



Universiteit
Leiden
The Netherlands

Optimizing the information extracted by a single qubit measurement

Polla, S.; Anselmetti, G.R.; O'Brien, T.E.

Citation

Polla, S., Anselmetti, G. R., & O'Brien, T. E. (2023). Optimizing the information extracted by a single qubit measurement. *Physical Review A*, 108(1). doi:10.1103/PhysRevA.108.012403

Version: Publisher's Version

License: [Leiden University Non-exclusive license](#)

Downloaded from: <https://hdl.handle.net/1887/3655602>

Note: To cite this publication please use the final published version (if applicable).

Optimizing the information extracted by a single qubit measurement

Stefano Polla^{1,2}, Gian-Luca R. Anselmetti,³ and Thomas E. O'Brien^{1,2}

¹Google Quantum AI, 80636 München, Germany

²Instituut-Lorentz, Universiteit Leiden, 2300 RA Leiden, The Netherlands

³Covestro Deutschland AG, Leverkusen 51373, Germany



(Received 26 May 2023; accepted 20 June 2023; published 5 July 2023)

We consider a quantum computation that only extracts one bit of information per N -qubit quantum state preparation. This is relevant for error mitigation schemes where the remainder of the system is measured to detect errors. We optimize the estimation of the expectation value of an operator by its linear decomposition into bitwise-measurable terms. We prove that optimal decompositions must be in terms of reflections with eigenvalues ± 1 . We find the optimal reflection decomposition of a fast-forwardable operator, and show a numerical improvement over a simple Pauli decomposition by a factor $N^{0.7}$.

DOI: [10.1103/PhysRevA.108.012403](https://doi.org/10.1103/PhysRevA.108.012403)

I. INTRODUCTION

The largest bottleneck in quantum algorithm design is the encoding and decoding of a quantum state. Although each full characterization of a quantum state requires an exponentially large amount of information, direct measurements of an N -qubit quantum state ρ extract only N bits of information, and collapse ρ to a state described by those N bits alone, erasing any other information. Performing this repeatedly allows the estimation of an expectation value $\langle O \rangle := \text{Tr}[O\rho]$ of any operator O that is diagonal in the measurement basis. The rate at which such a measurement converges is known as the standard quantum or shot-noise limit [1]; after M repeated preparations, $\langle O \rangle$ can be estimated with variance

$$\text{Var}[O] = M^{-1}(\langle O^2 \rangle - \langle O \rangle^2). \quad (1)$$

Although this rate can be improved upon [2–5], doing so requires implementing long coherent circuits or performing large correlated measurements, which are not feasible in the current noisy intermediate-scale quantum (NISQ) era [6].


Instead of using all N qubits to extract data from a quantum state, one may perform a partial measurement that extracts less than N bits, and use the remaining qubits to detect and mitigate errors [7–9]. Error mitigation is key in obtaining precise results from NISQ circuits, such as variational algorithms [10,11], where the output of the quantum algorithm is a set of estimates of expectation values. Echo verification (EV, see Sec. II B) [12–15] allows one to strongly mitigate errors in a wide class of algorithms, by recasting measurements as Hadamard tests. In each EV circuit, a single bit of information is extracted from the system register as a measurement, freeing up the remainder of the register for error detection and mitigation. One may combine results of multiple EV circuits (through classical postprocessing) into an error-mitigated estimator of any target quantity. However, the stringent requirement that only one bit of information be extracted from the device further tightens the bottleneck of quantum-classical I/O.

In this paper we study how we can optimize information extraction from a quantum system to estimate the expectation

value of an observable O , under the restriction that only a single bit of information is measured per state preparation. This matches the requirements of EV, the rest of the information being reserved for error mitigation. We do not focus in this work on the effectiveness of EV as an error mitigation strategy, and consider only the case of error-free quantum simulation. We define measurements with a single-bit outcome in terms of the Hadamard test, use these to construct an expectation value estimator for a more complicated operator via a linear decomposition, and calculate the variance of this resulting estimator. We prove necessary conditions for such a linear decomposition to be optimal, i.e., to minimize the cost of expectation value estimation. We construct a provably optimal (in some sense) decomposition for a fast-forwardable operator, and give a general (albeit expensive) method to implement this decomposition through quantum signal processing [16–18]. We analyze our methods numerically, comparing the variance of estimators based on our optimal method with other known approaches such as Pauli decompositions and the Dirichlet kernel measurements introduced in [19]. We find an asymptotic improvement between our optimal decomposition and a simple Pauli decomposition of a factor $N^{0.7}$, which at 13 qubits gives already an order of magnitude improvement.

