

# Preparing many copies of a quantum state in the black-box model

Yassine Hamoudi

*Simons Institute for the Theory of Computing, Berkeley, California 94720, USA\**

(Dated: April 23, 2024)

We describe a simple quantum algorithm for preparing  $K$  copies of an  $N$ -dimensional quantum state whose amplitudes are given by a quantum oracle. Our result extends a previous work of Grover, who showed how to prepare one copy in time  $O(\sqrt{N})$ . In comparison with the naive  $O(K\sqrt{N})$  solution obtained by repeating this procedure  $K$  times, our algorithm achieves the optimal running time of  $\Theta(\sqrt{KN})$ . Our technique uses a refinement of the quantum rejection sampling method employed by Grover. As a direct application, we obtain a similar speed-up for obtaining  $K$  independent samples from a distribution whose probability vector is given by a quantum oracle.

## I. INTRODUCTION

The preparation of a specific quantum state is an important building block and a critical bottleneck in many quantum algorithms [1–4]. The objective of the STATE PREPARATION problem is to find the minimum amount of resources needed to generate a quantum state given some description of it. In general, the complexity of this problem scales linearly with the dimension of the state to be prepared [5, 6]. Yet, it is possible to achieve sublinear bounds for particular states or input models. One such example is the *black-box model* where, given oracle access to a non-zero non-negative vector  $w = (w_1, \dots, w_N)$ , the objective is to load the associated normalized probability vector into the amplitudes of the  $\log(N)$ -qubit state  $|w\rangle$  defined as

$$|w\rangle := \frac{1}{\sqrt{\mathcal{W}}} \sum_{i=1}^N \sqrt{w_i} |i\rangle \quad (1)$$

where  $\mathcal{W} = \sum_{i=1}^N w_i$  is the (unknown) normalization factor. Grover adapted his celebrated quantum search algorithm to this problem in [7], where he showed that  $O(\sqrt{N})$  queries to  $w$  are sufficient to prepare  $|w\rangle$ . In practice, one can expect that *several* copies of the same quantum state are needed for further use. For instance,  $|w\rangle$  may be fed in an algorithm that fails with some probability and that must be repeated several times. The no-cloning theorem prevents the state  $|w\rangle$  from being easily duplicated. In fact, it is easy to show that additional queries to the input are required to prepare several copies of  $|w\rangle$ . The problem of adapting the state preparation procedure to the desired number  $K$  of copies has received little attention. Usually, it is possible to simply repeat  $K$  times the procedure used to prepare one copy, but the complexity grows linearly with  $K$ . For instance, the algorithm of Grover leads to a query complexity of  $O(K\sqrt{N})$  for preparing the  $K$ -fold state  $|w\rangle^{\otimes K}$ . In this paper, we investigate the question of whether a more efficient approach exists. We describe a two-phase

algorithm consisting of a preprocessing step that uses  $O(\sqrt{KN})$  queries, after which each copy of  $|w\rangle$  requires only  $O(\sqrt{N/K})$  queries to be prepared. Our result improves upon the previous approach by a factor of  $\sqrt{K}$ , and it is shown to be optimal.

### A. Related work

Our work is based on the *quantum rejection sampling* method, where a state that is easy to prepare (e.g. a uniform superposition) is mapped to a target state by amplitude amplification. This method was pioneered by Grover in [7] and subsequently studied in [8–11]. All of these works (except for [10]) take place in the quantum oracle model and they often require a number of queries that is polynomial in the dimension  $N$  of the state. In the non-oracular setting, the problem of loading an arbitrary vector  $(w_1, \dots, w_N)$  into the amplitudes of a quantum state can be done with a circuit of depth  $O(N)$  and width  $O(\log N)$  [5, 6]. It is possible to use only  $\text{polylog}(N)$  resources for specific cases such as efficiently integrable probability distributions [12–14] (Proposition 4), uniformly bounded amplitudes [15], Gaussian states [16] or probability distributions resulting from a Bayesian network [10]. A different line of work [17–22] studied the preparation of a quantum state that corresponds to the stationary distribution of a Markov chain. These algorithms use Markov chain Monte Carlo methods and quantum walk techniques to obtain a preparation time scaling with the spectral gap. Aharonov and Ta-Shma [17] also showed that the existence of an efficient procedure to convert any circuit into a coherent state encoding the output distribution of that circuit would imply that  $\text{SZK} \subseteq \text{BQP}$ .

