $\ell_2^2 + \ell_3^2 - c^1 + c_1^2$ to: $\ell_2^2 + \ell_3^2 - c^1 + c_1^2 + \frac{1}{2}\mathsf{pf}_1(a, a)$. By optimizing with this time complexity, we obtain a time exponent 0.2356 (the detailed parameters are given in Table 2). The corresponding memory is 0.2356 (given by the list L_1^3).

Table 2. Optimization results for the quantum asymmetric HGJ algorithm (in $\log_2$ and relative to n), rounded to four digits. The time complexity is an upper bound.

Variant	Time	a	b	c	ℓ_0^3	ℓ_1^3	ℓ_2^3	ℓ_4^3	ℓ_0^2	ℓ_0^1
Classical	0.3370	0.1249	0.11	0.1401	0.3026	0.2267	0.25	0.2598	0.3369	0.3114
Section 4.2	0.2374	0.0951	0.0951	0.2146	0.4621	0.2367	0.2267	0.2267	0.4746	0.4395
Section 4.3	0.2356	0.0969	0.0952	0.2110	0.4691	0.2356	0.2267	0.2296	0.4695	0.4368

Remark 7 (More improvements). We have tried increasing the tree depth or changing the tree structure, but it does not seem to bring any improvement. In theory, we could allow for more general representations involving "-1" and "2". However, computing the filtering probability, when merging two lists of subknapsacks in $D^n[\alpha, \beta, \gamma]$ *with different distributions* becomes much more technical. We managed to compute it for $D^n[\alpha, \beta]$, but the number of parameters was too high for our numerical optimizer, which failed to converge.

4.4 Quantum Time-Memory Tradeoff

In the original HGJ algorithm, the lists at level 3 contain full distributions $D^{n/2}[0, 1/8]$. By reducing their sizes to a smaller exponential, one can still run the merging steps, but the final list L^0 is of expected size exponentially small in n. Hence, one must redo the tree many times. This general time-memory tradeoff is outlined in [21] and is also reminiscent of Schroeppel and Shamir's algorithm [37], which can actually be seen as repeating $2^{n/4}$ times a merge of lists of size $2^{n/4}$, that yields $2^{-n/4}$ solutions on average.

Asymmetric Tradeoff. The tradeoff that we propose is adapted to the QRACM model. It consists in increasing the asymmetry of the tree: we reduce the sizes of the intermediate lists L_1^3, L_1^2, L_1^1 in order to use less memory; this in turn increases the size of L_0^3, L_0^2 and L_0^1 in order to ensure that a solution exists. We find that this tradeoff is close to the time-memory product curve $TM = 2^{n/2}$, and actually slightly better (the optimal point when $m = 0.2356$ has $TM = 2^{0.4712n}$). This is shown on Figure 4. At $m = 0$, we start at $2^{n/2}$, where L_0^3 contains all vectors of Hamming weight $n/2$.

Fact 1 *For any memory constraint* $m \leq 0.2356$ *(in* $\log_2$ *and proportion of* n*), the optimal time complexity in the quantum asymmetric HGJ algorithm of Section 4.3 is lower than* $\widetilde{\mathcal{O}}\left(2^{n/2-m}\right)$.

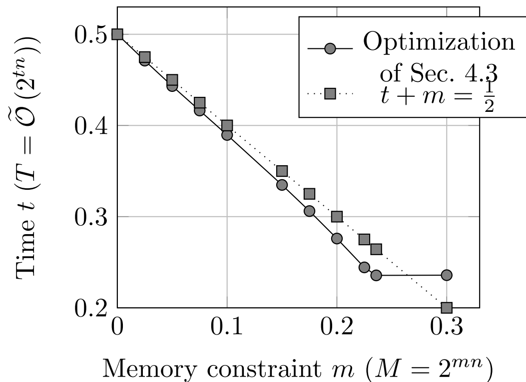

Fig. 4. Quantum time-memory tradeoff of the asymmetric HGJ algorithm

Improving the QRACM usage. In trying to reduce the quantum or quantum-accessible hardware used by our algorithm, it makes sense to draw a line between QRACM and classical RAM, *i.e.* between the part of the memory that is actually accessed quantumly, and the memory that is used only classically. We now try to enforce the constraint only on the QRACM, using possibly more RAM. In this context, we cannot produce the list L_1^1 via quantum filtering. The memory constraint on lists L_1^3, L_1^2, L_1^1 still holds; however, we can increase the size of lists $L_4^3, L_5^3, L_6^3, L_7^3, L_2^2, L_3^2$.

Fact 2 *For any QRACM constraint* $m \leq 0.2356$*, the optimal time complexity obtained by using more RAM is always smaller than the best optimization of Section 4.3.*

The difference remains only marginal, as can be seen in Table 3, but it shows a tradeoff between quantum and classical resources.

5 New Algorithms Based on Quantum Walks

In this section, we improve the algorithm by Helm and May [18] based on BCJ and the MNRS quantum walk framework. Our algorithm is a quantum walk on a product Johnson graph, as in Section 3.2. There are two new ideas involved.

Table 3. Time-memory tradeoffs (QRACM) for three variants of our asymmetric HGJ algorithm, obtained by numerical optimization, and rounded upwards. The last variant uses more classical RAM than the QRACM constraint.

QRACM	Section 4.2		Section 4.3		With more RAM	
bound	Time	Memory	Time	Memory	Time	Memory
0.0500	0.4433	0.0501	0.4433	0.0501	0.4412	0.0650
0.1000	0.3896	0.1000	0.3896	0.1000	0.3860	0.1259
0.1500	0.3348	0.1501	0.3348	0.1501	0.3301	0.1894
0.3000	0.2374	0.2373	0.2356	0.2356	0.2373	0.2373

5.1 Asymmetric 5th level

In our new algorithm, we can afford one more level than BCJ. We then have a 6-level merging tree, with levels numbered 5 down to 0. Lists at level i all have the same size ℓ_i, *except at level 5.* Recall that the merging tree, and all its lists, is the additional data structure attached to a node in the Johnson graph. In the original algorithm of [18], there are 5 levels, and a node is a collection of 16 lists, each list being a subset of size ℓ_4 among the $g(1/16 + \alpha_3, \alpha_3)/2$ vectors having the right distribution.