II. SINGLE QUBIT MEASUREMENTS

The most general measurement that extracts one bit of information from a N -qubit state $|\psi\rangle$ is a binary positive-operator valued measurement (binary POVM); this is defined by two positive operators $\Pi_+, \Pi_- > 0$ such that $\Pi_+ + \Pi_- = \mathbb{1}$. The outputs of such measurement, which we label $+1$ and -1 , have probabilities $p_{\pm} = \langle \psi | \Pi_{\pm} | \psi \rangle$. Schematically,




where we defined the unitary preparing the state $V|0\rangle := |\psi\rangle$. In Sec. II A, we review the Hadamard test and we show that there exists a one-to-one equivalence between outcomes of Hadamard tests and binary POVMs.

Extracting only a single bit allows further processing of the quantum information remaining in the state register. For instance, inverting the unitary that prepared $|\psi\rangle$ and measuring in the computational basis yields a powerful error mitigation technique, echo verification [12–14], which we review in Sec. II B. In another example, the Hadamard test may be used to estimate the gradient of a cost function with respect to a variational term $\exp(iA\theta)$ in a circuit, as $\frac{d}{d\theta} \exp(iA\theta) = iA \exp(iA\theta)$ [20,21]. Both these methods require operating on the system register after the binary measurement is performed, preventing further information extraction. (For the specific case of EV, we show in Appendix B that extracting more than one bit of information is counterproductive.) Furthermore, this restricted output model of quantum computation can be relevant in quantum-enhanced metrology settings [2,22], where a single qubit probe is used [23]. A similar restricted access model has been studied in the context of Hamiltonian learning [24,25]. Note that this single qubit access model is different to the one clean qubit model of computation (DQC-1) [26]; here we consider using a single qubit to extract information from a nontrivial quantum state.

A. Hadamard test

A Hadamard test (HT) is a binary measurement performed on a state $|\psi\rangle_s$ in the N -qubit system register s . It is implemented through a control qubit c initialized in the state $|+\rangle_c$, a controlled unitary CU , and a projective Pauli measurement X_c on the control qubit. As a quantum circuit this can be written



and the resulting state before measurement can be easily calculated to be

$$|\Phi\rangle = \frac{1}{\sqrt{2}}(|0\rangle|\psi\rangle + |1\rangle U|\psi\rangle). \tag{2}$$

Tracing out the system register then yields the following reduced density matrix on the control qubit:

$$\rho_c = \frac{1}{2} \begin{pmatrix} 1 & \langle U \rangle \\ \langle U \rangle^* & 1 \end{pmatrix}. \tag{3}$$

One may estimate the expectation value of $\text{Re}(U) := \frac{1}{2}(U + U^\dagger)$ by measuring the control qubit in the X basis, which returns $\text{Tr}[X\rho_c] = \langle \frac{1}{2}(U + U^\dagger) \rangle$.

To prove the equivalence between HT and binary POVM, we explicitly construct one from another. To construct the binary POVM corresponding to the HT, we define the measurement operators that represent the back-action of the

measurement on the system register

$$M_\pm = \langle \pm |_c CU | + \rangle_c = \frac{\mathbb{1} \pm U}{2}, \tag{4}$$

and the relative positive operators $\Pi_\pm = M_\pm^\dagger M_\pm$ used to compute probabilities $p_\pm = \langle \psi | \Pi_\pm | \psi \rangle$ of measuring ± 1 on the ancilla. Vice versa, given a binary POVM $\{\Pi_+, \Pi_-\}$, we can construct a corresponding Hadamard test by choosing a unitary U that satisfies $\text{Re}(U) = \Pi_+ - \Pi_-$,


$$U = \exp[i \arccos(\Pi_+ - \Pi_-)]. \tag{5}$$

This is always possible because $\Pi_+ - \Pi_-$ is Hermitian and $\|\Pi_+ - \Pi_-\| < \|\Pi_+ + \Pi_-\| = 1$. It is easy to check that the Hadamard test constructed from this unitary return the correct positive operators Π_\pm .

B. Echo verification

The name echo verification (EV) refers to a class of powerful error mitigation techniques [12–14], applicable in most algorithms that make use of a Hadamard test to perform measurements on a system register. This technique was originally introduced by the name of *verified phase estimation* [12] as it considered estimating expectation values of multiple unitaries $U_l = e^{iHt_l}$, with an archetypal application in the context of single-ancilla phase estimation. However, in this work, we consider the more general expectation-value estimation subroutine yielding $\langle \text{Re}(U) \rangle$. We prefer the name echo verification (used also in [15,27]) due to the similarities to a Loschmidt echo.

Echo verification relies on a key idea: exploiting the information left in the system register after the application of the controlled-unitary operator prescribed by the Hadamard test. This information is used to detect errors and mitigate their effect on estimated quantities. This is done by “echoing” the preparation unitary V , i.e., applying V^\dagger after the controlled evolution, and verifying whether the register s returns to the initial state $|0\rangle$. The corresponding circuit is



where the multiplication of the classical information channels (red double lines) sets the circuit output to zero upon failed verification (i.e., if the final system state is orthogonal to $|0\rangle_s$), and to the output of the Hadamard test otherwise.

