

# Decision Diagrams for Quantum Measurements with Shallow Circuits

Stefan Hillmich<sup>\*†</sup> Charles Hadfield<sup>\*‡</sup> Rudy Raymond<sup>\*§¶</sup> Antonio Mezzacapo<sup>‡</sup> Robert Wille<sup>†||</sup>

<sup>†</sup>Institute for Integrated Circuits, Johannes Kepler University Linz, 4040 Linz, Austria

<sup>‡</sup>IBM Quantum, IBM T.J. Watson Research Center, Yorktown Heights, NY 10598

<sup>§</sup>IBM Quantum, IBM Japan, 19-21 Nihonbashi Chuo-ku, Tokyo, 103-8510, Japan

<sup>¶</sup>Quantum Computing Center, Keio University, 3-14-1 Hiyoshi, Kohoku-ku, Yokohama, Kanagawa, 223-8522, Japan

<sup>||</sup>Software Competence Center Hagenberg GmbH (SCCH), 4232 Hagenberg, Austria

<sup>\*</sup>Corresponding authors: stefan.hillmich@jku.at, charles.hadfield@ibm.com, rudyhar@jp.ibm.com

**Abstract**—We consider the problem of estimating quantum observables on a collection of qubits, given as a linear combination of Pauli operators, with shallow quantum circuits consisting of single-qubit rotations. We introduce estimators based on randomised measurements, which use decision diagrams to sample from probability distributions on measurement bases. This approach generalises previously known uniform and locally-biased randomised estimators. The decision diagrams are constructed given target quantum operators and can be optimised considering different strategies. We show numerically that the estimators introduced here can produce more precise estimates on some quantum chemistry Hamiltonians, compared to previously known randomised protocols and Pauli grouping methods.

## I. INTRODUCTION

Variational quantum algorithms are based on a quantum-classical optimisation feedback loop, in which a trial parameterised quantum state is prepared on a quantum computer, a target quantum cost function is estimated on it, and a classical optimiser changes the quantum parameters to minimise the target observable. This machinery has been instrumental to find ground state energies of quantum chemistry systems, which are the smallest quantum systems believed to deliver quantum advantage in the field of quantum simulations.

Recent experiments on variational quantum algorithms [1]–[3] have shown that precise estimates of complex quantum operators are essential for a successful execution of the algorithms. A finite single-qubit measurement budget can hinder the performance of the quantum-classical optimisation cycle, by introducing stochastic noise on the quantum cost function to be optimised. This problem is particularly severe for quantum chemistry systems, whose molecular Hamiltonians are composed of a linear combination of Pauli operators that grows, at worst, with the fourth power of the system size [4].

To alleviate the measurement problem, a variety of algorithms have been proposed. They all ultimately aim at obtaining precise estimations of multi-qubit quantum operators, typically given as linear combination of Pauli operators, with the smallest amount of single-qubit measurements.

The idea of using the same single-qubit measurements to estimate grouped Pauli operators that qubit-wise commute, introduced in [1] as *tensor-product basis sets*, is at the core of

several measurements protocols. Some of them promise a reduction in the number of measurements for quantum chemistry systems, exploiting Pauli grouping heuristics, at the expense of an increase depth in the quantum circuits used to prepare the state to be measured [5]–[7]. However increased circuit depths can impair the execution on noisy quantum computers prone to decoherence. Furthermore, even in the fault-tolerant regime, bigger circuit depths can increase the overall runtime of the quantum algorithm. Addressing molecular systems, [8] shows a linear saving in the number of grouped Pauli operators when addressing molecular systems through unitary partitioning of a target Hamiltonian, while [9] finds a cubic reduction if the problem is expressed in plane wave basis, at the expense of a linear increase in circuit depth. Ref. [7] proposes *sorted insertion* to group Pauli operators based on their weights, preprocessing computations linear in the number of qubits and quadratic in the number of Pauli operators. Refs. [5] and [10] exploit simultaneous measurability of partitions of commuting Pauli strings. Exploiting the automated search for symmetries introduced in [11], [12] shows a linear scaling when applied to chemistry problems, again at the expense of increased circuit depth. Ref. [13] considers random Pauli sets, and uses greedy graph colouring algorithms to determine partition of Pauli operators of Hamiltonians, conjecturing a linear saving in number of measurements if arbitrary Clifford operators before measurement are allowed.

Other approaches address the measurement problem while not increasing circuit depths. We refer to this specific case as the *shallow-circuit measurement problem*. It has been tackled so far by formulating it in terms of graph colouring, which was solved with a variety of heuristics [12], [14]. Hybrid architectures made of quantum computers in conjunction with trained neural-network quantum states have been employed to reduce measurement variances [15].

The results contained in this work build on recent techniques based on randomised sequences of single-qubit measurements. A framework for efficiently estimating properties of reduced subsystems of quantum states was introduced in [16]. There, collections of randomised measurement outcomes, labelled *classical shadows* [16], [17], are classically stored to retrieve at a later stage expectation values of local observables.

While this procedure is very well suited for retrieving many generic local observables, the uniform distribution used to draw measurement bases is not optimal in estimating with high precision specific observables such as molecular Hamiltonians. Building on this result, Ref. [18] uses *derandomisation* to deterministically change a sequence of measurement bases drawn at random, with the goal of improving precision in estimating specific sets of Pauli operators.

Improving on uniform distribution sampling, biased randomised measurement protocols can be used to improve estimated precision of given observables. The *locally-biased classical shadows* introduced in [19] are collections of random measurements generated by probability distributions optimised locally at the single-qubit level. While a bias on single-qubit product probability distributions can outperform on uniform random distributions and Pauli grouping heuristics, it still misses on improvements that can come by considering generic measurement probability distributions on a set of qubits. Here we introduce a framework to sample from probability distributions on measurement bases that generalises the local product probability distributions considered in [19]. These probability distributions are generated using decision diagrams, constructed from a target quantum observable, given as a linear combination of Pauli operators.

Decision diagrams are a well-known graph-based data structure used in many disciplines of computer science to enable compact representation of data in many cases. Example applications include binary decision diagrams representing Boolean functions [20], zero-suppressed binary decision diagrams with a focus on sets [21], tagged binary decision diagrams as a combination of both [22],  $\pi$ DDs representing permutations [23], as well as decision diagrams representing quantum states and quantum operations [24]–[28]. At their core, decision diagrams decompose the given data into smaller parts by successively making decisions to remove degrees of freedom, recording these decisions, and exploiting the emergence of parts that are equal. For probability distributions for the measurement problem, decision diagrams provide a natural way to bias the selection of the next measurement basis based on the previous decisions. Decision diagrams have been used for efficient sampling on large sets by applying dynamic programming methods [29].

We propose and discuss different strategies to build decision diagrams and to optimise them in order to reduce the variance of quantum observable estimators that rely on them. We show numerically that they can improve estimation precision on some molecular Hamiltonians, compared to locally biased probability distributions. The implementation we base the results on is available [github.com/iic-jku/dd-quantum-measurements](https://github.com/iic-jku/dd-quantum-measurements).

### Outline of the paper

We introduce the shallow-circuit measurement problem in Section II together with the general probabilistic measurement framework, reviewing the locally-biased approach. We define the decision diagrams and present the main idea in Section III,

followed by strategies to construct decision diagrams given target Hamiltonian operators in Section IV. In Section V, we numerically benchmark the estimators bases on decision diagrams versus existing approaches, on quantum Hamiltonians of increasing sizes, representing quantum chemistry models. We summarize the results obtained and draw conclusions in Section VI. Some technical details on the variances of the estimation by decision diagrams, their implementation, optimisation, and relation with previous work are given in the Appendix of [30].