In our new algorithm, at level 5, we separate the lists into "left" lists of size ℓ_5^l and "right" lists of size ℓ_5^r. The quantum walk will only be performed on the left lists, while the right ones are full enumerations. Each list at level 4 is obtained by merging a "left" and a "right" list. The left-right-split at level 5 is then asymmetric: vectors in one of the left lists L_5^l are sampled from $D^{\eta n}[\alpha_4, 1/32, \gamma_4] \times \{0^{(1-\eta)n}\}$ and the right lists L_5^r contain *all* the vectors from $\{0^{\eta n}\} \times D^{(1-\eta)n}[\alpha_4, 1/32, \gamma_4]$. This yields a new constraint: $\ell_5^r = f(1/32 + \alpha_4 - 2\gamma_4, \alpha_4, \gamma_4)(1 - \eta)$.

While this asymmetry does not bring any advantage classically, it helps in reducing the update time. We enforce the constraint $\ell_5^r = c_4$, so that for each element of L_5^l, there is on average one matching element in L_5^r. So updating the list L_4 at level 4 is done on average time 1. Then we also have $\ell_4 = \ell_5^l$.

With this construction, ℓ_5^r and ℓ_5^l are actually unneeded parameters. We only need the constraints $c_4 (= \ell_5^r) = f(1/32 + \alpha_4 - 2\gamma_4, \alpha_4, \gamma_4)(1 - \eta)$ and $\ell_4 (= \ell_5^l) \leq f(1/32 + \alpha_4 - 2\gamma, \alpha_4, \gamma_4)\eta$. The total setup time is now:

$$\mathsf{S} = \max\Big(\underbrace{c_4, \ell_4}_{\text{Lv. 5 and 4}}, \underbrace{2\ell_4 - (c_3 - c_4)}_{\text{Level 3}}, \underbrace{2\ell_3 - (c_2 - c_3)}_{\text{Level 2}}, \underbrace{2\ell_2 - (c_1 - c_2)}_{\text{Level 1}}, \underbrace{\ell_1 + \max(\ell_1 - (1 - c_1), 0)}_{\text{Level 0}} \Big)$$

and the expected update time for level 5 (inserting a new element in a list L_5^l at the bottom of the tree) and at level 4 (inserting a new element in L_4) is 1.

The spectral gap of the graph is $\delta = -\ell_5^l$ and the proportion of marked vertices is $\epsilon = -\ell_0$.

Saturation Constraints. In the quantum walk, we have $\ell_0 < 0$, since we expect only some proportion of the nodes to be marked (to contain a solution). This proportion is hence ℓ_0. The saturation constraints are modified as follows:

$$\ell_5^l \leq \frac{\ell_0}{16} + f(\tfrac{1}{32} + \alpha_4 - 2\gamma_4, \alpha_4, \gamma_4)\eta, \qquad \ell_4 \leq \frac{\ell_0}{16} + f(\tfrac{1}{32} + \alpha_4 - 2\gamma_4, \alpha_4, \gamma_4) - c_4$$
$$\ell_3 \leq \frac{\ell_0}{8} + f(\tfrac{1}{16} + \alpha_3 - 2\gamma_3, \alpha_3, \gamma_3) - c_3, \; \ell_2 \leq \frac{\ell_0}{4} + f(\tfrac{1}{8} + \alpha_2 - 2\gamma_2, \alpha_2, \gamma_2) - c_2$$
$$\ell_1 \leq \frac{\ell_0}{2} + f(1/4 + \alpha_1 - 2\gamma_1, \alpha_1, \gamma_1) - c_1$$

Indeed, the *classical* walk will go through a total of $-\ell_0$ trees before finding a solution. Hence, it needs to go through $-\ell_0/16$ different lists at level 5 (and 4), which is why we need to introduce ℓ_0 in the saturation constraint: there must be enough elements, not only in L_5^l, but in the whole search space that will be spanned by the walk. These constraints ensure the existence of marked vertices in the walk.

5.2 Better Setup and Updates using quantum search

Along the lines of Lemma 4 and corollary 1, we now show how to use a quantum search to speed up the Setup and Update steps in the quantum walk. As the structure of the graph is unchanged, we still have $\epsilon = -\ell_0$ and a spectral gap $\delta = -\ell_5^l$.

Setup. Let $p_i, (1 \leq i \leq 3)$ be the filtering probabilities at level i, *i.e.* the (logarithms of the) probabilities that an element that satisfies the modulo condition resp. at level i also has the desired distribution of 0s, 1s, -1s and 2s, and appears in list L_i. Notice that $p_i \leq 0$. Due to the left-right split, there is no filtering at level 4.

We use quantum filtering (Corollary 1) to speed up the computation of lists at levels 3, 2 and 1 in the setup, reducing in general a time $2\ell - c$ to $2\ell - c + \mathsf{pf}/2$. It does not apply for level 0, since L_0 has a negative expected size. At this level, we will simply perform a quantum search over L_1. If there is too much constraint, *i.e.* $(1 - c_1) > \ell_1$, then for a given element in L_1, there is on average less than one modular candidate. If $(1 - c_1) < \ell_1$, there is on average more than one (although less than one with the filter) and we have to do another quantum search on them all. This is why the setup time at level 0, in full generality, becomes $(\ell_1 + \max(\ell_1 - (1 - c_1), 0))/2$. The setup time can thus be improved to:

$$\mathsf{S} = \max\Big(\underbrace{c_4, \ell_4}_{\text{Lv. 5 and 4}}, \underbrace{2\ell_4 - (c_3 - c_4) + p_3/2}_{\text{Level 3}}, \underbrace{2\ell_3 - (c_2 - c_3) + p_2/2}_{\text{Level 2}}, \underbrace{2\ell_2 - (c_1 - c_2) + p_1/2}_{\text{Level 1}}, \underbrace{(\ell_1 + \max(\ell_1 - (1 - c_1), 0))/2}_{\text{Level 0}} \Big) .$$

Update. Our update will also use a quantum search. First of all, recall that the updates of levels 5 and 4 are performed in (expected) time 1. Having added an element in L_4, we need to update the upper level. There are on average $\ell_4 - (c_3 - c_4)$ candidates satisfying the modular condition. To avoid a blowup in the time complexity, we forbid to have more than one element inserted in L_3 on average, which means: $\ell_4 - (c_3 - c_4) + p_3 \leq 0 \iff \ell_3 \leq \ell_4$. We then find this element, if it exists, with a quantum search among the $\ell_4 - (c_3 - c_4)$ candidates.