Let us denote the combined state after the controlled unitary as $|\Phi\rangle$, and let $\Pi_\psi = |\psi\rangle\langle\psi| = V^\dagger|0\rangle\langle 0|V$ be the projector on the state $|\psi\rangle_s$. The estimate of $\langle \text{Re}(U) \rangle$ can be obtained by measuring the operators $X^{\text{EV}} := X \otimes \Pi_\psi$ on $|\Phi\rangle$ (EV circuit), as opposed to $X_c := X \otimes \mathbb{1}$ (HT circuit). One can confirm that, in the absence of error, these operators have identical expectation values on the state at the end of the circuit [12]

$$\langle \Phi | X^{\text{EV}} | \Phi \rangle = \langle \Phi | X_c | \Phi \rangle = \langle \psi | \text{Re}(U) | \psi \rangle. \tag{6}$$

For an intuitive explanation, note that if the controlled unitary changes the state of the system register, the ancilla qubit must have been in the $|1\rangle$ state, and $\langle 1|X|1\rangle = 0$. This implies that the expectation value of $X \otimes (\mathbb{1}_N - \Pi_\psi)$ is 0.

In the presence of a circuit error, verification is likely to fail. This decreases the expectation value measured by the error probability, which can be measured separately. Rescaling the result by the error probability yields a noise-mitigated estimate of the expectation value $\langle \psi | \text{Re}(U) | \psi \rangle$. The error mitigation power of this method is explored in [12–14,27,28] and experimentally tested in [15]. In this work, we only consider noiseless circuits.

The EV circuit implements a ternary measurement, with outputs $+1, -1, 0$. Compared to a standard HT defined by the same unitary, the probabilities p_+ and p_- are reduced by the same amount ($\frac{p_0}{2}$), yielding a result with the same expected value. As a consequence, the variance of an EV measurement is always smaller than that of the corresponding HT (this is formalized in Appendix A).


An extension of Echo verification allows extracting more than one qubit of information per circuit run by using multiple auxiliary qubits. However, as the measurement is quadratic in $|\psi\rangle\langle\psi|$ (resulting by the use of two copies of V_ψ in the circuit), reconstructing the desired expectation values requires non-linear processing of the measurement results. Furthermore, as each measurement interferes with the verification of the others, all the variances of estimated expectations increase. In Appendix B we explore this, and we prove that measuring more than one bit of information per EV experiment is always counterproductive in terms of final variance, for a fixed total number of shots.

C. Ancilla-free echo verification

The direct (control-based) measurement via the HT may often be replaced by an indirect measurement using an altered circuit [12,29,30], allowing control-free implementations of these single-bit measurements. We review briefly the control-free echo verification scheme.

In the Hadamard test, the control qubit provides a clock-reference state $|0\rangle|\psi\rangle$, which is not changed by the application of CU. This clock-reference state is necessary to give physical meaning to the phase U induces on the system register states, thus making it measurable. If U has a known eigenstate $U|\psi_r\rangle = e^{i\phi_r}|\psi_r\rangle$ orthogonal to $|\psi\rangle$, this state can be used as a clock-reference removing the need for a control qubit. In quantum simulation, this state can often be found thanks to the symmetries of the system. For example, in second-quantized simulation of particle systems the vacuum state $|0\rangle$ is an eigenstate of any particle-number preserving operator.

The control-free EV scheme prescribes preparing a cat state $\frac{1}{\sqrt{2}}(|\psi\rangle + |\psi_r\rangle)$, applying U , and measuring $X^{\text{CFEV}} = (|\psi\rangle\langle\psi_r| + |\psi_r\rangle\langle\psi|)$. This can be done with the circuit



where $V_{\text{cat}}|0 \dots 00\rangle = |\psi_r\rangle$ and $V_{\text{cat}}|0 \dots 01\rangle = |\psi\rangle$. After the application of U , the state is $|\Phi\rangle = \frac{1}{\sqrt{2}}(U|\psi\rangle + e^{i\phi_r}|\psi_r\rangle)$, thus

$$\langle \Phi | X^{\text{CFEV}} | \Phi \rangle = \langle \psi | \text{Re}(U e^{-i\phi_r}) | \psi \rangle. \quad (7)$$

If $\phi_r \neq 0$, the desired result $\langle \psi | \text{Re}(U) | \psi \rangle$ can be obtained by substituting $U \rightarrow U e^{i\phi_r}$ or applying a phase gate $e^{-i\phi_r/2Z}$ to the first qubit before measurement.