The STATE PREPARATION problem is also related to the task of preparing samples from a discrete distribution. We refer the reader to [23–25] for a general introduction on the latter topic. In particular, the IMPORTANCE SAMPLING problem (also called WEIGHTED SAMPLING or  $L_1$  SAMPLING) asks to generate  $K$  independent samples from the probability vector  $(\frac{w_1}{\mathcal{W}}, \dots, \frac{w_N}{\mathcal{W}})$  associated with a non-negative weight vector  $w$ . The *alias method* [26–29] solves this problem with a preprocessing cost of  $O(N)$  operations, and a generating cost

---

\* hamoudi@berkeley.edu

of  $O(1)$  operations per sample. Grover [7, 30] suggested a quadratically faster algorithm for obtaining one sample, based on preparing the state  $|w\rangle$  and measuring it in the computational basis. Our state preparation algorithm extends the work of Grover to the case of generating  $K$  independent samples with a total cost of  $O(\sqrt{KN})$  operations. An alternative quantum algorithm for (approximately) generating  $K$  such samples was proposed before in [31], where it was combined with the stochastic gradient descent method to address the submodular function minimization problem.

## B. Overview

The two parts of our state preparation algorithm are described in Theorems 2 and 3. Combined together, these results lead to the following main theorem.

**Theorem 1.** *There is a quantum algorithm with the following properties. Consider two integers  $1 \leq K \leq N$ , a real  $\delta \in (0, 1)$  and a non-zero vector  $w \in \mathbb{R}_{\geq 0}^N$ . Then, with probability at least  $1 - \delta$ , the algorithm outputs  $K$  copies of the state  $|w\rangle$  and it uses  $O(\sqrt{KN} \log(1/\delta))$  queries to  $w$  in expectation.*

We now provide a high-level description of how the algorithm works. Our starting point is the result from Grover [7] for preparing one copy of the state  $|w\rangle$  in the black-box model. Given an upper bound  $h$  on the largest value  $\max_i w_i$ , this algorithm uses two queries and one controlled rotation to implement a unitary  $U$  such that,

$$U|\vec{0}\rangle = \frac{1}{\sqrt{N}} \sum_{i=1}^N |i\rangle \left( \sqrt{\frac{w_i}{h}} |0\rangle + \sqrt{1 - \frac{w_i}{h}} |1\rangle \right). \quad (2)$$

The state  $|w\rangle|0\rangle$  has amplitude  $\sqrt{\frac{w_i}{Nh}}$  in  $U|\vec{0}\rangle$ , thus it can be extracted by using the amplitude amplification algorithm with  $O\left(\sqrt{\frac{Nh}{W}}\right)$  applications of  $U$  and  $U^{-1}$ . In particular, if the largest coordinate of  $w$  is smaller than  $h = W/K$  then we can prepare one copy of  $|w\rangle$  in time  $O(\sqrt{N/K})$ , and  $K$  copies in time  $O(\sqrt{KN})$ . We use this observation to construct a new circuit  $\mathcal{C}$  (Figure 2) such that  $|w\rangle|0\rangle$  has amplitude at least  $\sqrt{K/N}$  in  $\mathcal{C}|\vec{0}\rangle$ , even if  $w$  contains large coordinates. This circuit uses only two queries to  $w$ , but it requires  $O(\sqrt{KN})$  queries to be constructed during a preprocessing phase that is executed only once (Theorem 2). The preprocessing phase consists first of computing the set  $H$  that contains the positions of the  $K$  largest coordinates in  $w$ , by using a variant of the quantum maximum finding algorithm (Proposition 2). The circuit  $\mathcal{C}$  is then defined to proceed in two stages. First, it prepares a state whose amplitudes depend only on the values in  $\{w_i : i \in H\}$  (Equation (4)). Next, it modifies this state by querying the set  $\{w_i : i \notin H\}$  in a way that is similar to that of  $U$ . The crucial observation is that the values in  $\{w_i : i \notin H\}$

must be smaller than  $W/K$  by definition of  $H$ , thus they can be amplified at a smaller cost. Finally, each copy of  $|w\rangle$  can be obtained by one application of the amplitude amplification algorithm on  $\mathcal{C}$  (Theorem 3).