## II. THE SHALLOW-CIRCUIT MEASUREMENT PROBLEM

This section precisely establishes the problem of shallow-circuit measurement and a general framework of its solutions as search problems over probability distributions. The general framework allows us not only to describe the existing probabilistic solutions, such as, the Pauli grouping via graph colouring and the (locally-biased) classical shadows, but also to derive the new probabilistic solution using decision diagrams that overcome the drawbacks of existing solutions. This requires us to introduce some notation which we do progressively in the following subsection.

### A. Problem Definition

Consider an  $n$ -qubit Hamiltonian

$$H = \sum_{P \in \{I, X, Y, Z\}^n} \alpha_P P \quad (1)$$

with  $\text{poly}(n)$  number of real coefficients  $\alpha_P$  acting on a quantum processor, we say  $P$  is a *Pauli operator* consisting of  $n$  *single-qubit* Pauli operators and write  $P = (\otimes_{i \in [n]} P_i) \in \{I, X, Y, Z\}^n$  where  $I, X, Y, Z$  are  $2 \times 2$  Pauli matrices. The Hilbert space is  $\mathcal{H} := (\mathbb{C}^2)^{\otimes n} = \mathbb{C}^{2^n}$ . Let  $\mathcal{D}(\mathcal{H})$  denote the space of quantum densities and fix some unknown  $\rho \in \mathcal{D}(\mathcal{H})$ . Our task is to estimate  $\text{Tr}(H\rho)$  to some additive accuracy  $\varepsilon > 0$ .

We restrict our attention to algorithms for  $\text{Tr}(H\rho)$  which are compatible with quantum processors of the current generation (see Remark 1). Specifically we assume that the measurement bases in which we may measure  $\rho$  are of the form  $B = \otimes_{i \in [n]} B_i$  where  $B_i = x_i X + y_i Y + z_i Z$  and  $x_i^2 + y_i^2 + z_i^2 = 1$ . If we then prepare  $\rho$  many times, say  $S \in \mathbb{N}$ , and for each  $s \in [S]$  choose a measurement basis  $B^{(s)}$  in which to measure  $\rho$  we can estimate, with progressively increasing accuracy, the value of  $\text{Tr}(H\rho)$ . We will make two further assumptions: the choice of  $B^{(s)}$  is independent of  $B^{(s')}$  for  $s' < s$ ; any such basis  $B$  is a Pauli operator  $B \in \mathcal{P}^n$  where  $\mathcal{P} = \{X, Y, Z\}$ . The *shallow-circuit measurement problem* is how to best choose the measurement bases in order to estimate  $\text{Tr}(H\rho)$  within accuracy  $\varepsilon$  with as few preparations of  $\rho$  as possible.

*Remark 1.* Quantum computers of the current generation are not fault-tolerant. This naturally imposes restrictions on the type of measurement schedules we should consider. The precision with which we can implement non-entangling gates (single-qubit unitaries) is significantly greater than that of

entangling gates (multi-qubit unitaries). It is therefore natural to consider the restriction of the problem of estimating  $\text{Tr}(H\rho)$  to the constrained setting of measuring  $\rho$  in bases which do not require entangling gates. In general, reducing entangling gates is beneficial in computational time even for fault-tolerant quantum computers.

*Remark 2.* The two further assumptions have been chosen to ease the exposition. The first assumption, that the choice of a basis is independent of previous choices is important. It prevents the situation where a potentially expensive classical computation is performed between each shot. This is similar to the distinction between adaptive and non-adaptive scheduling in annealing schedules. The second assumption, that measurement bases are full-weight Pauli operators, allows us, in the following section, to talk about probability distributions which are over finite sets. This assumption should be considered minor.

### B. The General Probabilistic Measurement Framework

We can solve the shallow-circuit measurement problem by viewing it as the problem of how to best pick a probability distribution  $\beta$  over the measurement bases  $\mathcal{P}^n$ . In order to see the relationship we fix some notation. First, for a fixed Pauli operator  $P$ , let

$$\text{Cover}(P) := \{B \in \mathcal{P}^n \mid B_i = P_i \text{ whenever } P_i \neq I\}. \quad (2)$$

This is the set of measurement bases which allow us to estimate  $\text{Tr}(P\rho)$ . (We shall say that any such  $B$  covers  $P$ .) Next, if  $B = \otimes_{i \in [n]} B_i \in \mathcal{P}^n$ , then measuring qubit  $i$  in the  $B_i$  basis returns an eigenvalue  $\mu(B, i) \in \{\pm 1\}$ . For a subset  $A \subseteq [n]$  let us declare  $\mu(B, A) := \prod_{i \in A} \mu(B, i)$  with the convention that  $\mu(P, \emptyset) = 1$ . If we set  $\text{supp}(P) := \{i \in [n] \mid P_i \neq I\}$  then we find that  $\text{Tr}(P\rho)$  is estimated by  $\mu(B, \text{supp}(P))$  whenever  $B$  covers  $P$ . Penultimately, let  $\mathbb{1}_\Omega$  represent the indicator function of a set  $\Omega$ . That is  $\mathbb{1}_\Omega(x)$  returns 1 if  $x \in \Omega$  and 0 if  $x \notin \Omega$ . Finally, if  $\beta : \mathcal{P}^n \rightarrow \mathbb{R}^+$  is a probability distribution then the probability that a basis  $B$  is chosen such that  $\text{Tr}(P\rho)$  may be estimated is

$$\zeta(P, \beta) = \sum_{B \in \mathcal{P}^n} \mathbb{1}_{\text{Cover}(P)}(B) \cdot \beta(B) = \sum_{B \in \text{Cover}(P)} \beta(B). \quad (3)$$

The shallow-circuit measurement problem reduces to finding  $\beta : \mathcal{P}^n \rightarrow \mathbb{R}^+$  which minimizes the variance of the estimator  $\nu$  produced by Algorithm 1.

Let us say that  $\beta$  is *compatible* with the Hamiltonian  $H$  if  $\zeta(P, \beta) > 0$  whenever  $\alpha_P \neq 0$ . In [30, Appendix A] we prove two results: we show that this algorithm returns an unbiased estimator for  $\text{Tr}(H\rho)$  provided  $\beta$  is compatible with  $H$ ; we also calculate the variance of the estimator. We record these statements below.

**Lemma 1.** *If  $\beta$  is compatible with  $H$  then the first moment of the estimator  $\nu^{(s)}$  from Algorithm 1 is*

$$\mathbb{E}(\nu^{(s)}) = \sum_P \alpha_P \text{Tr}(P\rho). \quad (4)$$

---

**Algorithm 1** Shallow-measurement estimation of  $\text{Tr}(H\rho)$  given  $\beta$ .

---

**for** shot  $s \in [S]$  **do**

  Prepare  $\rho$

  Select basis  $B \in \mathcal{P}^n$  from  $\beta$ -distribution

**for** qubit  $i \in [n]$  **do**

    Measure qubit  $i$  in basis  $B_i$  giving  $\mu(P, i) \in \{\pm 1\}$

  Estimate observable expectation

$$\nu^{(s)} = \sum_P \alpha_P \cdot \frac{\mathbb{1}_{\text{Cover}(P)}(B)}{\zeta(P, \beta)} \cdot \mu(B, \text{supp}(P))$$

**return**  $\nu = \frac{1}{S} \sum_{s \in [S]} \nu^{(s)}$

---

**Lemma 2.** *If  $\beta$  is compatible with  $H$  then the second moment of the estimator  $\nu^{(s)}$  from Algorithm 1 is*

$$\mathbb{E}(\nu^{(s)} \nu^{(s)}) = \sum_{P, Q} \alpha_P \alpha_Q g(P, Q, \beta) \text{Tr}(PQ\rho) \quad (5)$$

where

$$g(P, Q, \beta) := \frac{1}{\zeta(P, \beta)} \cdot \frac{1}{\zeta(Q, \beta)} \cdot \sum_{B \in \mathcal{P}^n} \mathbb{1}_{\text{Cover}(P)}(B) \cdot \mathbb{1}_{\text{Cover}(Q)}(B) \cdot \beta(B). \quad (6)$$

**Proposition 1.** *If  $\beta$  is compatible with  $H$  then the estimator in Algorithm 1 is an unbiased estimator of  $\text{Tr}(H\rho)$  and has variance*

$$\text{Var}(\nu) = \frac{1}{S} \left( \left( \sum_{P, Q} \alpha_P \alpha_Q g(P, Q, \beta) \text{Tr}(PQ\rho) \right) - (\text{Tr}(H\rho))^2 \right). \quad (7)$$

The task is to prescribe  $\beta$  in order to obtain a measurement schedule with small variance. In the subsequent sections we show how existing solutions, i.e., the Pauli grouping and (locally-biased) classical shadows can be cast as instances of choosing specific probability distributions and point out their drawbacks. We then show how decision diagrams offer an excellent prescription for prescribing  $\beta$ .