D. Variance of a binary POVM

The Hadamard test differs from the projective measurement of $\text{Re}(U) := \frac{1}{2}(U + U^\dagger)$ (the Hermitian part of U). Each instance of the Hadamard test can only output $+1$ or -1 , whereas the spectrum of $\text{Re}(U)$ can have up to 2^N distinct eigenvalues in the range $[-1, 1]$. This has a direct impact on the estimation uncertainty: performing the Hadamard test M times and measuring the control qubit in the X basis yields an estimator of $\langle \text{Re}(U) \rangle = \text{Re}(\langle U \rangle)$ with a variance

$$\text{Var}^*[\langle \text{Re}(U) \rangle] = \frac{1 - \langle \text{Re}(U) \rangle^2}{M}, \quad (8)$$

which can be seen to be strictly larger than the variance one would obtain by performing a projective measurement of $\text{Re}(U)$ on M copies of $|\psi\rangle$ [Eq. (1)],

$$\text{Var}[\langle \text{Re}(U) \rangle] \leq \text{Var}^*[\langle \text{Re}(U) \rangle], \quad (9)$$

as $\langle \text{Re}(U) \rangle^2 \leq 1$. Our goal is to optimize estimators of expectation values $\langle O \rangle$ of a given operator, which use data from multiple HTs with different unitaries U [each with the given variance Eq. (8)], and assuming one test per state preparation. We want to minimize the total number of state preparations (distributed over different choices of U) needed to achieve an estimator of $\langle O \rangle$ with error smaller than a fixed ϵ .

III. OPERATOR DECOMPOSITIONS

It is common in quantum computing to estimate the expectation value of an operator O by writing O as a linear combination of simpler terms (i.e., a *decomposition*) which have their expectation values estimated independently [10,31,32]. In this work, we make use of this method, and consider estimating these simpler terms via Hadamard tests. Let us fix a decomposition¹ X ,

$$O = \sum_{x \in X} c_x \text{Re}(U_x) \leftrightarrow \langle O \rangle = \sum_{x \in X} c_x \langle \text{Re}(U_x) \rangle, \quad (10)$$

and consider estimating $\langle O \rangle$ by estimating each $\langle \text{Re}(U_x) \rangle$ independently and summing the results. As $\text{Re}(U_x)$ and O are Hermitian operators we may assume c_x to be real without loss of generality, and we may further assume $c_x \geq 0$ by absorbing a minus sign onto U_x . Note that the arrow in Eq. (10) points both ways as the set of expectation values on all states $|\psi\rangle$ uniquely defines an operator.

¹In a slight abuse of notation, throughout this work we will use the same label (e.g., X) to represent the entire linear decomposition defined by the set $\{c_x, U_x\}$ in Eq. (10), and the set of labels x that we sum over.

Once a suitable decomposition X of an operator O [Eq. (10)] has been chosen, to calculate the total cost of the algorithm we must allocate a number m_x of repeated single-shot HT experiments to estimate individual $\langle \text{Re}(U_x) \rangle$. We assume a single-bit measurement per state preparation, i.e., each HT requires resetting the circuit and re-preparing $|\psi\rangle$, and the total number of re-preparations $M_X = \sum_{x \in X} m_x$ is the relevant cost of implementing our measurement scheme. If each $\langle \text{Re}(U_x) \rangle$ is estimated independently, the variance on a final estimate of $\langle O \rangle$ can be calculated by standard propagation of variance

$$\text{Var}_X^*[\langle O \rangle] = \sum_{x \in X} c_x^2 \text{Var}^*[\langle \text{Re}(U_x) \rangle] \quad (11)$$

$$= \sum_{x \in X} \frac{c_x^2 [1 - \langle \text{Re}(U_x) \rangle^2]}{m_x}. \quad (12)$$

Equation (9) implies that under the same decomposition of O

$$\text{Var}_X[\langle O \rangle] := \sum_{x \in X} c_x^2 \text{Var}[\langle \text{Re}(U_x) \rangle] \leq \text{Var}_X^*[\langle O \rangle], \quad (13)$$

for all states ρ .

A. Adaptive shot allocation

Given a decomposition X and a total shot budget M_X , an optimal choice for the m_x may be found using Lagrange multiplier methods [33]

$$m_x = M_X \frac{c_x \sqrt{1 - \langle \text{Re}(U_x) \rangle^2}}{\sum_{y \in X} c_y \sqrt{1 - \langle \text{Re}(U_y) \rangle^2}}, \quad (14)$$

recalling that $c_x \geq 0$. This yields a bound on the required M_X to estimate $\langle O \rangle$ with $\text{Var}^*[\langle O \rangle] = \epsilon^2$:

$$M_X \geq \mathcal{M}_X := \epsilon^{-2} \left[\sum_{x \in X} c_x \sqrt{1 - \langle \text{Re}(U_x) \rangle^2} \right]^2. \quad (15)$$

We call \mathcal{M}_X the *cost* of the decomposition X . This may be compared to well-known results for measurement bounds using standard tomography methods [9,32–34] by substituting Var^* for Var in Eq. (11). Although exact values of $\langle U_x \rangle$ will not be known in advance, these can be estimated using a small initial fraction of measurements before a final distribution of measurements is allocated.