Similarly, as at most one element is updated in L_3, we can move on to the upper levels 2, 1 and 0 and use the same argument. We forbid to have more than one element inserted in L_2 on average: $\ell_3 - (c_2 - c_3) + p_2 \leq 0 \iff \ell_2 \leq \ell_3$, and in L_1: $\ell_1 \leq \ell_2$. At level 0, a quantum search may not be needed, hence a time $\max(\ell_1 - (1 - c_1), 0)/2$. The expected update time becomes:

$$\mathsf{U} = \max\Bigg(0, \underbrace{(\ell_4 - (c_3 - c_4))/2}_{\text{Level 3}}, \underbrace{(\ell_3 - (c_2 - c_3))/2}_{\text{Level 2}}, \underbrace{(\ell_2 - (c_1 - c_2))/2}_{\text{Level 1}}, \underbrace{(\ell_1 - (1 - c_1))/2}_{\text{Level 0}}\Bigg) .$$

5.3 Parameters

Using the following parameters, we found an algorithm that runs in time $\widetilde{\mathcal{O}}\left(2^{0.2156n}\right)$:

$$\ell_0 = -0.1916, \ell_1 = 0.1996, \ell_2 = 0.2030, \ell_3 = 0.2110, \ell_4 (= \ell_5^l) = 0.2110$$
$$c_1 = 0.6190, c2 = 0.4445, c3 = 0.2506, c_4 (= \ell_5^r) = 0.0487$$
$$\alpha_1 = 0.0176, \alpha_2 = 0.0153, \alpha_3 = 0.0131, \alpha_4 = 0.0087$$
$$\gamma_1 = 0.0019, \gamma_2 = \gamma_3 = \gamma_4 = 0, \eta = 0.8448$$

There are many different parameters that achieve the same time. The above set achieves the lowest memory that we found, at $\widetilde{\mathcal{O}}\left(2^{0.2110n}\right)$. Note that time and memory complexities are different in this quantum walk, contrary to previous works, since the update procedure has now a (small) exponential cost.

Remark 8 (Time-memory tradeoffs). Quantum walks have a natural time-memory tradeoff which consists in reducing the vertex size. Smaller vertices have a smaller chance of being marked, and the walk goes on for a longer time. This is also applicable to our algorithms, but requires a re-optimization with a memory constraint.

6 Mitigating Quantum Walk Heuristics for Subset-Sum

In this section, we provide a modified quantum walk NEW-QW for any quantum walk subset-sum algorithm QW, including [6,18] and ours, that will no longer rely on Heuristic 2. In NEW-QW, the Johnson graph is the same, but the vertex data structure and the update procedure are different (Section 6.2). It allows us to guarantee the update time, at the expense of losing some marked vertices. In Section 6.3, we will show that most marked vertices in QW remain marked.

6.1 New Data Structure for Storing Lists

The main requirement of the vertex data structure is to store lists of subknapsacks with modular constraints in QRAQM. For each list, we will use two data structures. The first one is the combination of a hash table and a skip list given in [1] (abbreviated *skip list* below) and the second one is a *Bucket-modulus list* data structure, adapted from Definition 5, that we define below.

Hash Table and Skip List. We use the data structure of [1] to store lists of entries $(\mathbf{e}, \mathbf{e} \cdot \mathbf{a})$, sorted by knapsack value $\mathbf{e} \cdot \mathbf{a}$. The data structure for M entries, that we denote $\mathcal{SL}(M)$, uses $\widetilde{\mathcal{O}}(M)$ qRAM memory cells and supports the following operations: inserting an entry in the list, deleting an entry from the list and producing the uniform superposition of entries in the list. All these operations require time $\mathsf{polylog}(M)$.

We resort to this data structure because the proposal of "radix trees" in [6] is less detailed. It is defined relatively to a choice of $\mathsf{polylog}(M) = \text{poly}(n)$ hash functions selected from a family of independent hash functions of the entries (we refer to [1] for more details). For a given choice of hash functions, the insertion or deletion operations can fail. Thus, the data structure is equipped with a superposition of such choices. Instead of storing $\mathcal{SL}(M)$, we store: $\sum_h |h\rangle |\mathcal{SL}_h(M)\rangle$ where $\mathcal{SL}_h$ is the data structure flavored with the choice of hash functions h. Insertions and deletions are performed depending on h. This allows for a globally negligible error: if sufficiently many hash functions are used, the insertion and deletion of *any element* add a global error vector of amplitude $o(2^{-n})$ *regardless of the current state of the data.* The standard "hybrid argument" from [5] and [1, Lemma 5] can then be used in the context of an MNRS quantum walk.

Proposition 1 ([1], Lemma 5, adapted). *Consider an MNRS quantum walk with a "perfect" (theoretical) update unitary U, managing data structures, and an "imperfect" update unitary U' such that, for any basis state $|x\rangle$:*

$$U' |x\rangle = U |x\rangle + |\delta_x\rangle$$

where $|\delta_x\rangle$ is an error vector of amplitude bounded by $o(2^{-n})$ for any x*. Then running the walk with U' instead of U, after T steps, the final "imperfect" state $|\psi'\rangle$ deviates from the "perfect" state $|\psi\rangle$ by: $\| |\psi'\rangle - |\psi\rangle \| \leq o(2^{-n}T)$.*

This holds as a general principle: in the update unitary, any perfect procedure can be replaced by an imperfect one as long as its error is negligible (with respect to the total number of updates) *and data-independent.* In contrast, the problem with Heuristic 2 is that a generic constant-time update induces data-dependent errors (bad cases) that do not seem easy to overcome.