*Remark 3.* Although we have derived a formula for the variance of the estimator, we would like to point out the role of  $\zeta(P, \beta)$  to the variance. Clearly, the square root of the variance (or, the standard deviation of)  $\nu^{(s)}$  at Algorithm 1 is at most  $\max_B \sum_P \mathbb{1}_{\text{Cover}(P)}(B) \cdot |\alpha_P| / \zeta(P, \beta)$ . Thus, a rule of thumb to minimise the variance is by choosing the distribution  $\beta$  so that  $\max_P 1/\zeta(P, \beta)$  is minimised. This intuition is indeed true as we show later for problems of estimating a set of Pauli terms: our inconfidence of the correctness of the estimator is proportional to  $\max_P 1/\zeta(P, \beta)$ .

*Remark 4.* How much knowledge of  $\rho$  should we allow ourselves in order to prescribe  $\beta$ ? The obvious application of our subroutine is in the context of finding the ground energy

of a Hamiltonian using variational quantum eigensolvers. In this context, it is natural to assume we already have an ansatz for the ground state hence we may assume  $\rho$  is always *close to* the ground state of the system. Quantifying this phrase is subtle, however in order to illustrate the strength of our proposal, it is not unreasonable to do the following: Prescribe the distribution  $\beta$  using *no* knowledge of  $\rho$  and then calculate the variance of  $\nu$  when given access to the true ground state of the system. Already in this regime, we show significant improvements over existing protocols.

### C. Existing Solutions and Drawbacks

In the following two subsections, we review briefly some existing solutions to the shallow measurement problem and their drawbacks that motivate us to develop a new framework with decision diagrams. It is useful to observe that they can be easily seen as instances of Algorithm 1.

1) *Pauli Grouping via Graph Colouring*: Pauli grouping via graph colouring has been used experimentally in [1]. A detailed explanation may be found in [19, Section A.2]. The core idea is to group the Pauli operators  $\{P\}_{\alpha_P \neq 0}$  occurring in  $H$  into  $K$  collections and assign one measurement basis  $B^{(k)} \in \mathcal{P}^n$  to each collection  $k \in [K]$  such that  $B^{(k)}$  allows all Pauli operators in the  $k^{\text{th}}$  collection to be estimated. That is, if  $P$  belongs to the  $k^{\text{th}}$  collection then  $B^{(k)} \in \text{Cover}(P)$ .

The grouping is performed by colouring a specific graph using any graph-colouring heuristic. The graph is constructed first by assigning vertices to each Pauli operator appearing in the Hamiltonian. Second, if two vertices represent Pauli operators  $P = \otimes_{i \in [n]} P_i$ , and  $Q = \otimes_{i \in [n]} Q_i$ , then an edge is added to the graph precisely when there exists a qubit  $i \in [n]$  such that  $P_i, Q_i \in \mathcal{P}$  and  $P_i \neq Q_i$ . It follows that any vertices with the same colour may be assigned a single measurement basis. Graph colouring heuristics lead to collections for which  $K \ll |\mathcal{P}^n|$ .

The best assignment of the probability  $\beta: \{B^{(k)}\}_{k \in [K]} \rightarrow \mathbb{R}^+$  has not been rigorously studied. In [1], the measurement bases  $\{B^{(k)}\}_{k \in [K]}$  are effectively sampled uniformly, but this is due to hardware considerations. In [19, Section A.2] an improved sampling is proposed which is based on the  $\ell_1$  weight of the coefficients of the Pauli operators appearing in each collection. This proposal may be improved slightly by observing that some Pauli operators may be assigned to several collections. That is, sometimes, several bases from  $\{B^{(k)}\}_{k \in [n]}$  may be used to estimate a single Pauli observable. This will decrease slightly the variances obtained in [19, Tables 1, 2].

The main drawback of this approach is observational. That is, on Hamiltonians thus far studied in the literature, the variances are large compared to other proposals. This is despite the preprocessing steps to solve a graph colouring problem to reduce the choices of measurement bases. This may be caused by two reasons. First, the Pauli grouping stage makes no reference to the coefficients  $\{\alpha_P\}$ . Second, the current proposals may not be optimally assigning the distribution over  $\{B^{(k)}\}_{k \in [K]}$ .

2) *Locally-Biased Classical Shadows*: One method of choosing the distribution  $\beta: \mathcal{P}^n \rightarrow \mathbb{R}^+$  has been proposed in [19] and is called locally-biased classical shadows (LBCS). It may be seen as an extension of a proposal to perform tomography, called classical shadows using random Pauli measurements [16], to the problem of estimating a single observable. LBCS fits precisely into the framework discussed in the preceding subsection. First, the class of distributions from which  $\beta$  may be chosen is restricted: only *product* probability distributions are considered, which we may write as  $\beta = \prod_{i \in [n]} \beta_i$  where  $\beta_i: \mathcal{P} \rightarrow \mathbb{R}^+$  is the probability distribution for choosing to measure the  $i^{\text{th}}$  qubit in a basis  $B_i \in \mathcal{P}$ . Second, after assuming this restriction on the class of probability distributions, the choice of  $\beta$  is made by optimising a convex cost function associated with the Hamiltonian:

$$\text{cost}_{\text{diag}}(\beta) := \sum_P \alpha_P^2 \frac{1}{\prod_{i \in \text{supp}(P)} \beta_i(P_i)} \quad (8)$$

A motivation for this choice of cost function can be obtained from [30, Appendix A]. Note that since  $\beta$  is assumed to be a product distribution, the denominator  $\prod_{i \in \text{supp}(P)} \beta_i(P_i)$  is precisely  $\zeta(P, \beta)$ . In [19], it was shown that this method leads to significant reductions in the variance of energy estimation in the context of quantum chemistry over the method of Pauli grouping via graph colouring.

*Remark 5.* Classical shadows, as originally proposed, may be seen as the uniform version of LBCS. That is,  $\beta_i(P) = \frac{1}{3}$  for  $P \in \mathcal{P}$ . A recent result derandomises this idea and shows promise for estimating single observables [18]. Derandomisation should also prove useful to the general framework of decision diagrams. See [30, Appendix F] where the main result of [18] is generalised to the setting of decision diagrams.

Both Classical Shadows and LBCS appear to be attractive when Pauli operators in the Hamiltonian are low-weight. That is, when  $|\text{supp}(P)|$  is small relative to  $n$ , since in this case the denominator of Eq. (8) remains relatively small. This intuition leads to the following setup showing a shortcoming of LBCS. Consider the toy Hamiltonian  $H = \otimes_{i \in [n]} X_i + \otimes_{i \in [n]} Z_i$ . LBCS would assign, for each qubit, the probability  $1/2$  to measure it either in the  $X$  or the  $Z$  basis. The lack of correlation in these choices implies that the only two bases  $X^n$  and  $Z^n$  which are useful for energy estimation are rarely chosen. In the parameter  $n$  it is with exponentially vanishing probability that such bases are chosen. Pauli grouping via graph colouring would perform much better on this example. This example motivates the search for distributions  $\beta$  from a wider class of probabilities. Although the preceding Hamiltonian is a toy example, there does occur Hamiltonians with similar structure. For example, for the purpose of self-testing quantum devices, the Hamiltonian  $H = X_1 Z_2 \cdots Z_n + \frac{1}{n-1} \sum_{i \geq 2} Z_1 X_i$  is evaluated to test Bell-type inequalities on star graphs [31], [32].