B. Decomposition hierarchy

We have shown above how to optimize measurement allocation given a linear decomposition X [Eq. (10)]. Let us now consider how to optimize X to minimize Eq. (15).

We first consider the effect of possible rescalings of $\text{Re}(U_x)$. If any term $c_x \text{Re}(U_x)$ has $\|\text{Re}(U_x)\| < 1$,² one can find some unitary $U_{x'}$ for which $\text{Re}(U_{x'}) = \text{Re}(U_x) / \|\text{Re}(U_x)\|$; substituting $U_x \rightarrow U_{x'}$ (and $c_x \rightarrow c_{x'}$ accordingly) will always improve the bound in Eq. (15). (For now we do not worry about how the unitaries may be implemented as quantum circuits; we will consider this issue later.)

²Unless stated otherwise, all norms in this work are the spectral norm.

One may next consider subdividing individual terms $\text{Re}(U_x)$ of X , by writing

$$c_x \text{Re}(U_x) = c_{x,0} \text{Re}(U_{x,0}) + c_{x,1} \text{Re}(U_{x,1}), \quad (16)$$

where $U_{x,0}$ and $U_{x,1}$ are both unitary, and $c_x, c_{x,0}, c_{x,1} > 0$. As we can assume $\|\text{Re}(U_x)\| = 1$, such a decomposition requires $c_{x,0} + c_{x,1} \geq c_x$, to preserve the spectral norm of $\text{Re}(U_{x,0})$ and $\text{Re}(U_{x,1})$. When this inequality is saturated, we call the subdecomposition *norm preserving*. It turns out that this condition is sufficient for the subdecomposition to be nonincreasing in the cost \mathcal{M} of estimation [Eq. (15)], for all states $|\Psi\rangle$; formally:

Lemma 1. Given a linear decomposition X of a target operator O [Eq. (10)], a subdecomposition X' [Eq. (16)] that is norm preserving has nonincreasing cost, $\mathcal{M}_{X'} \leq \mathcal{M}_X$ [Eq. (15)], for any state $|\Psi\rangle$.

We give a proof of this lemma in Appendix C 1

We would like to extend the above lemma to a statement that norm-increasing subdecompositions of a linear decomposition X are always suboptimal in some sense. To achieve this, note that as a corollary to Lemma 1, we can improve on all terms $c_x \text{Re}(U_x)$ in a linear decomposition X by a norm-preserving identity shift

$$c_x \text{Re}(U_x) = c_x (1 - \bar{\lambda}_x) \text{Re}(U_{\bar{x}}) + c_x \bar{\lambda}_x \mathbb{1}, \quad (17)$$

where $\bar{\lambda}_x = \frac{1}{2}(\lambda_x^{\min} + \lambda_x^{\max})$, λ_x^{\min} and λ_x^{\max} are the lowest and highest eigenvalues of $\text{Re}(U_x)$, respectively, and $\text{Re}(U_{\bar{x}})$ has the same eigenvectors of $\text{Re}(U_x)$ (with its spectrum shifted and rescaled). We call the outcome decomposition \tilde{X} of the procedure above the *center* of X . Although a norm-increasing subdecomposition of X may not be suboptimal relative to X , it is suboptimal relative to this center:

Lemma 2. Let X be a linear decomposition of O with all $\|\text{Re}(U_x)\| = 1$; let \tilde{X} be the center of X and let X' be a strictly norm-increasing subdecomposition. There exists at least one state $|\Psi\rangle$ for which the cost $\mathcal{M}_{\tilde{X}} < \mathcal{M}_{X'}$.

We give a proof of this lemma in Appendix C 2.

To recap, the above two lemmas show (a) that norm-preserving subdecompositions do not increase the cost of estimating expectation values via Hadamard tests on any given state, and (b) norm-increasing subdecompositions not only can increase expectation value estimation costs on some states, but are guaranteed to do so on at least one. This result is in direct contrast to standard expectation value estimation, where independent estimation of $\langle A \rangle$ and $\langle B \rangle$ is suboptimal to joint estimation of $\langle A + B \rangle$ whenever the latter is possible. This suggests a path towards optimizing HT expectation value estimation, by repeatedly dividing terms $\text{Re}(U_x)$ in a norm-preserving manner, until no further subdecomposition can reduce the cost any state. It turns out that not all choices of division lead to the same end point, however, all end points of this procedure have one common property (proven in Appendix C 3):

Lemma 3. A decomposition X of an operator O has no nontrivial norm-preserving subdecompositions if and only if all operators $\text{Re}(U_x)$ in X are reflections: $\text{Re}(U_x)^2 = 1$.