We show in the next proposition that our algorithm is optimal by a simple reduction from the  $K$ -SEARCH problem.

**Proposition 1.** *Any bounded-error quantum algorithm that can output  $K$  copies of the quantum state  $|w\rangle$  given oracle access to any non-zero vector  $w \in \mathbb{R}_{\geq 0}^N$  must perform at least  $\Omega(\sqrt{KN})$  quantum queries to  $w$ .*

*Proof.* We consider a variant of the  $K$ -SEARCH problem where the objective is to find  $K$  preimages of 1 in an oracle  $\mathcal{O} : [N] \rightarrow \{0, 1\}$  containing at least  $2K$  such preimages. The bounded-error quantum query complexity of this problem is  $\Theta(\sqrt{KN})$  (the proof can easily be derived from [32, Appendix A] for instance). On the other hand, by a coupon collector argument [33], if we prepare and measure in the computational basis  $\Theta(K)$  copies of  $|w\rangle$  where  $w = (\mathcal{O}(1), \dots, \mathcal{O}(N)) \in \{0, 1\}^N$ , then we obtain the positions of at least  $K$  different preimages of 1 with constant success probability. It implies that generating  $\Omega(K)$  such copies requires using at least  $\Omega(\sqrt{KN})$  quantum queries to  $w$ .  $\square$

## II. PRELIMINARIES

### A. Computational model

We use the quantum circuit model over a universal gate set made of the CNOT gate and of all one-qubit gates. We suppose that the real numbers manipulated by our algorithms (such as the coordinates of  $w$ ) can be encoded over  $c$  bits, for a fixed value of  $c$ . In particular, these numbers can be stored in quantum registers of size  $c$ . We also add the three gates described in Figure 1. The indicator gate  $\mathbb{1}_H$  is specified by a subset  $H \subseteq [N]$ . It operates on a Boolean value  $b$  and an index  $i \in [N]$ . The Boolean value  $[i \notin H]$  is equal to 1 if and only if  $i \notin H$ . The query gate  $\mathcal{O}_w$  is specified by the input vector  $w$  to the problem. It operates on an index  $i \in [N]$  and a real  $v$  (encoded over  $c$  bits). Finally, the controlled rotation gate  $\text{Rot}_h$  is specified by a real  $h > 0$  and it operates on a Boolean value  $b$  and a real  $v$ . We refer the reader to [9, 34] and references therein for efficient implementations of the arithmetic gates and controlled rotation gates with a given precision. We will also use  $|\vec{0}\rangle$  in our notations to represent a multi-qubit basis state  $|0\rangle^{\otimes \ell}$  for some  $\ell > 1$ .

The *query complexity* of an algorithm is the number of times it uses the oracle gate  $\mathcal{O}_w$  to access the input  $w$ . If not specified, the total number of elementary gates used by our algorithms will be larger than their respective query complexity by at most a polylog( $N$ ) factor.

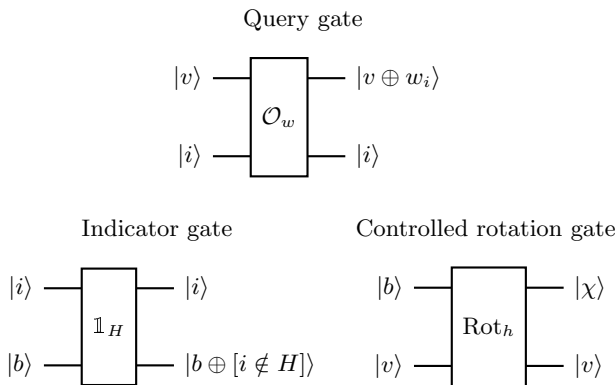


FIG. 1. Three gates used in our circuits. The state  $|\chi\rangle$  is defined as  $\sqrt{\frac{v}{h}}|b\rangle + \sqrt{1-\frac{v}{h}}|1-b\rangle$  if  $0 < v \leq h$ , and  $|b\rangle$  otherwise.