### III. ESTIMATORS OF QUANTUM OBSERVABLES WITH DECISION DIAGRAM

This section introduces a new type of decision diagrams. Any such decision diagram (DD) provides a compact representation of a probability distribution over full-weight Pauli operators associated with a given Hamiltonian. To this end, we first briefly review the main idea of decision diagrams in general, before we define and detail the proposed type. Based on that, we show how the new decision diagram can be used to improve the sampling process when measuring quantum Hamiltonians on quantum processors of the current generation. Overall, this section provides the motivation and main idea, while details on the technical implementation are covered in the next section.

#### A. Decision Diagrams in General

Decision diagrams are a tried and tested data structure in many areas of computer science to provide a compact representation of entities in various domains. Example applications include binary decision diagrams representing conventional Boolean functions [20], zero-suppressed binary decision diagrams with a focus on sets [21], tagged binary decision diagrams as a combination of both [22], and  $\pi$ DDs representing permutations [23]. Also in the domain of quantum computing, decision diagrams representing quantum states and quantum operations received interest [24]–[28] and found application, e.g., in the synthesis [24], [33], simulation [26], [34], and verification [25], [35] of quantum circuits.

The common idea of all decision diagram-based representations is to decompose a given original representation (e.g., of a Boolean function or quantum state) in a structured fashion that recognizes and exploits redundancies of the decomposed data in order to provide a more compact representation. The repeatedly conducted decompositions are represented by means of a directed and acyclic multi-graph, where vertices represent the decomposed data and redundancy is exploited through *shared vertices*.

**Example 1.** Consider a Boolean function  $f : \{0, 1\}^4 \rightarrow \{0, 1\} = \bar{x}_1\bar{x}_2\bar{x}_3x_4 + \bar{x}_1x_2x_3\bar{x}_4 + x_1\bar{x}_2x_3\bar{x}_4 + x_1x_2\bar{x}_3x_4$ . A straightforward complete representation of this function would require the representation of a total of  $2^4$  input-output mappings, e.g., in terms of a truth table. Encoding the function as decision diagram results in a graph with only 9 nodes as illustrated in Figure 1. Here, the overall function  $f$  is first decomposed with respect to variable  $x_1$  into two sub-functions  $f_6$  (assuming  $x_1 = 0$ ) and  $f_5$  (assuming  $x_1 = 1$ ). This is recursively continued for all remaining variables until only terminals 0 and 1 result. Whenever this decomposition yields equivalent (and, hence, redundant) sub-functions (as it is the case, e.g., for  $f_4$ ), the sub-function is represented by a single shared node only—providing a more compact representation.

*Remark 6.* Albeit decision diagrams are *directed* graphs, edges in illustrations do not include arrow tips since, by convention,

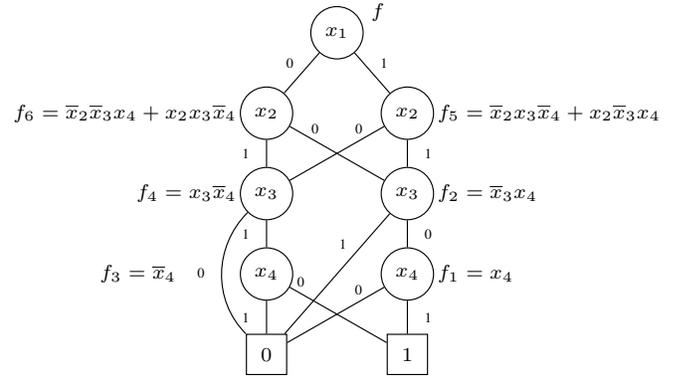


Figure 1: Decision diagram representing the Boolean function  $f = \bar{x}_1\bar{x}_2\bar{x}_3x_4 + \bar{x}_1x_2x_3\bar{x}_4 + x_1\bar{x}_2x_3\bar{x}_4 + x_1x_2\bar{x}_3x_4$

the direction is fixed—most commonly from the top strictly to the bottom.

#### B. Proposed Decision Diagram

In this work, we propose a type of decision diagrams aiming for a compact representation of a probability distribution over a given Hamiltonian. To this end, we first start by providing the definition of the proposed type:

**Definition 1.** The decision diagram we propose is a rooted directed acyclic multi-graph  $G = (V, E)$  such that all maximal directed paths consist of precisely  $n$  edges. Each edge  $e \in E$  is equipped with two pieces of data: A traceless Pauli operator  $B(e) \in \mathcal{P}$  and a weight  $w(e) \in (0, 1]$  such that

- 1) for each vertex  $v \in V$ , there is at most one out-going edge for each traceless Pauli operator, and
- 2) for each vertex  $v \in V$  and outgoing edges  $e \in \text{out}(v)$ , the weights are probabilistic, i.e.,  $\sum_{e \in \text{out}(v)} w(e) = 1$  (therefore, each vertex except for the terminal has at least one out-going edge).

Having this structure, the edge weights in the decision diagram provide the probabilities by which a random walk should follow each edge. Multiplying the edge weight along a path gives the probability of encountering this path in a random walk. Intuitively, the probability of a path corresponds to the weight of the Pauli operators it covers. More precisely, the sum of the absolute values of the coefficients  $|\alpha_P|$  with  $P$  covered by the path is used as relative probability and is encoded in the decision diagram. An example illustrates the idea.

**Example 2.** Consider the Hamiltonian  $H$  for the hydrogen molecule  $H_2$  with 4 qubits and Bravyi-Kitaev encoding, namely:

$$\begin{aligned} H = & -0.811IIII + 0.120IZII - 0.045XZ XI + 0.045XIXZ \\ & + 0.045XIXI - 0.045XZ XZ + 0.120IZIZ + 0.172ZIII \\ & - 0.225IZZZ - 0.228ZZII + 0.172IIZI + 0.168ZIZI \\ & + 0.166ZZZZ + 0.166ZZZI + 0.174ZIZZ \end{aligned}$$

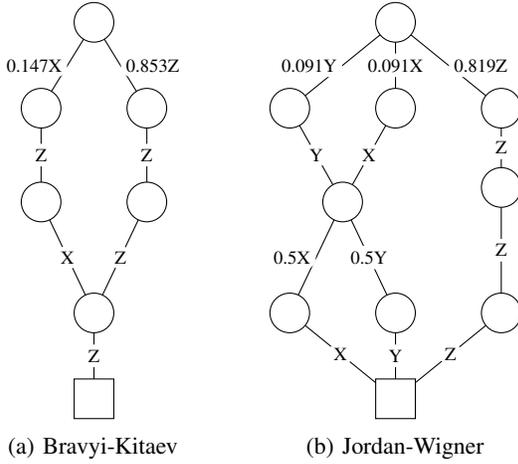


Figure 2: Decision diagrams for  $H_2$  with 4 qubits

The full-weight terms  $XZZX$  or  $ZZZZ$  cover every Pauli term in the Hamiltonian, hence, the corresponding decision diagram encoding the probability distribution only has to include these two terms.

Figure 2a illustrates a compatible decision diagram that has only two maximal paths  $XZZX$  and  $ZZZZ$ . The edges are labelled with a probability (edge weights of 1 are omitted for the sake of readability). Using the decision diagram, one can obtain the respectively desired probabilities by traversing the decision diagram starting at the top-most node and following the edges strictly downwards until the terminal vertex (depicted as rectangle) is reached. Generating measurements from this decision diagram will result in  $XZZX$  with a probability of 0.147 or  $ZZZZ$  with 0.853.