It should be no surprise that we find reflection operators $\text{Re}(U_x)^2 = 1$ to be a crucial ingredient to optimize HT tomography, as these are the only operators that saturate the bound in Eq. (9) for all states $|\Psi\rangle$. We call a decomposition X that

consists of reflection operators only a *reflection decomposition*. We give some simple examples of these in Appendix C 4.

C. Optimizing reflection decompositions

Above we demonstrated that, for a decomposition X of an operator O to be optimal with regards to the cost \mathcal{M}_X of estimating expectation values on a set of states [Eq. (15)], all terms in X must be reflection operators. Otherwise, we demonstrated a means of subdividing single terms in the distribution to generate a new distribution with lower cost. However, this is not to say that all reflection decompositions X have the same cost \mathcal{M}_X . (These two statements are consistent as we cannot transform between reflection decompositions using subdivision.) The set of reflection decompositions of O form a convex set that is $2^{2^N - N}$ dimensional if all U_x are diagonal in the eigenbasis of O . This raises two questions: Is there an optimal decomposition amongst the set of reflection decompositions, and does it achieve the von Neumann bound [Eq. (13)]?

Lemma 4. Let O be an operator and Π_j be projectors onto the eigenvalues of O ; $O\Pi_j = \Pi_j O = \lambda_j \Pi_j$. The Ξ -decomposition of O , given by

$$O = \frac{\lambda_0 + \lambda_J}{2} \mathbb{1} + \sum_{x=1}^{J-1} \frac{\delta\lambda_x}{2} \Xi_x, \quad (18)$$

$$\Xi_x = \mathbb{1} - \sum_{j < x} 2\Pi_j, \quad \delta\lambda_x = \lambda_x - \lambda_{x-1}, \quad (19)$$

uniquely achieves the bound $\text{Var}_{\Xi}^*[O] = \text{Var}[O]$ on all states $|\Psi\rangle$ with support on up to two eigenstates of O . No such decomposition achieves this bound on all states $|\Psi\rangle$ with support on three or more eigenstates of O .

We prove this lemma in Appendix C 5. Note that the Ξ decomposition can be immediately restricted to any subspace of the full 2^N -dimensional Hilbert space containing $|\Psi\rangle$ (i.e., if we knew that due to a symmetry or by virtue of being a low-energy state, $|\Psi\rangle$ had support only on such a space), and the optimality result still holds. This implies in turn that no linear decomposition X can achieve the von Neumann variance bound even for as small as a three-dimensional subspace. This makes sense, as our restriction to measure one bit of information per state preparation forms a bottleneck with respect to the three nonzero-probability outcomes of a von Neumann measurement on this space.


D. Implementing the optimal decomposition

In order to realize the Ξ -decomposition estimator, we need to implement HT circuits that (approximately) estimate $\langle \Xi_x \rangle$. This may be achieved by realizing that

$$\Xi_x = \text{sgn}[O - \mu_x], \quad \mu_x = \frac{\lambda_{x-1} + \lambda_x}{2}, \quad (20)$$

where sgn is the sign function. An approximation of this unitary operator can then be realized using quantum signal processing (QSP) [16–18] of the sign function [35], requiring only one additional ancillary qubit. The QSP circuit is

given by



where $R_X(\phi_r) = e^{-i\frac{\phi_r}{2} \sigma_x}$ implements a unitary block encoding Q_ϕ of a degree- R trigonometric polynomial S_ϕ of the operator $(O - \mu_x)t$:

$$\langle 1|_c Q_\phi |0\rangle_c = \sum_{r=0}^R c_r(\phi) e^{-ir(O - \mu_x)t} := S_\phi[(O - \mu_x)t]. \quad (21)$$

Here, ϕ is a vector containing the individual angles ϕ_r implemented during the QSP circuit. We can then sample $\langle \text{Re}\{S_\phi[(O - \mu_x)t]\} \rangle$ through HT (or EV), using another qubit controlling all gates in the QSP circuit. To approximate Eq. (20) with our block-encoded operator S_ϕ , we must choose $t < \frac{\pi}{\|O - \mu_x\|}$ to avoid aliasing, and find the optimal ϕ :

$$\phi = \arg \min_{(\phi_r = -\phi_{R-r})} \int_{0+\delta}^{\pi-\delta} d\omega [\text{sgn}(\omega) - \text{Im}[S_\phi(\omega)]]. \quad (22)$$

Here, the constraint $\phi_r = -\phi_{R-r}$ ensures $\text{Im}[S_\phi(\omega)]$ is an odd function of ω . A resolution parameter $\delta \geq 0$ can be introduced to improve the approximation away from the nodes $\omega = \{0, \pm\pi\}$ of $S_\phi(\omega)$. In Appendix D we give further details of this decomposition, and analyze the approximation error numerically. We find that this error converges exponentially in the number of circuit blocks R .