### B. Quantum subroutines

The next algorithm generalizes the quantum minimum finding of Dürr and Høyer [35] to finding  $K$  largest entries in a vector  $w$ . The algorithm succeeds if it outputs a set  $H \subset [N]$  of  $K$  coordinates that dominate all other entries. There is not necessarily a unique choice for  $H$  since different coordinates of  $w$  may be equal.

**Proposition 2** (TOP- $K$  MAXIMUM FINDING – Theorem 3.4 in [36]). *There exists a quantum algorithm with the following properties. Consider two integers  $1 \leq K \leq N$ , a real  $\delta \in (0, 1)$  and a vector  $w \in \mathbb{R}_{\geq 0}^N$ . Then, the algorithm outputs the positions of  $K$  largest entries in  $w$  with success probability at least  $1 - \delta$ , and it performs  $O(\sqrt{KN} \log(1/\delta))$  queries to  $w$ .*

We also need the well-known amplitude amplification algorithm.

**Proposition 3** (AMPLITUDE AMPLIFICATION – Theorem 3 in [37]). *Let  $\mathcal{C}$  be a quantum circuit that prepares the state  $\mathcal{C}|\vec{0}\rangle = \sqrt{p}|\varphi\rangle|0\rangle + \sqrt{1-p}|\varphi^\perp\rangle|1\rangle$  for some  $p \in [0, 1]$  and two unit states  $|\varphi\rangle, |\varphi^\perp\rangle$ . Then, the amplitude amplification algorithm outputs the state  $|\varphi\rangle$  by using  $O(1/\sqrt{p})$  applications of  $\mathcal{C}$  and  $\mathcal{C}^\dagger$  in expectation.*

Finally we will use the next quantum state preparation algorithm that requires having an efficient procedure to compute the partial sum  $\sum_{\ell=i}^j w_\ell$  for any  $1 \leq i \leq j \leq N$ .

**Proposition 4** (STATE PREPARATION BY INTEGRATION – [12–14]). *There is a quantum algorithm with the following properties. Consider an integer  $N$  and a non-zero vector  $w \in \mathbb{R}_{\geq 0}^N$  such that there is a (classical) reversible circuit with  $T$  gates that computes  $\sum_{\ell=i}^j w_\ell$  given  $i \leq j$ . Then, the algorithm outputs  $|w\rangle$  and it uses  $O(T \log N)$  elementary gates.*

## III. MAIN ALGORITHM

We describe in details our state preparation algorithm for preparing  $K$  copies of  $|w\rangle$  given oracle access to  $w = (w_1, \dots, w_N)$ . The first step of the algorithm is a preprocessing phase (Algorithm 1) that constructs a particular circuit  $\mathcal{C}$  described in Figure 2.

**Algorithm 1** Preprocessing phase.

- 1: Compute a set  $H \subseteq [N]$  of the positions of  $K$  largest entries in  $w$  by using the top- $K$  maximum finding algorithm (Proposition 2) with failure probability  $\delta$ .
- 2: Compute  $h = \min_{i \in H} w_i$  and

$$\mathcal{Z} = (N - K)h + \sum_{i \in H} w_i. \quad (3)$$

- 3: Use the state preparation algorithm of Proposition 4 to construct a circuit  $\mathcal{D}$  such that, on input  $|\vec{0}\rangle_{\text{out}}$ , it prepares the state

$$\mathcal{D}|\vec{0}\rangle_{\text{out}} = \sum_{i \in H} \sqrt{\frac{w_i}{\mathcal{Z}}} |i\rangle_{\text{out}} + \sum_{i \notin H} \sqrt{\frac{h}{\mathcal{Z}}} |i\rangle_{\text{out}}. \quad (4)$$

- 4: Output the circuit  $\mathcal{C}$  represented in Figure 2.




FIG. 2. Circuit  $\mathcal{C}$  output at the end of the preprocessing phase (Algorithm 1).