In Figure 2a, the encoded terms end with the  $Z$  operator regardless of previous choices, hence the last decision is presented by a single vertex and edge. This may seem like a small gain, but generally larger instances have more potential for sharing.

In a similar fashion, a decision diagram representing the Jordan-Wigner encoding can be generated—yielding the structure as shown in Figure 2b.

### C. Sampling Using the Proposed Decision Diagrams

The definition in the previous subsection does not make explicit reference to the prescribed Hamiltonian. In order to use one such instance of these decision diagrams we augment the definition with the following

**Definition 2.** The decision diagram is called compatible with the Hamiltonian  $H$  if all Pauli observables  $P$  with  $\alpha_P \neq 0$  are able to be estimated. Precisely, if  $\alpha_P \neq 0$  then we require at least one directed path  $(e_1, \dots, e_n)$  such that  $B \in \text{Cover}(P)$  where  $B$  is the full-weight Pauli operator  $\otimes_{i \in [n]} B(e_i)$ .

Let us make two observations that bring decision diagrams into the probabilistic setup of the shallow measurement problem as explained in the previous section. First, the decision diagram provides a probability distribution over full-weight

Pauli operators  $\beta : \mathcal{P}^n \rightarrow \mathbb{R}^+$ . For  $B \in \mathcal{P}^n$  such that  $B = \otimes_{i \in [n]} B(e_i)$  for some maximal directed path  $(e_1, \dots, e_n)$  then we set  $\beta(B) = \prod_{i \in [n]} w(e_i)$ . If no such maximal directed path exists, then we set  $\beta(B) = 0$ . Condition 2 in Definition 1 ensures that  $\sum_B \beta(B) = 1$ . Indeed, using the decision diagrams proposed above, samples can be drawn by performing a random walk. Starting at the root vertex, a successor vertex is randomly selected according to the weights on the out-going edges. This process is repeated at the selected vertex until the terminal vertex is reached. Second, if the decision diagram is compatible with  $H$  then, for any quantum density  $\rho \in \mathcal{D}(\mathcal{H})$ , the estimator  $\nu$  of Algorithm 1 is an unbiased estimator of the energy. That is, Lemma 1 establishes  $\mathbb{E}(\nu) = \text{Tr}(H\rho)$ .

**Example 2** (continued). Consider again Figure 2a. Sampling from this decision diagram, one starts at the root vertex and randomly chooses the  $X$  or  $Z$  edge, according to the edge weights. Continuing from successor vertex of the chosen edge, the remaining decisions are fixed since each following vertex only has one out-going edge, again, resulting in either  $XZZX$  or  $ZZZZ$ .

Decision diagrams provide a more powerful way to solve the shallow measurement problem. Indeed Pauli grouping via graph-colouring from Section II-C may be seen as one instance of a decision diagram according to Definition 1 and 2. Therefore any Pauli term which is covered by multiple bases under that proposal will be estimated more often under the framework here. Also, LBCS from Section II-C is a very simplistic instance of a decision diagram. We establish the decision diagram link in [30, Appendix C]. As previously commented, LBCS (which observationally is better than Pauli grouping) suffers from the lack of correlation between choices of measurement bases on each qubit. The more general decision diagram framework presented here allows such correlated choices.

The larger class of distributions therefore allows us to ultimately reduce the variance associated with our estimator. Importantly, our proposal for building decision diagrams also leads to an efficient proposal for assigning weights locally such that an attractive distribution  $\beta$  is ultimately found. We reemphasize the possibility that such an algorithm could be extended further with the techniques of derandomisation.

### D. Optimising the Probability Distribution

It is a computationally difficult problem to find the optimal  $\beta$ . This would be the *minimum-variance unbiased estimator* (MVUE) over all such distributions  $\beta : \mathcal{P}^n \rightarrow \mathbb{R}^+$ . It would be interesting to understand how close decision diagrams get to approximating the MVUE. Nevertheless, following the method of Lagrange multipliers used in [19] to optimise the probability distribution of LBCS, we can derive similar iterative procedure to fine tune the probability distribution  $\beta$  based on the diagonal cost function described in [30, Appendix E]. The computational cost of the iterative updates is proportional to the size of the decision diagram.

---

**Algorithm 2** Construction of a decision diagram (DD) from Hamiltonian  $H$ 

---

Take absolute values of coefficients in  $H$   
Merge compatible terms to get reduced positive Pauli list  $\mathcal{R}(H)$  ▷ Preprocessing  
**for** Each term and coefficient in  $\mathcal{R}(H)$  **do** ▷ Initialisation of DD  
    Take existing path covering the longest prefix of term  
    Create new edges for remaining Pauli operators up to the last  
    Create edge to terminal with the last Pauli op and coefficient as edge weight  
**for** Vertex in decision diagram in breadth-first order from terminal **do** ▷ Normalisation of DD  
    Calculate sum of weights on out-going edges  
    Divide weights on out-going edges by sum and multiply sum to in-coming edge weights  
**for** Vertex in decision diagram in breadth-first order from terminal **do** ▷ Merge equivalent vertices in DD  
    Calculate hash of vertex and if equivalent vertex exists, merge both  
Remove identities in DD  
    Replace “lonely” identity edges with virtual edges  
    Remove identity edges where other edge with same source and target exists  
    Merge targets of identity edges with target vertices of other edge  
**for** Vertex in decision diagram in breadth-first order from terminal **do** ▷ Merge equivalent vertices in DD  
    Calculate hash of vertex and if equivalent vertex exists, merge both  
**return** decision diagram

---

A direct implication of the use of probability-optimised decision diagrams is to generalise and improve previous results (Theorem 3 in [36] and Theorem 1 in [18]) on estimating the expectation values of a collection of Pauli operators (e.g., for partial tomography [37]) thanks to the ability to compute  $\zeta(P, \beta)$  efficiently from a decision diagram. For example, the error bounds of previous results provide non-trivial bounds when the size of  $\text{supp}(P)$  for all Pauli  $P$  is small (or, low-weight Paulis), while those of ours can give non-trivial bounds even when some of the Paulis are of full weight. We give the details in [30, Appendix F] because the focus of this paper is on different topics of measuring quantum Hamiltonians.

#### IV. EFFICIENT CONSTRUCTION OF THE PROPOSED DECISION DIAGRAMS

The decision diagrams as introduced in the previous section promise to give suitable probability distribution. Still, the question remains how to efficiently transform the Hamiltonian consisting of coefficients and Pauli terms into a decision diagram. In the following, we describe the main ideas. The full implementation is available at [github.com/iic-jku/dd-quantum-measurements](https://github.com/iic-jku/dd-quantum-measurements).

In a high-level view, the construction of the proposed decision diagram consists of multiple steps as shown in Algorithm 2. The first stage is a preprocessing step of the Hamiltonian to reduce the number of identity-terms. The second stage is the initialisation and refinement of the decision diagram. This second stage has several steps. Only after these steps have been performed are we guaranteed that the decision diagram conforms to both Definition 1 and Definition 2. These steps include normalising the information present in the preprocessed Hamiltonian in order to maximise sharing and also removing identity-edges through *merging*. The following paragraphs explain the individual steps in more detail.

*Preprocessing:* Consider the Hamiltonian as presented in Eq. (1). Immediately, we remove the term  $\alpha_{I^n} I^n$  as this term does not need to be estimated using the quantum processor. We also map coefficients to their absolute values ( $\alpha_P \mapsto |\alpha_P|$ ) giving what we will call the *positive Pauli list*. Then, compatible Pauli terms in this positive Pauli list are merged to reduce the number of paths in the initial decision diagram. More precisely, for each of Pauli terms the number of compatible terms is determined. The Pauli term  $P_{\text{high}}$  with the highest number of compatible terms  $n_{\text{comp}}$  is subsequently merged into these compatible terms and the fraction  $\frac{\alpha_{P_{\text{high}}}}{n_{\text{comp}}}$  is added to the coefficient of each compatible term. This merging procedure is repeated until no further merging is possible. This preprocessing provides what we shall refer to as the *reduced positive Pauli list* and shall denote it by  $\mathcal{R}(H)$ .