Bucket-modulus List. Let $B = \text{poly}(n)$ be a "bucket size" that will be chosen later. The bucket-modulus list is a tool for making our update time data-independent: it limits the number of vectors that can have a given modulus (where moduli are of the same order as the list size).

Definition 6 (Bucket-modulus list). *A Bucket-modulus list* $\mathcal{BL}(B,M)$ *is a qRAM data structure that stores at most* $B \times M$ *entries* $(\mathbf{e}, \mathbf{e}\cdot\mathbf{a})$*, with at most* B *entries sharing the same modulus* $\mathbf{e}\cdot\mathbf{a} \mod M$*. Thus,* $\mathcal{BL}(B,M)$ *contains* M *"buckets". Buckets are indexed by moduli, and kept sorted. It supports the following operations:*

- *Insertion: insert* $(\mathbf{e}, \mathbf{e}\cdot\mathbf{a})$*. If the bucket at index* $\mathbf{e}\cdot\mathbf{a} \mod M$ *contains* B *elements,* empty the bucket*. Otherwise, sort it using a simple sorting circuit.*
- *Deletion: remove an entry from the corresponding bucket.*
- *Query in superposition: similar as in Definition 5.*

In our new quantum walks, each list will be stored in a skip list $\mathcal{SL}(M)$ associated with a bucket-modulus $\mathcal{BL}(B,M)$. Each time we insert or delete an element from $\mathcal{SL}(M)$, we update the bucket-modulus list accordingly, according to the following rules.

Upon deletion of an element $\mathbf{e}$ in $\mathcal{SL}(M)$, let $\mathbf{e}\cdot\mathbf{a} = T \mod M$, there are three cases for $\mathcal{BL}(B,M)$:

- If $|\{\mathbf{e}' \in \mathcal{SL}(M), \mathbf{e}'\cdot\mathbf{a} = T\}| > B+1$, then bucket number T in $\mathcal{BL}(B,M)$ stays empty;
- If $|\{\mathbf{e}' \in \mathcal{SL}(M), \mathbf{e}'\cdot\mathbf{a} = T\}| = B+1$, then removing $\mathbf{e}$ makes the number of elements reach the bound B, so we add them all in the bucket at index T;
- If $|\{\mathbf{e}' \in \mathcal{SL}(M), \mathbf{e}'\cdot\mathbf{a} = T\}| \leq B$, then we remove $\mathbf{e}$ from its bucket.

Upon insertion of an element $\mathbf{e}$ in $\mathcal{SL}(M)$, there are also three cases for $\mathcal{BL}(B,M)$:

- If $|\{\mathbf{e}' \in \mathcal{SL}(M), \mathbf{e}'\cdot\mathbf{a} = T\}| = B$, then we empty the bucket at index T;
- If $|\{\mathbf{e}' \in \mathcal{SL}(M), \mathbf{e}'\cdot\mathbf{a} = T\}| < B$, then we add $\mathbf{e}$ to the bucket at index T in $\mathcal{BL}(B,M)$;
- If $|\{\mathbf{e}' \in \mathcal{SL}(M), \mathbf{e}'\cdot\mathbf{a} = T\}| > B$, then the bucket is empty and remains empty.

In all cases, there are at most B insertions or deletions in a single bucket. Note that $\mathcal{BL}(B,M) \subseteq \mathcal{SL}(M)$ but that some elements of $\mathcal{SL}(M)$ will be dropped.

Remark 9. The mapping from a skip list of size M (considered as perfect), which does not "forget" any of its elements, to a corresponding bucket-modulus list with M buckets, which forgets some of the previous elements, is deterministic. Given a skip list L, a corresponding bucket modulus list L' can be obtained by inserting all elements of L into an empty bucket modulus list.

6.2 New Data Structure for Vertices

The algorithms that we consider use multiple levels of merging. However, we will focus only on a single level. Our arguments can be generalized to any constant number of merges (with an increase in the polynomial factors involved). Recall that the product Johnson graph on which we run the quantum walk is unchanged, only the data structure is adapted.

In the following, we will consider the merging of two lists L_l and L_r of subknapsacks of respective sizes ℓ_l and ℓ_r, with a modular constraint c and a filtering probability pf. The merged list is denoted $L^c = L_l \bowtie_c L_r$ and the filtered list is denoted L^f. We assume that pairs $(\mathbf{e_1}, \mathbf{e_2})$ in L^c must satisfy $(\mathbf{e_1} + \mathbf{e_2}) \cdot \mathbf{a} = 0 \mod 2^{cn}$ (the generalization to any value modulo any moduli is straightforward).

On the positive side, our new data structure can be updated, *by design*, with a fixed time that is data-independent. On the negative side, we will not build the complete list L^f, and miss some of the solutions. As we drop a fraction of the vectors, some nodes that were previously marked will potentially appear unmarked, but this fraction is polynomial at most. We defer a formal proof of this fact to Section 6.3 and focus on the runtime.

We will focus on the case where $\ell_l = \ell_r$ and either L_l or L_r are updated, which happens at all levels in our quantum walk, except the first level. Because there is no filtering at the first level, it is actually much simpler to study with the same arguments. In previous quantum walks, we had $\ell^c = 2\ell - c \leq \ell$, *i.e.* $\ell \leq c$; now we will have $2\ell - c \geq \ell$ and $2\ell - c + \mathsf{pf} \leq \ell$.