**Theorem 2** (PREPROCESSING PHASE). *Consider two integers  $1 \leq K \leq N$ , a real  $\delta \in (0, 1)$  and a non-zero vector  $w \in \mathbb{R}_{\geq 0}^N$ . Then, Algorithm 1 outputs with probability at least  $1 - \delta$  the description of a quantum circuit  $\mathcal{C}$  such that, on input  $|\vec{0}\rangle$ , it prepares the state*

$$|\psi\rangle = \sqrt{p_w} |w\rangle |0\rangle + \sqrt{1-p_w} |w^\perp\rangle |1\rangle \quad (5)$$

where  $p_w \geq K/N$  and  $|w^\perp\rangle$  is some unit state. The algorithm performs  $O(\sqrt{KN} \log(1/\delta))$  queries to  $w$ . The circuit  $\mathcal{C}$  performs two queries to  $w$  and it uses  $O(K \log N)$  elementary gates.

*Proof.* We assume that the top- $K$  maximum finding algorithm returns a correct set  $H$  at step 1 of Algorithm 1, which is the case with probability at least  $1 - \delta$ . The circuit  $\mathcal{C}$  represented in Figure 2 operates on four registers: *rot* and *ind* that contain a Boolean value, *qry* that contains a real  $v \geq 0$ , and *out* that contains an integer  $i \in [N]$ . The indicator gate  $\mathbb{1}_H$  flips the content

of  $\text{ind}$  when the  $\text{out}$  register contains  $i \notin H$ , which allows the rotation gate  $\text{Rot}_h$  to be activated only when  $i \notin H$ . The gates  $\mathcal{O}_w$  and  $\mathbb{1}_H$  are applied a second time at the end of  $\mathcal{C}$  to uncompute the registers  $\text{qry}$  and  $\text{ind}$ . A simple calculation shows that the final state is  $\mathcal{C}|\vec{0}\rangle = |\vec{0}\rangle_{\text{qry}}|0\rangle_{\text{ind}}|\psi\rangle_{\text{out,rot}}$  where

$$\begin{aligned} |\psi\rangle_{\text{out,rot}} &= \sum_{i \in H} \sqrt{\frac{w_i}{\mathcal{Z}}} |i\rangle_{\text{out}} |0\rangle_{\text{rot}} \\ &+ \sum_{i \notin H} \sqrt{\frac{h}{\mathcal{Z}}} |i\rangle_{\text{out}} \left( \sqrt{\frac{w_i}{h}} |0\rangle_{\text{rot}} + \sqrt{1 - \frac{w_i}{h}} |1\rangle_{\text{rot}} \right) \\ &= \sqrt{\frac{\mathcal{W}}{\mathcal{Z}}} |w\rangle_{\text{out}} |0\rangle_{\text{rot}} + \sqrt{1 - \frac{\mathcal{W}}{\mathcal{Z}}} |w^\perp\rangle_{\text{out}} |1\rangle_{\text{rot}} \end{aligned} \quad (6)$$

for some irrelevant unit state  $|w^\perp\rangle_{\text{out}}$ . In order to lower bound the coefficient  $p_w := \frac{\mathcal{W}}{\mathcal{Z}}$ , we first observe that the smallest value  $h = \min_{i \in H} w_i$  over  $H$  must satisfy  $h \leq \frac{\mathcal{W}}{K}$  since otherwise  $\sum_{i \in H} w_i$  would exceed  $\mathcal{W}$ . Thus,  $p_w^{-1} = \frac{\mathcal{Z}}{\mathcal{W}} = \frac{(N-K)h + \sum_{i \in H} w_i}{\mathcal{W}} \leq \frac{N-K}{K} + 1 = \frac{N}{K}$ .