IV. NUMERICAL EXPERIMENTS

To investigate performance of various decompositions on states that have support on more than two eigenstates of O , and therefore are not covered by Lemma 4, we perform numerical simulations using random variationally generated states and a simple toy operator $O = \sum_j Z_j$. (In Appendix F, we report this scaling for other systems.) We measure the variances on states generated by a hardware-efficient *Ansatz* [36] with random input parameters using PENNYLANE [37]. For each data point 100 random states are generated. We consider estimating $\langle O \rangle$ in a realistic scenario where the $\langle \text{Re}(U_x) \rangle$ values will not be known in advance to optimally choose m_x via Eq. (14). Instead, for each random state we generate a prior estimate of each $\langle \text{Re}(U_x) \rangle$ from 10^5 measurements of the state, and use these to determine m_x (which are then only approximately optimal). This leaves the total shot count M_X as a free parameter; we resolve this in Fig. 1 by calculating $M_X \text{Var}_X^*[\langle O \rangle]$. (This gives a quantity that is relevant regardless of the number of the shots actually used to estimate $\langle O \rangle$.)

An average of $M_X \text{Var}_X^*[\langle O \rangle]$ over the 100 states is formed and plotted in Fig. 1 for each grouping method. This is compared to the von Neumann measurement variance $\text{Var}[O]$, which does not require any shot allocation, and sets a lower limit to the other estimators [see Appendix C 5, Eq. (C15)]. The Ξ decomposition (orange, “ Ξ ”) has the best asymptotic scaling of all decompositions, being suboptimal to $\text{Var}[O]$ by

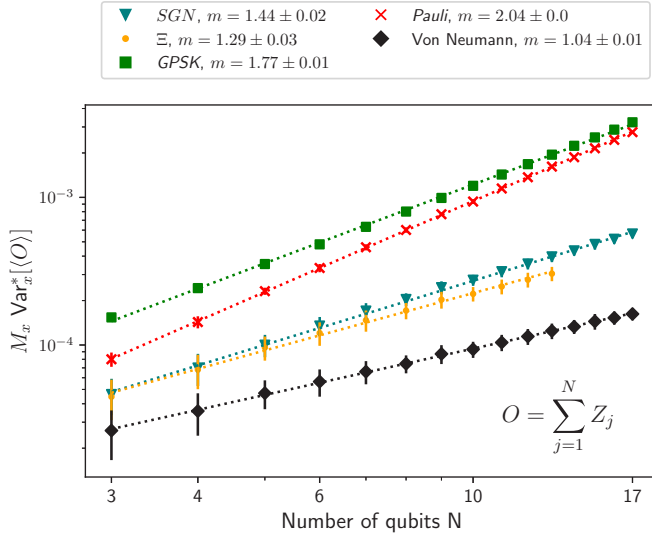


FIG. 1. Comparison study of variances of different decompositions on random states generated by a hardware-efficient *Ansatz* (see text for details). Different colors correspond to different decompositions [Eq. (10)] of the target operator O (see text for the description of all decompositions). Dashed lines are power-law fits to the data (obtained exponents are given in legend).

a factor $\approx N^{1/3}$. The QSP approximation of Ξ , (teal, “SGN”), has a slightly worse asymptotic scaling, which we associate to the error in approximating $\text{sgn}(O - \mu_j)$. At the largest considered $N = 13$, these two decompositions suffer approximately a factor 2 penalty in their total cost compared to $\text{Var}[\langle O \rangle]$. The generalized parameter-shift kernel decomposition [19] (green, “GPKS,” described in Appendix E) has the worst overall performance out of the investigated estimators, due to the constant factor. It has, however, a better asymptotic scaling than a simple Pauli decomposition $U_x = Z_j$ (red, “Pauli,” Appendix C 4). In Appendix F we investigate the scaling of different sets of observables. We observe that the order of the performance of the different decompositions remains consistent throughout, but the relative gains and losses in performance can be significantly different.