*Initialisation of DD:* From the reduced positive Pauli list  $\mathcal{R}(H)$ , an initial decision diagram whose maximal paths are all of length  $n$  is constructed. Each term in  $\mathcal{R}(H)$  is associated with a unique maximal path and the coefficients of  $\mathcal{R}(H)$  are assigned to the final edges in the respective path, i.e., to the edge pointing to the terminal vertex.

*Normalisation of DD:* Afterwards the edge weights are normalised such that the sum of weight of out-going edges equals 1. This decision diagram at this stage has sharing for common prefixes (but not suffixes) and at this point may include edges with the identity operator.

**Example 3.** Consider the 4 qubit Hamiltonian for the hydrogene molecule in Jordan-Wigner encoding:

$$\begin{aligned} H = & -0.810IIII + 0.045YYXX + 0.045YYYY \\ & + 0.045XXXX + 0.045XXYY + 0.172ZIII \\ & - 0.225IZII + 0.172IIZI - 0.225IIIZ \\ & + 0.120ZZII + 0.168ZIZI + 0.166ZIIZ \\ & + 0.166IZZI + 0.174IZIZ + 0.120IIZZ \end{aligned}$$

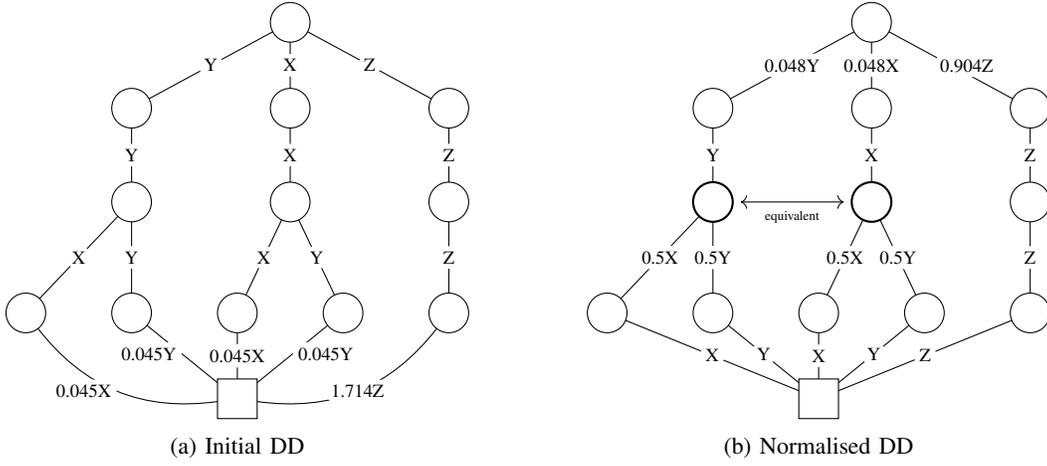


Figure 3: First steps of decision diagram construction from reduced positive Hamiltonian of  $H_2$  JW (4 qubits)

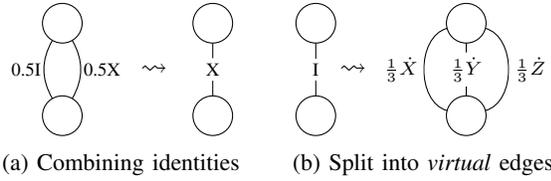


Figure 4: Removing identity edges from the decision diagram

The reduced positive Pauli list from  $H$  is

$$\mathcal{R}(H) = 0.045YYXX + 0.045YYYY + 0.045XXXX + 0.045XXYY + 1.714ZZZZ.$$

Figure 3a illustrates the initial decision diagram created from  $\mathcal{R}(H)$  and Figure 3b illustrates the same decision diagram with normalised edge weights.

Before considering removing potentially remaining identity edges in the decision diagram, functionally equivalent vertices are merged. Two vertices are equivalent if they have the same successors considering the Pauli operator and respective weight. Given a suitable hash function, finding equivalent nodes is linear in the number of nodes [28].

**Example 3** (continued). Consider again Figure 3b. The two vertices highlighted by a bold border are equivalent and, thus, are combined to exploit the redundancy. This results in the decision diagram as already shown before in Figure 2b.

After the previous steps, the decision diagram may still have identity edges, which have to be removed to get a proper probability distribution over the Hamiltonian. The potentially remaining identity edges are removed in three steps.

The first two steps are local operations. Given two fixed vertices  $u, v \in V$  the following checks are performed:

- 1) For fixed  $u, v \in V$ , if there is an edge  $u \xrightarrow{I} v$  and any  $u \xrightarrow{\{X,Y,Z\}} v$ : Remove  $u \xrightarrow{I} v$  and add the weight to the remaining edge from  $u \xrightarrow{\{X,Y,Z\}} v$  with the smallest weight. Figure 4a illustrates this step in an example with one further edge.

- 2) Again, for fixed  $u, v \in V$ , if there is only one edge  $u \xrightarrow{I} v$  (referred to as “lonely” in Algorithm 2): Split  $u \xrightarrow{I} v$  into *virtual edges* (denoted by a dot above the operator)  $u \xrightarrow{\dot{X}} v$ ,  $u \xrightarrow{\dot{Y}} v$ , and  $u \xrightarrow{\dot{Z}} v$  with weights  $\frac{1}{3}$  each. Figure 4b illustrates this step.

The remaining identity edges cannot be removed by only considering individual pairs of nodes, but require a more global approach. Recall that at this point there are no two vertices with only an identity edge between them. So for  $u \xrightarrow{I} v$  and any  $u \xrightarrow{\{X,Y,Z\}} v'$  we merge  $v$  into  $v'$  and adjust the weights accordingly. More precisely, the merging is handled by checking the following list for each out-going edge of  $v$ :

- 1) If the target vertex  $v'$  does not have an out-edge with the same Pauli operator as the currently considered out-edge of  $v$ , add this edge to  $v'$ .
- 2) If the currently considered out-edge of  $v$  and the out-edge of  $v'$  with the same operator point to same vertex (which may be the terminal vertex), the weights stay the same.
- 3) Otherwise the merging process has to recurse to merge the successors of  $v$  and  $v'$  with corresponding Pauli operators.

During the merging of two edges, the resulting edge is only *virtual* if both previous edges were virtual. Performing the merging process from the terminal vertex upwards ensures that the recursive applications never encounter an identity edge. At this point there may be superfluous virtual edges left, which are removed to reduce the number of paths. After the merging process is completed, the decision diagram is renormalised to ensure the sum of weights of out-going edges on a node equals 1. An example illustrates the idea.

**Example 4.** Consider the left-hand side of Figure 5. This decision diagram has a top vertex with two out-going edges with the operators  $I$  and  $Y$ . To remove the identity edge, the target vertex of the identity edge (1) is merged into the target vertex of the Pauli- $Y$  edge (2) (indicated by a solid arrow labelled “merge into”). Since both (1) and (2) have an out-going Pauli- $X$  edge to different target vertices, these targets

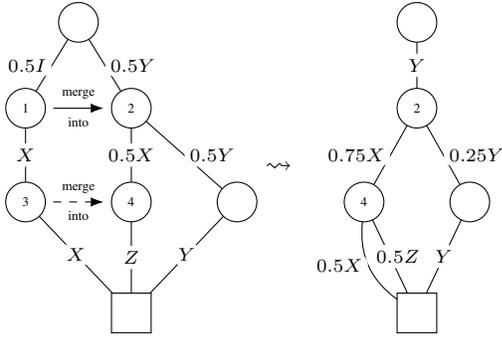


Figure 5: Merging two vertices

(3) and (4) have to merged as well in a recursive fashion (indicated by a dashed arrow).