The algorithm uses  $O(\sqrt{KN} \log(1/\delta))$  queries at step 1 by Proposition 2. The set  $\{w_i : i \in H\}$  can be computed with  $K$  queries, after which steps 2–4 do not need to perform any new query. For any  $i \leq j$ , the partial amplitude sum  $\sum_{\ell=i}^j \langle \ell | \mathcal{D} | \vec{0} \rangle^2$  is equal to  $\frac{1}{\mathcal{Z}} (\sum_{\ell \in H \cap \{i, \dots, j\}} w_i + h \cdot |\{i, \dots, j\} \setminus H|)$ , which can be computed by a classical circuit with  $O(K)$  gates since  $H$  is of size  $K$ . Thus, by Proposition 4, the circuit  $\mathcal{D}$  constructed at step 3 requires  $O(K \log N)$  elementary gates. Finally, the number of gates needed to implement the circuit  $\mathcal{C}$  at step 4 is

dominated by the number of gates needed in  $\mathcal{D}$  since the other gates are included in the computational model (see Section II A).  $\square$

We use the circuit  $\mathcal{C}$  constructed during the above preprocessing phase, together with the amplitude amplification algorithm, to obtain the next state preparation phase that generates one copy of  $|w\rangle$  in time  $O(\sqrt{N/K})$ .

**Theorem 3** (STATE PREPARATION PHASE). *Consider two integers  $1 \leq K \leq N$ , a real  $\delta \in (0, 1)$  and a non-zero vector  $w \in \mathbb{R}_{\geq 0}^N$ . Let  $\mathcal{C}$  denote a quantum circuit obtained with Algorithm 1 on input  $K, \delta, w$  that correctly prepares the state  $|\psi\rangle$  described in Theorem 2. Then, given the description of  $\mathcal{C}$ , one can prepare the state  $|w\rangle$  by using  $O(\sqrt{N/K})$  queries to  $w$  in expectation.*

*Proof.* This is a direct application of the amplitude amplification algorithm (Proposition 3) on  $\mathcal{C}$ , where the complexity is derived from the fact that  $|w\rangle|0\rangle$  has amplitude at least  $\sqrt{K/N}$  in  $|\psi\rangle$  by Theorem 2.  $\square$

#### IV. DISCUSSION

We did not address the precision errors in our analysis. In particular, it can be relevant to replace the controlled rotation gate (which requires to calculate the arcsine function) by the comparison-based circuit defined in [9] that avoids arithmetic. We also restricted ourselves to preparing states with non-negative real amplitudes. Arbitrary phase factors can be introduced by using the techniques discussed in [9, 14].

- 
- [1] F. Magniez, A. Nayak, J. Roland, and M. Santha, Search via quantum walk, *SIAM Journal on Computing* **40**, 142 (2011).
  - [2] A. W. Harrow, A. Hassidim, and S. Lloyd, Quantum algorithm for linear systems of equations, *Physical Review Letters* **103**, 150502 (2009).
  - [3] D. W. Berry, High-order quantum algorithm for solving linear differential equations, *Journal of Physics A: Mathematical and Theoretical* **47**, 105301 (2014).
  - [4] S. Aaronson, Read the fine print, *Nature Physics* **11**, 291 (2015).
  - [5] M. Möttönen, J. J. Vartiainen, V. Bergholm, and M. M. Salomaa, Transformation of quantum states using uniformly controlled rotations, *Quantum Information & Computation* **5**, 467 (2005).
  - [6] M. Plesch and v. Brukner, Quantum-state preparation with universal gate decompositions, *Physical Review A* **83**, 032302 (2011).
  - [7] L. K. Grover, Synthesis of quantum superpositions by quantum computation, *Physical Review Letters* **85**, 1334 (2000).
  - [8] M. Ozols, M. Roetteler, and J. Roland, Quantum rejection sampling, *ACM Transactions on Computation Theory* **5**, 11:1 (2013).
  - [9] Y. R. Sanders, G. H. Low, A. Scherer, and D. W. Berry, Black-box quantum state preparation without arithmetic, *Physical Review Letters* **122**, 020502 (2019).
  - [10] G. H. Low, T. J. Yoder, and I. L. Chuang, Quantum inference on Bayesian networks, *Physical Review A* **89**, 062315 (2014).
  - [11] N. Wiebe and C. Grandade, Can small quantum systems learn, *Quantum Information & Computation* **17**, 568 (2017).
  - [12] C. Zalka, Simulating quantum systems on a quantum computer, *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **454**, 313 (1998).
  - [13] L. K. Grover and T. Rudolph, Creating superpositions that correspond to efficiently integrable probability distributions (2002), [arXiv:quant-ph/0208112](https://arxiv.org/abs/quant-ph/0208112).
  - [14] P. Kaye and M. Mosca, Quantum networks for generating arbitrary quantum states, in *Proceedings of the International Conference on Quantum Information (ICQI)* (2001) p. PB28.
  - [15] A. N. Soklakov and R. Schack, Efficient state preparation for a register of quantum bits, *Physical Review A* **73**, 012307 (2006).