Recall that our heuristic time complexity analysis assumes an update time $(\ell - c)/2$. Indeed, the update of an element in L_l or L_r modifies on average $(\ell - c)$ elements in $L_l \bowtie_c L_r$, among which we expect at most one filtered pair $(\mathbf{e_1}, \mathbf{e_2})$ (by the inequality $2\ell - c + \mathsf{pf} \leq \ell$). We find this solution with a quantum search. In the following, we modify the data structure of vertices in order to guarantee the best update time possible, up to additional polynomial factors. We will see however that it does not reach $(\ell - c)/2$. We now define our intermediate lists and sublists, before giving the update procedure and its time complexity.

Definitions. Both lists L_l, L_r are of size $M \simeq 2^{\ell n}$. We store them in skip lists. In both L_r and L_l, for each $T \leq M$, we expect on average only one element $\mathbf{e}$ such that $\mathbf{e} \cdot \mathbf{a} = T \mod M$. We introduce two *Bucket-modulus lists* (Definition 6) $L'_l(B, M)$ and $L'_r(B, M)$ that we will write as L'_l and L'_r for simplicity, indexed by $\mathbf{e} \cdot \mathbf{a} \mod M$, with an arbitrary bound $B = \mathrm{poly}(n)$ for the bucket sizes. They are attached to L_l and L_r as detailed in Section 6.1. When an element in L_l or L_r is modified, they are modified accordingly.

In L'_l and L'_r, we consider the sublists of subknapsacks having the same modulo $C \mod 2^{cn}$, and we denote by $L'_{l,C}$ and $L'_{r,C}$ these sublists. They can be easily considered separately since the vectors are sorted by knapsack weight. By design of the bucket-modulus lists, $L'_{l,C}$ and $L'_{r,C}$ both have size at most $B2^{(\ell - c)n}$. We have:

$$L'_l \bowtie_c L'_r = \bigcup_{0 \leq C \leq 2^{cn} - 1} L'_{l,C} \times L'_{r,C} \quad .$$

Next, we have a case disjunction to make. The most complicated case is when $2\ell - 2c + \mathsf{pf} > 0$, that is, each product $L'_{l,C} \times L'_{r,C}$ for a given C yields more than one filtered pair on average. In that case, we define sublists $L'_{l,C,i}$ of $L'_{l,C}$ and sublists $L'_{r,C,j}$ of $L'_{r,C}$ using a new arbitrary modular constraint, so that each of

these sublists is of size $-\mathsf{pf}/2$ (at most). There are $\ell - c + \mathsf{pf}/2$ sublists (exactly). The rationale of this cut is that a product $L'_{l,C,i} \times L'_{r,C,j}$ for a given i, j now yields on average a single filtered pair (or less). When $2\ell - 2c + \mathsf{pf} \leq 0$, we don't perform this last cut and consider the product $L'_{l,C} \times L'_{r,C}$ immediately. By a slight abuse of notation, we denote: $(L'_{l,C,i} \times L'_{r,C,j})^f$ the set of filtered pairs from $L'_{l,C,i} \times L'_{r,C,j}$, and we have:

$$L^f = \bigcup_{0 \leq C \leq 2^{cn}-1} \bigcup_{i,j} (L'_{l,C,i} \times L'_{r,C,j})^f \ .$$

Algorithm 1 Update algorithm: given L_l, L_r of size ℓ, we insert or delete an element in L_l and update the filtered list L^f accordingly. We focus here on the case $2\ell - 2c + \mathsf{pf} > 0$.

Data: skip lists for L_l, L_r, L^f, bucket-modulus lists L'_l, L'_r
1: ▷ The bucket-modulus list for L^f will be updated later
Input: an insertion / deletion instruction for L_l
Output: updates L_l, L'_l, L^f accordingly
2: Insert or delete in L_l ▷ only one element to update
3: Update the bucket-modulus structure L'_l ▷ at most B elements to update
4: **for** each element $\mathbf{e}$ to insert / delete in L'_l **do** ▷ $B = \mathrm{poly}(n)$ iterations
5: Select its corresponding sublist $L'_{l,C,i}$
6: Let $L''_{l,C,i} = L'_{l,C,i} \cup \{\mathbf{e}\}$ or $L'_{l,C,i} \backslash \{\mathbf{e}\}$
7: **for** each sublist $L'_{r,C,j}$ **do** ▷ $\ell - c + \mathsf{pf}/2$ iterations
8: Estimate $s = |(L'_{l,C,i} \times L'_{r,C,j})^f|$ ▷ time $\widetilde{\mathcal{O}}\left(B \times 2^{-\mathsf{pf}n/2}\right)$
9: Estimate $s' = |(L''_{l,C,i} \times L'_{r,C,j})^f|$ ▷ time $\widetilde{\mathcal{O}}\left(B \times 2^{-\mathsf{pf}n/2}\right)$
▷ In the case of an insertion, $s' \geq s$ and $s' \leq s$ for a deletion
10: **if** $s > B$ and $s' \leq B$
▷ The removal of $\mathbf{e}$ makes the number of filtered pairs acceptable
11: **then** $L^f \leftarrow L^f \cup (L''_{l,C,i} \times L'_{r,C,j})^f$
12: **if** $s > B$ and $s' > B$
13: **then** do nothing
14: **if** $s \leq B$ and $s' > B$
▷ The insertion of $\mathbf{e}$ overflows the filtered pairs
15: **then** remove all $(L'_{l,C,i} \times L'_{r,C,j})^f$ from L^f
16: **if** $s \leq B$ and $s' \leq B$
17: **then** update L^f with the (at most) B new or removed pairs
18: **end for**
19: **end for**

Algorithm and Complexity. Algorithm 1 details our update procedure. We now compute its time complexity and explain why it remains data-independent. Recall that we want to avoid the "bad cases" where an update goes on for too long:

this is the case where an update in L_l (or L_r) creates too many updates in L^f. In Algorithm 1, we avoid this by deliberately limiting the number of elements that can be updated. We can see that L^f will be smaller than the "perfect" one for two reasons: • the bucket-modulus data structure loses some vectors, since the buckets are dropped when they overflow. • filtered pairs are lost. Indeed, the algorithm ensures that in L^f, at most B solutions $\mathbf{e}_l + \mathbf{e}_r$ come from a cross-product $L'_{l,C,i} \times L'_{r,C,j}$.