The following section provides the experimental evaluation on decision diagrams constructed as described above.

## V. NUMERICAL EXPERIMENTS

In this section, we give numerical experiments demonstrating the efficacy of decision diagrams (DDs) to solve the shallow-circuit measurement problem using the implementation available at [github.com/iic-jku/dd-quantum-measurements](https://github.com/iic-jku/dd-quantum-measurements). In particular, the experiments are to confirm that DDs allow to encode probability distributions  $\beta$  from a wider class of probabilities than with known approaches, such as, the Pauli grouping and LBCS. Following the experimental setting of [19], we consider five molecular Hamiltonians that range from 4 to 14 spin orbitals. Excepting the 8-qubit  $\text{H}_2$  (hydrogen) that uses a 6-31G basis, all Hamiltonians are represented in a minimal STO-3G basis. Each of the molecular Hamiltonians is turned into a qubit Hamiltonian using three encodings: Jordan-Wigner (JW), Parity, and Bravyi-Kitaev (BK) [11]. The resulting qubit Hamiltonians are those requiring 4 to 14 qubits that are defined in the molecular basis, and therefore their Hartree-Fock states are those of computational basis. We compare the variances of the measurements obtained from decision diagrams against those from the Pauli grouping and LBCS as shown in Table I.

In the table we list the variances of estimators according to Eq. (7) with  $S = 1$  of two established approaches (LDF, LBCS) and four proposed DD-based approaches:

- LDF Pauli grouping [19]
- Locally-biased classical shadows [19]
- Decision diagrams constructed as described in Section III.
- Decision diagrams with the same construction but up to ten optimisation passes as described in Section III-D (denoted as  $DD+opt10$ )
- Decision diagrams with paths constructed from the LDF Pauli grouping (denoted as  $DD+LDF$ )
- The  $DD+LDF$  with ten optimisation passes as described in Section III-D (denoted as  $DD+LDF+opt10$ )

As shown in Table I, the variances of the DD-based probability distributions are lower than those of LBCS for  $\text{H}_2$ , LiH,

and  $\text{BeH}_2$ , but show deterioration for larger molecules such as  $\text{H}_2\text{O}$ . For the smallest considered molecule (4-qubit  $\text{H}_2$ ), DDs provide a slightly lower variance for the JW encoding compared to LDF grouping but consistently beat LBCS. Also, we can confirm the efficacy of the DDs from the variances of  $DD+LDF$  which are always better than those of Pauli grouping for all considered molecules, except 4-qubit  $\text{H}_2$  in the Parity and BK encodings.

The variances for DD-based probability distributions can be further reduced by applying the optimisation procedure from Section III-D, as can be seen from the column of  $DD+opt10$  and  $DD+LDF+opt10$ : Even with limited steps of optimising the probability distribution, the variances can be reduced below those of LBCS for almost all considered molecules and encodings (except  $\text{H}_2\text{O}$  in the JW encoding). From the column of  $DD+LDF+opt10$  we can observe that although the variances are improved by tuning the probability distribution, the variances are mostly worse than those of  $DD+opt10$ , thus demonstrating the effectiveness of our proposed construction of the DD against the LDF-based one. Apart from the attained variances, decision diagrams are commonly characterized by their number of vertices, number of edges and number of paths. Table II provides these details for the respective molecules, encoding, and (optimised) approaches. In particular, the presented  $DD$ -approach generates more compact decision diagrams (i.e., lower number of vertices) compared to  $DD+LDF$ , but  $DD$  generates more or at least the same number of paths (excepting for LiH at BK encoding). For LDF-based decision diagrams the number of nodes scales with the product of the number of qubits and the number of Pauli terms; for LBCS the decision diagram will have its number of nodes equal to the number of qubits. This hints at the importance of designing compact DDs for shallow measurements.

## VI. CONCLUSION

We have introduced a new estimator for measuring quantum operators defined as linear combination of tensor products of single-qubit Pauli operators. The estimator is defined within a probabilistic measurement framework, where single-qubit measurement bases are drawn from probability distributions obtained using decision diagrams. The decision diagrams used to sample from measurement bases are constructed from target quantum operators, typically Hamiltonians, by associating paths in the diagrams with Pauli operators present in the Hamiltonians. The diagrams can then be simplified by removing paths with identities operators, and merging equivalent sub-paths.

We have shown that representing probability distributions with decision diagrams generalises previous classical-shadow randomises approaches to the measurement problem, namely the uniform [16] and the locally-biased one [19]. This generalisation comes with additional degrees of freedom that characterize each diagram, and introduce correlations between measurement bases for each qubit. We presented different strategies to optimise these additional degrees of freedom,

Table I: Variance for different estimators computed on the ground states

Molecule	Encoding	Variance					
		LDF grouping	LBCS	DDs	DD+opt10	DD+LDF	DD+LDF+opt10
H <sub>2</sub> (4 qubits)	JW	0.402	1.860	0.361	0.398	0.361	0.398
	Parity	0.193	0.541	0.307	0.300	0.292	0.285
	BK	0.193	0.541	0.307	0.300	0.292	0.285
H <sub>2</sub> (8 qubits)	JW	22.3	17.7	8.7	6.2	13.2	7.6
	Parity	38.0	18.9	10.2	8.5	13.6	8.9
	BK	38.4	19.5	7.5	6.4	16.2	9.1
LiH (12 qubits)	JW	54.2	14.8	13.5	8.5	33.1	16.3
	Parity	85.8	26.5	24.2	12.5	53.8	27.2
	BK	75.5	68.0	31.0	14.2	72.0	37.4
BeH <sub>2</sub> (14 qubits)	JW	135	67.6	51.8	32.8	68.4	35.1
	Parity	239	130	72.5	37.2	96.0	55.8
	BK	197	238	200.1	64.2	147.9	73.2
H <sub>2</sub> O (14 qubits)	JW	1 040	258	616.7	294.4	829.4	336.3
	Parity	2 670	429	915.6	425.3	660.9	368.0
	BK	2 090	1 360	1 084.8	527.0	1 403.4	628.1

Table II: Metrics on the generated decision diagrams

Molecule	Encoding	Approach	Vertices	Edges	Paths
H <sub>2</sub> (4 qubits)	JW	DD	10	12	5
		DD+LDF	10	12	5
	Parity	DD	7	7	2
		DD+LDF	7	7	2
	BK	DD	7	7	2
		DD+LDF	7	7	2
H <sub>2</sub> (8 qubits)	JW	DD	77	133	77
		DD+LDF	84	130	61
	Parity	DD	66	119	55
		DD+LDF	54	85	34
	BK	DD	60	101	46
		DD+LDF	56	88	34
LiH (12 qubits)	JW	DD	191	326	273
		DD+LDF	321	459	151
	Parity	DD	289	523	315
		DD+LDF	381	550	182
	BK	DD	413	666	673
		DD+LDF	518	749	1 012
BeH <sub>2</sub> (14 qubits)	JW	DD	537	761	289
		DD+LDF	552	693	147
	Parity	DD	570	872	1 508
		DD+LDF	636	816	192
	BK	DD	494	746	2 192
		DD+LDF	858	1 051	222
H <sub>2</sub> O (14 qubits)	JW	DD	556	894	1 013
		DD+LDF	759	983	234
	Parity	DD	759	1 314	5 711
		DD+LDF	762	1 029	290
	BK	DD	800	1 271	3 914
		DD+LDF	1 083	1 395	618

and have shown numerically that they can outperform locally-biased approaches as well as Pauli grouping strategies, on selected molecular Hamiltonian models. We foresee that future refined approaches in the construction and optimisation of the diagrams could further improve on the improvements in estimation precision reported here, especially considering problem-specific decision diagram construction methods.