- [16] A. Kitaev and W. A. Webb, Wavefunction preparation and resampling using a quantum computer (2009), [arXiv:0801.0342](https://arxiv.org/abs/0801.0342) [quant-ph].
- [17] D. Aharonov and A. Ta-Shma, Adiabatic quantum state generation, *SIAM Journal on Computing* **37**, 47 (2007).
- [18] R. D. Somma, S. Boixo, H. Barnum, and E. Knill, Quantum simulations of classical annealing processes, *Physical Review Letters* **101**, 130504 (2008).
- [19] P. Wocjan and A. Abeyesinghe, Speedup via quantum sampling, *Physical Review A* **78**, 042336 (2008).
- [20] D. Orsucci, H. J. Briegel, and V. Dunjko, Faster quantum mixing for slowly evolving sequences of Markov chains, *Quantum* **2**, 105 (2018).
- [21] S. Apers, Quantum walk sampling by growing seed sets, in *Proceedings of the 27th European Symposium on Algorithms (ESA)* (2019) pp. 9:1–9:12.
- [22] A. W. Harrow and A. Y. Wei, Adaptive quantum simulated annealing for bayesian inference and estimating partition functions, in *Proceedings of the 31st Symposium on Discrete Algorithms (SODA)* (2020) pp. 193–212.
- [23] L. Devroye, *Non-Uniform Random Variate Generation* (Springer-Verlag, 1986).
- [24] P. Bratley, B. L. Fox, and L. E. Schrage, *A Guide to Simulation*, 2nd ed. (Springer-Verlag, 1987).
- [25] D. E. Knuth, *The Art of Computer Programming, Volume II: Seminumerical Algorithms*, 3rd ed. (Addison-Wesley, 1998).
- [26] A. J. Walker, New fast method for generating discrete random numbers with arbitrary frequency distributions, *Electronics Letters* **10**, 127 (1974).
- [27] A. J. Walker, An efficient method for generating discrete random variables with general distributions, *ACM Transactions on Mathematical Software* **3**, 253 (1977).
- [28] R. A. Kronmal and A. V. Peterson, On the alias method for generating random variables from a discrete distribution, *The American Statistician* **33**, 214 (1979).
- [29] M. D. Vose, A linear algorithm for generating random numbers with a given distribution, *IEEE Transactions on Software Engineering* **17**, 972 (1991).
- [30] L. K. Grover, Rapid sampling through quantum computing, in *Proceedings of the 32nd Symposium on Theory of Computing (STOC)* (2000) pp. 618–626.
- [31] Y. Hamoudi, P. Reberntrost, A. Rosmanis, and M. Santha, Quantum and classical algorithms for approximate submodular function minimization, *Quantum Information & Computation* **19**, 1325 (2019).
- [32] Y. Hamoudi and F. Magniez, Quantum time-space trade-off for finding multiple collision pairs, in *Proceedings of the 16th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC)* (2021) pp. 1:1–1:21.
- [33] Coupon collector’s problem to obtain at least half the coupons (2017), available at <https://math.stackexchange.com/q/2117546/256631>.
- [34] T. Häner, M. Roetteler, and K. M. Svore, Optimizing quantum circuits for arithmetic (2018), [arXiv:1805.12445](https://arxiv.org/abs/1805.12445) [quant-ph].
- [35] C. Dürr and P. Høyer, A quantum algorithm for finding the minimum (1996), [arXiv:quant-ph/9607014](https://arxiv.org/abs/quant-ph/9607014).
- [36] C. Dürr, M. Heiligman, P. Høyer, and M. Mhalla, Quantum query complexity of some graph problems, *SIAM Journal on Computing* **35**, 1310 (2006).
- [37] G. Brassard, P. Høyer, M. Mosca, and A. Tapp, Quantum amplitude amplification and estimation, *Contemporary Mathematics* **305**, 53 (2002).