This makes the update procedure *history-independent* and its time complexity *data-independent*. Indeed:

Lemma 5. *The state of the data structures L_l, L_r, L^f after Algorithm 1 depends only on L_l, L_r, L^f before and on the element that was inserted / deleted.*

We omit a formal proof, as it follows from our definition of the bucket-modulus list and of Algorithm 1.

Lemma 6. *With a good choice of B, Algorithm 1 runs with a data-independent error in $o(2^n)$. The time complexity of Algorithm 1 is $\widetilde{\mathcal{O}}\left(2^{(\ell-c)n}\right)$ and an update modifies $\widetilde{\mathcal{O}}\left(2^{\max(\ell-c+\mathsf{pf}/2,0)n}\right)$ elements in the filtered list L^f at the next level (respectively $\ell - c$ and $\max(\ell - c + \mathsf{pf}/2, 0)$ in log scale).*

Proof. We check step by step the time complexity of Algorithm 1:

- Insertion and deletion from the skip list for L_l is done in $\mathrm{poly}(n)$, with a global error that can be omitted.
- The bucket-modulus list L'_l is updated in time $\mathcal{O}(B) = \mathrm{poly}(n)$ without errors. At most B elements must be inserted or removed.
- For each insertion or removal in L'_l, we select the corresponding sublist $L'_{l,C,i}$ (or simply $L'_{l,C}$ if $2\ell - 2c + \mathsf{pf} \leq 0$). We look at the sublists $L'_{r,C,j}$ and we estimate the number of filtered pairs in the products $L'_{l,C,i} \times L'_{r,C,j}$ (of size $-\mathsf{pf}$), checking whether it is smaller or bigger than B. We explain in [8, Appendix C] how to do that reversibly in time $\widetilde{\mathcal{O}}\left(B \times 2^{-\mathsf{pf}n/2}\right)$ ($-\mathsf{pf}/2$ in log scale). There are $\ell - c + \mathsf{pf}/2$ classical iterations, thus the total time is $\ell - c$.
- Depending whether we have found more or less than B filtered pairs, we will have to remove or to add all of them in L^f. This means that $B \times 2^{(\ell-c+\mathsf{pf}/2)n}$ insertion or deletion instructions will be passed over to L^f.

There are two sources of data-independent errors: first, the skip list data structure (see Section 6.1). Second, the procedure of [8, Appendix C]. Both can be made exponentially small at the price of a polynomial overhead. Note that B will be set in order to get a sufficiently small probability of error (see the next section), and can be a global $\mathcal{O}(n)$. However, the polynomial overhead of our update unitary grows with the number of levels. ◻

6.3 Fraction of Marked Vertices

Now that we have computed the update time of NEW-QW, it remains to compute its fraction ϵ_{new} of marked vertices. We will show that $\epsilon_{new} = \epsilon \left(1 - \frac{1}{\mathrm{poly}(n)}\right)$ with overwhelming probability on the random subset-sum instance, where ϵ is the previous fraction in QW.

Consider a marked vertex in QW. There is a path in the data structure leading to the solution, hence a constant number of subknapsacks $\mathbf{e_1}, \ldots, \mathbf{e_t}$ such that the vertex will remain marked *if and only* if none of them is "accidentally" discarded by our new data structure. Thus, if G is the graph of the walk, we want to upper bound:

$$\Pr_{v \in G}\begin{pmatrix} v \text{ is marked in QW and} \\ \text{not marked in NEW-QW} \end{pmatrix} \leq \sum_{\mathbf{e_i}, 1 \leq i \leq t} \Pr_{v \in G}\begin{pmatrix} \mathbf{e_i} \in v \text{ in QW} \\ \mathbf{e_i} \notin v \text{ in NEW-QW} \end{pmatrix} .$$

We focus on some level in the tree, on a list L of average size $2^{\ell n}$, and on a single vector $\mathbf{e_0}$ that must appear in L. Subknapsacks in L are taken from $\mathcal{B} \subseteq D^n[\alpha, \beta, \gamma]$. We study the event that $\mathbf{e_0}$ is accidentally discarded from L. This can happen for two reasons:

- we have $|\{\mathbf{e} \in L, \mathbf{e} \cdot \mathbf{a} = \mathbf{e_0} \cdot \mathbf{a} \mod 2^{\ell n}\}| > B$: the vector is dropped at the bucket-modulus level;
- at the next level, there are more than B pairs from some product of lists $L'_{l,C,i} \times L'_{r,C,j}$ to which the vector $\mathbf{e_0}$ belongs, that will pass the filter.

We remark the following to make our computations easier.

Fact 3 *We can replace the L from our new data structure* NEW-QW *by a list of exact size $2^{\ell n}$, which is a sublist from the list L in* QW*.*

At successive levels, our new data structure discards more and more vectors. Hence, the actual lists are smaller than in QW. However, removing a vector $\mathbf{e}$ from a list, if it does not unmark the vertex, does not increase the probability of unmarking it at the next level, since $\mathbf{e}$ does not belong to the unique solution.

Fact 4 *When a vertex in* NEW-QW *is sampled uniformly at random, given a list L at some merging level, we can assume that the elements of L are sampled uniformly at random from their distribution $\mathcal{B}$ (with a modular constraint).*

This fact translates Heuristic 1 as a global property of the Johnson graph. At the first level, nodes contain lists of exponential size which are sampled without replacement. However, when sampling with replacement, the probability of collisions is exponentially low. Thus, we can replace $\Pr_{v \in G}$ by $\Pr_{v \in G'}$ where G' is a "completed" graph containing all lists sampled uniformly at random with replacement. This adds only a negligible number of vertices and does not impact the probability of being discarded.