## ACKNOWLEDGEMENTS

SH and RW have received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 101001318), the LIT Secure and Correct Systems Lab funded by the State of Upper Austria, as well as the BMK, BMDW, and the State of Upper Austria in the context of the COMET program (managed by the FFG). RR would like to thank Sergey Bravyi and Shigeru Yamashita for technical discussion.

## REFERENCES

- [1] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, “Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets,” *Nature*, vol. 549, no. 7671, p. 242, 2017.
- [2] P. J. O’Malley *et al.*, “Scalable quantum simulation of molecular energies,” *Phys. Rev. X*, vol. 6, no. 3, 2016.
- [3] C. Hempel *et al.*, “Quantum chemistry calculations on a trapped-ion quantum simulator,” *Phys. Rev. X*, vol. 8, no. 3, p. 031 022, 2018.
- [4] D. Wecker, M. B. Hastings, and M. Troyer, “Progress towards practical quantum variational algorithms,” *Phys. Rev. A*, 2015.
- [5] P. Gokhale *et al.*, *Minimizing state preparations in variational quantum eigensolver by partitioning into commuting families*, arXiv:1907.13623, 2019.
- [6] A. F. Izmaylov, T.-C. Yen, R. A. Lang, and V. Verteletskyi, “Unitary partitioning approach to the measurement problem in the variational quantum eigensolver method,” *Journal of chemical theory and computation*, vol. 16, no. 1, pp. 190–195, 2019.
- [7] O. Crawford, B. v. Straaten, D. Wang, T. Parks, E. Campbell, and S. Brierley, “Efficient quantum measurement of pauli operators in the presence of finite sampling error,” *Quantum*, vol. 5, p. 385, 2021.

- [8] A. Zhao, A. Tranter, W. M. Kirby, S. F. Ung, A. Miyake, and P. J. Love, "Measurement reduction in variational quantum algorithms," *Phys. Rev. A*, vol. 101, no. 6, p. 062322, 2020.
- [9] W. J. Huggins, J. R. McClean, N. C. Rubin, Z. Jiang, N. Wiebe, K. B. Whaley, and R. Babbush, "Efficient and noise resilient measurements for quantum chemistry on near-term quantum computers," *npj Quantum Information*, vol. 7, no. 1, 2021.
- [10] I. Hamamura and T. Imamichi, "Efficient evaluation of quantum observables using entangled measurements," *npj Quantum Information*, vol. 6, no. 1, pp. 1–8, 2020.
- [11] S. Bravyi, J. M. Gambetta, A. Mezzacapo, and K. Temme, *Tapering off qubits to simulate fermionic Hamiltonians*, arXiv:1701.08213, 2017.
- [12] T.-C. Yen, V. Verteletskyi, and A. F. Izmaylov, "Measuring all compatible operators in one series of single-qubit measurements using unitary transformations," *Journal of chemical theory and computation*, vol. 16, no. 4, 2020.
- [13] A. Jena, S. Genin, and M. Mosca, *Pauli partitioning with respect to gate sets*, arXiv:1907.07859, 2019.
- [14] V. Verteletskyi, T.-C. Yen, and A. F. Izmaylov, "Measurement optimization in the variational quantum eigensolver using a minimum clique cover," *The Journal of Chemical Physics*, vol. 152, no. 12, p. 124114, 2020.
- [15] G. Torlai, G. Mazzola, G. Carleo, and A. Mezzacapo, "Precise measurement of quantum observables with neural-network estimators," *Phys. Rev. Research*, vol. 2, p. 022060, 2020.
- [16] H.-Y. Huang, R. Kueng, and J. Preskill, "Predicting many properties of a quantum system from very few measurements," *Nature Physics*, 2020.
- [17] S. Aaronson, "Shadow tomography of quantum states," *SIAM Journal on Computing*, vol. 49, no. 5, STOC18–368, 2020.
- [18] H.-Y. Huang, R. Kueng, and J. Preskill, *Efficient estimation of Pauli observables by derandomization*, arXiv:2103.07510, 2021.
- [19] C. Hadfield, S. Bravyi, R. Raymond, and A. Mezzacapo, *Measurements of Quantum Hamiltonians with Locally-Biased Classical Shadows*, arXiv:2006.15788, 2020.
- [20] R. E. Bryant, "Symbolic manipulation of Boolean functions using a graphical representation," in *Design Automation Conference*, 1985, pp. 688–694.
- [21] S. Minato, "Zero-suppressed BDDs and their applications," *Int. J. Softw. Tools Technol. Transf.*, vol. 3, no. 2, pp. 156–170, 2001.
- [22] T. van Dijk, R. Wille, and R. Meolic, "Tagged BDDs: Combining reduction rules from different decision diagram types," in *Formal Methods in Computer Aided Design*, 2017, pp. 108–115.
- [23] S. Minato, " $\pi$ DD: A new decision diagram for efficient problem solving in permutation space," in *Theory and Applications of Satisfiability Testing*, ser. Lecture Notes in Computer Science, vol. 6695, Springer, 2011, pp. 90–104.
- [24] A. Abdollahi and M. Pedram, "Analysis and synthesis of quantum circuits by using quantum decision diagrams," in *Design, Automation and Test in Europe*, 2006, pp. 317–322.
- [25] S.-A. Wang, C.-Y. Lu, I.-M. Tsai, and S.-Y. Kuo, "An XQDD-based verification method for quantum circuits," *IEICE Trans. Fundam. Electron. Commun. Comput. Sci.*, vol. 91-A, no. 2, pp. 584–594, 2008.
- [26] G. F. Viamontes, I. L. Markov, and J. P. Hayes, *Quantum Circuit Simulation*. Springer, 2009.
- [27] P. Niemann, R. Wille, D. M. Miller, M. A. Thornton, and R. Drechsler, "QMDDs: Efficient quantum function representation and manipulation," *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.*, vol. 35, no. 1, pp. 86–99, 2016.
- [28] A. Zulehner, S. Hillmich, and R. Wille, "How to efficiently handle complex values? Implementing decision diagrams for quantum computing," in *International Conference on Computer-Aided Design*, 2019.
- [29] S. Sakaue, M. Ishihata, and S. Minato, "Efficient bandit combinatorial optimization algorithm with zero-suppressed binary decision diagrams," in *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, vol. 84, 2018, pp. 585–594.
- [30] S. Hillmich, C. Hadfield, R. Raymond, A. Mezzacapo, and R. Wille, *Decision diagrams for quantum measurements with shallow circuits*, arXiv:2105.06932, 2021.
- [31] F. Baccari, D. Cavalcanti, P. Wittek, and A. Acín, "Efficient device-independent entanglement detection for multipartite systems," *Phys. Rev. X*, vol. 7, p. 021042, 2017.
- [32] B. Yang, R. Raymond, H. Imai, H. Chang, and H. Hiraishi, *Testing scalable Bell inequalities for quantum graph states on IBM quantum devices*, arXiv:2101.10307, 2021.
- [33] A. Zulehner and R. Wille, "One-pass design of reversible circuits: Combining embedding and synthesis for reversible logic," *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.*, vol. 37, no. 5, pp. 996–1008, 2018.
- [34] A. Zulehner and R. Wille, "Advanced simulation of quantum computations," *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.*, vol. 38, no. 5, pp. 848–859, 2019.
- [35] L. Burgholzer and R. Wille, "Advanced equivalence checking for quantum circuits," *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.*, 2021.
- [36] T. J. Evans, R. Harper, and S. T. Flammia, *Scalable Bayesian Hamiltonian learning*, arXiv:1912.07636, 2019.
- [37] J. Cotler and F. Wilczek, "Quantum overlapping tomography," *Phys. Rev. Lett.*, vol. 124, no. 10, p. 100401, 2020.