Number of Vectors Having the Same Modulus. Let $N \simeq 2^n$ and M be a divisor of N. Given a particular $\mathbf{e}_0 \in \mathcal{B}$ and a vector $\mathbf{a} \in \mathbb{Z}_N^n$,

$$\text{For } \mathbf{e} \in \mathcal{B}, \text{ define } X_{\mathbf{e}}(a) = \begin{cases} 1 & \text{if } \mathbf{e} \cdot \mathbf{a} = \mathbf{e}_0 \cdot \mathbf{a} \pmod{M} \\ 0 & \text{otherwise} \end{cases}$$

We prove the following Lemma in the full version of the paper [8].

Lemma 7. *If $|\mathcal{B}| \gg M \simeq |L|$, then for a $1 - \mathrm{negl}(n)$ proportion of $\mathbf{a} \in \mathbb{Z}_N^n$, and with an appropriate $B = \mathcal{O}(n)$:*

$$\Pr_{\mathbf{e}_1, \cdots, \mathbf{e}_{|L|} \sim Unif(\mathcal{B})} \left[\sum_{i=1}^{|L|} X_{\mathbf{e}_i}(\mathbf{a}) < B - 1 \right] > 1 - \frac{1}{\mathrm{poly}(n)} \quad (1)$$

For the number of filtered pairs, we use the fact that the vectors at each level are sampled uniformly at random from their distribution. If this is the case, then a Chernoff bound (similar to the proof of Lemma 7) limits the deviation of the number of filtered pairs in $L'_{l,C,i} \times L'_{r,C,j}$ from its expectation (which is 1 by construction): the probability that there are more than $B+1$ pairs is smaller than $e^{-(B+1)/3}$. By taking a sufficiently big $B = \mathcal{O}(n)$, we can take a union bound over all products of lists $L'_{l,C,i} \times L'_{r,C,j}$ in which $\mathbf{e_0}$ intervenes. We also take a union bound over the intermediate subknapsacks that we are considering. The loss of vertices remains inverse polynomial.

6.4 Time Complexities without Heuristic 2

Previous quantum subset-sum algorithms [6,18] have the same time complexities without Heuristic 2, as they fall in parameter ranges where the bucket-modulus data structure is enough. However, this is not the case of our new quantum walk. We keep the same set of constraints and optimize with a new update time. Although using the extended $\{-1, 0, 1, 2\}$ representations brings an improvement, neither do the fifth level, nor the left-right split. This simplifies our constraints. Let $\widehat{\max}(\cdot) = \max(\cdot, 0)$. The guaranteed update time becomes:

$$\mathsf{U} = \widehat{\max}\Bigg(\underbrace{\ell_3 - (c_2 - c_3)}_{\text{Level 2}}, \underbrace{\widehat{\max}(\ell_3 - (c_2 - c_3) + \frac{p_2}{2})}_{\substack{\text{Number of elements} \\ \text{to update at level 1}}} + \widehat{\max}(\ell_2 - (c_1 - c_2)),$$

$$\underbrace{\frac{1}{2}\left(\widehat{\max}\Big(\ell_3 - (c_2 - c_3) + \frac{p_2}{2}\Big) + \widehat{\max}\Big(\ell_2 - (c_1 - c_2) + \frac{p_1}{2}\Big) + \widehat{\max}(\ell_1 - (1 - c_1)) \right)}_{\text{Final quantum search among all updated elements}} \Bigg)$$

We obtain the time exponent 0.2182 (rounded upwards) with the following parameters (rounded). The memory exponent is 0.2182 as well.

$$\ell_0 = -0.2021, \ell_1 = 0.1883, \ell_2 = 0.2102, \ell_3 = 0.2182, \ell_4 = 0.2182$$
$$c_3 = 0.2182, c_2 = 0.4283, c_1 = 0.6305, p_0 = -0.2093, p_1 = -0.0298, p_2 = -0.0160$$
$$\alpha_1 = 0.0172, \alpha_2 = 0.0145, \alpha_3 = 0.0107, \gamma_1 = 0.0020$$

7 Conclusion

In this paper, we proposed improved classical and quantum heuristic algorithms for subset-sum, building upon several new ideas. First, we used extended representations ($\{-1, 0, 1, 2\}$) to improve the current best classical and quantum algorithms. In the quantum setting, we showed how to use a quantum search to speed up the process of *filtering* representations, leading to an overall improvement on existing work. We built an "asymmetric HGJ" algorithm that uses a nested quantum search, leading to the first quantum speedup on subset-sum in the model of *classical memory with quantum random access*. By combining all our ideas, we obtained the best quantum walk algorithm for subset-sum in the MNRS framework. Although its complexity still relies on Heuristic 2, we showed how to partially overcome it and obtained the first quantum walk that requires only the classical subset-sum heuristic, and the best to date for this problem.

Open Questions. We leave as open the possibility to use representations with "-1"s (or even "2"s) in a quantum asymmetric merging tree, as in Section 4.3. Another question is how to bridge the gap between heuristic and non-heuristic quantum walk complexities. In our work, the use of an improved vertex data structure seems to encounter a limitation, and we may need a more generic result on quantum walks, similar to [2]. Finally, it would be of interest to study representations with a larger set of integers.

Acknowledgments. The authors want to thank André Chailloux, Stacey Jeffery, Antoine Joux, Frédéric Magniez, Alexander May, Amaury Pouly, Nicolas Sendrier for helpful discussions and comments. Thanks to Zhenzhen Bao and the anonymous CRYPTO and ASIACRYPT referees for their detailed comments. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 714294 - acronym QUASYModo). Research also supported in part by the ERA-NET Cofund in Quantum Technologies project QuantAlgo and the French ANR Blanc project RDAM.

References

1. Ambainis, A.: Quantum Walk Algorithm for Element Distinctness. SIAM J. Comput. 37(1), 210–239 (2007)
2. Ambainis, A.: Quantum search with variable times. Theory Comput. Syst. 47(3), 786–807 (2010)
3. Becker, A., Coron, J., Joux, A.: Improved generic algorithms for hard knapsacks. In: EUROCRYPT. LNCS, vol. 6632, pp. 364–385. Springer (2011)
4. 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: EUROCRYPT. LNCS, vol. 7237, pp. 520–536. Springer (2012)
5. Bennett, C.H., Bernstein, E., Brassard, G., Vazirani, U.V.: Strengths and weaknesses of quantum computing. SIAM J. Comput. 26(5), 1510–1523 (1997)

6. Bernstein, D.J., Jeffery, S., Lange, T., Meurer, A.: Quantum algorithms for the subset-sum problem. In: PQCrypto. LNCS, vol. 7932, pp. 16–33. Springer (2013)
7. Bonnetain, X.: Improved low-qubit hidden shift algorithms. CoRR (2019)
8. Bonnetain, X., Bricout, R., Schrottenloher, A., Shen, Y.: Improved classical and quantum algorithms for subset-sum. IACR Cryptol. ePrint Arch. 2020, 168 (2020), `https://eprint.iacr.org/2020/168`
9. Bonnetain, X., Naya-Plasencia, M.: Hidden shift quantum cryptanalysis and implications. In: ASIACRYPT (1). LNCS, vol. 11272, pp. 560–592. Springer (2018)
10. Bonnetain, X., Naya-Plasencia, M., Schrottenloher, A.: Quantum security analysis of AES. IACR Trans. Symmetric Cryptol. 2019(2), 55–93 (2019)
11. Bonnetain, X., Schrottenloher, A.: Quantum security analysis of CSIDH. In: EUROCRYPT 2020. LNCS, Springer (May 2020)
12. Brassard, G., Hoyer, P., Mosca, M., Tapp, A.: Quantum amplitude amplification and estimation. Contemporary Mathematics 305, 53–74 (2002)
13. Bricout, R., Chailloux, A., Debris-Alazard, T., Lequesne, M.: Ternary syndrome decoding with large weight. In: SAC 2019. LNCS, vol. 11959, pp. 437–466. Springer (2019)
14. Böhme, E.: Verbesserte Subset-Sum Algorithmen. Master's thesis, Ruhr Universität Bochum (2011)
15. Esser, A., May, A.: Better sample - random subset sum in $2^{0.255n}$ and its impact on decoding random linear codes. CoRR abs/1907.04295 (2019), withdrawn.
16. Garey, M.R., Johnson, D.S.: Computers and Intractability: A Guide to the Theory of NP-Completeness. W. H. Freeman (1979)
17. Grover, L.K.: A Fast Quantum Mechanical Algorithm for Database Search. In: Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing 1996. pp. 212–219. ACM (1996)
18. Helm, A., May, A.: Subset sum quantumly in 1.17^{n}. In: TQC. LIPIcs, vol. 111, pp. 5:1–5:15. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2018)
19. Helm, A., May, A.: The power of few qubits and collisions - subset sum below grover's bound. In: PQCrypto. LNCS, vol. 12100, pp. 445–460. Springer (2020)
20. Horowitz, E., Sahni, S.: Computing partitions with applications to the knapsack problem. J. ACM 21(2), 277–292 (1974)
21. Howgrave-Graham, N., Joux, A.: New generic algorithms for hard knapsacks. In: EUROCRYPT. LNCS, vol. 6110, pp. 235–256. Springer (2010)
22. Kachigar, G., Tillich, J.: Quantum information set decoding algorithms. In: PQCrypto. LNCS, vol. 10346, pp. 69–89. Springer (2017)
23. Kirshanova, E., Mårtensson, E., Postlethwaite, E.W., Moulik, S.R.: Quantum algorithms for the approximate k-list problem and their application to lattice sieving. In: ASIACRYPT. LNCS, vol. 11921, pp. 521–551. Springer (2019)
24. Kuperberg, G.: Another subexponential-time quantum algorithm for the dihedral hidden subgroup problem. In: TQC. LIPIcs, vol. 22, pp. 20–34. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2013)
25. Laarhoven, T.: Search problems in cryptography. Ph.D. thesis, PhD thesis, Eindhoven University of Technology (2015)
26. Laarhoven, T., Mosca, M., van de Pol, J.: Finding shortest lattice vectors faster using quantum search. Des. Codes Cryptogr. 77(2-3), 375–400 (2015)
27. Lagarias, J.C., Odlyzko, A.M.: Solving low-density subset sum problems. In: FOCS. pp. 1–10. IEEE Computer Society (1983)
28. Lyubashevsky, V.: The parity problem in the presence of noise, decoding random linear codes, and the subset sum problem. In: APPROX-RANDOM. LNCS, vol. 3624, pp. 378–389. Springer (2005)

29. Lyubashevsky, V., Palacio, A., Segev, G.: Public-key cryptographic primitives provably as secure as subset sum. In: TCC. LNCS, vol. 5978, pp. 382–400. Springer (2010)
30. Magniez, F., Nayak, A., Roland, J., Santha, M.: Search via quantum walk. SIAM Journal on Computing 40(1), 142–164 (2011)
31. May, A., Meurer, A., Thomae, E.: Decoding random linear codes in $\widetilde{\mathcal{O}}\left(2^{0.054n}\right)$. In: ASIACRYPT. LNCS, vol. 7073, pp. 107–124. Springer (2011)
32. May, A., Ozerov, I.: On computing nearest neighbors with applications to decoding of binary linear codes. In: EUROCRYPT (1). LNCS, vol. 9056, pp. 203–228. Springer (2015)
33. Naya-Plasencia, M., Schrottenloher, A.: Optimal merging in quantum k-xor and k-sum algorithms. In: EUROCRYPT 2020. LNCS, Springer (May 2020)
34. Newman, D.J., Shepp, L.: The double dixie cup problem. The American Mathematical Monthly 67(1), 58–61 (1960)
35. Nielsen, M.A., Chuang, I.: Quantum computation and quantum information (2002)
36. Ozerov, I.: Combinatorial Algorithms for Subset Sum Problems. Ph.D. thesis, Ruhr Universität Bochum (2016)
37. Schroeppel, R., Shamir, A.: A $T = O(2^{n/2})$, $S = O(2^{n/4})$ algorithm for certain np-complete problems. SIAM Journal on Computing 10(3), 456–464 (1981)