V. CONCLUSION

In this work we studied the optimization of expectation value estimation for a quantum state in the case where we are only allowed to measure a single qubit per state preparation (e.g., through Hadamard tests, with relevant application to echo verification). We calculated the cost of estimating the expectation value of an operator O by linearly decomposing O into a linear combination of subunitary terms, assuming an optimal shot allocation. We demonstrated that this cost is strictly nonincreasing when terms are further subdivided, under the constraint that this subdivision preserves the induced 1-norm of the term coefficients. We showed that the end points of this procedure of repeated division are linear decompositions of O where all terms are reflection operators: a so-called “reflection decomposition.” We identified one such decomposition, the Ξ decomposition, as unique in its ability to estimate $\langle O \rangle$ with a variance matching the von Neumann

measurement limit on any linear combination of up to 2 eigenstates of O . We demonstrated how the Ξ decomposition may be approximately implemented through quantum signal processing. Numerical results demonstrate that on simple systems, the Ξ decomposition and its approximate counterpart demonstrate clear constant and asymptotic improvements over other reflection decompositions (in the cost of estimating $\langle O \rangle$ on random states), with up to a factor 10 \times improvement for estimation on 20 qubits.

Although these results are encouraging, the significant discrepancy between $\text{Var}_{\Xi}^*[O]$ and $\text{Var}[O]$ is worrying for NISQ algorithms that already incur a significant cost to tomograph complex Hamiltonians [7,9,32,38–40]; either one incurs a large overhead for measurement due to the need to invoke quantum signal processing or incur the clear asymptotic scaling cost that comes with measuring single Pauli terms per state preparation. Given that echo verification has a sampling cost scaling as $1/F^2$ (for a circuit fidelity F) [12], this result adds to the unlikelihood of beyond-classical NISQ variational algorithms in chemistry. Finding reflection decompositions with lower circuit depth is a clear avenue for future work.

ACKNOWLEDGMENTS

The authors wish to acknowledge F. Gkritsis for providing shot allocation code, and W. Huggins, R. Babbush, J. Tura, A. Dutkiewicz, D. Wierichs, and C. Gogolin for useful discussions and advice on this work. Covestro acknowledges funding from the German Ministry for Education and Research (BMBF) under the funding program quantum technologies as part of project HFAK (Project No. 13N15630).

APPENDIX A: ECHO VERIFICATION ESTIMATORS

The estimator used for echo verification is not identical to the one studied in the main text, and so its variance is not quite identical. In particular, we have $(X^{\text{EV}})^2 = \Pi_{\psi}$ which implies that the variance on an estimate of $\langle \Phi | X_{\text{EV}} | \Phi \rangle$ is

$$\text{Var}_{\text{EV}}^*[\langle \text{Re}(U) \rangle] = \frac{\langle \Phi | \mathbb{1} \otimes \Pi_{\psi} | \Phi \rangle - \langle \psi | \text{Re}(U) | \psi \rangle^2}{M}. \quad (\text{A1})$$

Clearly $|\langle \Phi | \psi \rangle|^2 \leq 1$, which implies $\text{Var}_{\text{EV}}^*[\langle \text{Re}(U) \rangle] \leq \text{Var}^*[\langle \text{Re}(U) \rangle]$ [by comparison with Eq. (8)]. In other words, the variance of the EV estimator is always smaller or equal to the variance of the relative HT estimator. It is easy to calculate from the circuit above that

$$\langle \Phi | \Pi_{\psi} \otimes \mathbb{1} | \Phi \rangle = \frac{1}{2} |1 + \langle \psi | U | \psi \rangle|^2 \quad (\text{A2})$$

[noting that $\langle \text{Re}(U) \rangle = \text{Re}(\langle U \rangle)$, which can be substituted back into our variance estimate to obtain

$$\text{Var}_{\text{EV}}^*[\langle \text{Re}(U) \rangle] \geq \frac{1 - \langle \psi | \text{Re}(U) | \psi \rangle^2}{2M}. \quad (\text{A3})$$

Thus, we have

$$\text{Var}^*[\langle \text{Re}(U) \rangle] \geq \text{Var}_{\text{EV}}^*[\langle \text{Re}(U) \rangle] \geq \frac{\text{Var}^*[\langle \text{Re}(U) \rangle]}{2}. \quad (\text{A4})$$

This justifies our focus in the main text on optimizing the estimator from a standard Hadamard test; this estimator is

simpler to analyze, more general, and differs from the EV estimator (that motivated this work) by at most a factor 2.

APPENDIX B: PARALLELIZING ECHO VERIFICATION

In absence of echo verification, we can trivially parallelize Hadamard tests measuring K commuting operators $\{\text{Re}(U_0), \dots, \text{Re}(U_{K-1})\}$ using K ancillary qubits, one controlling each U_k . If each U_k is controlled by a separate ancillary qubit (labeled k , where $C_k U_k$ represents the k th unitary controlled by the k th control qubit), the combined state of the system register s and ancillary qubits after all the unitaries are applied will be

$$\bigotimes_k C_k U_k |+\rangle_k |\psi\rangle. \quad (\text{B1})$$

The probabilities of obtaining ± 1 when measuring X on the j th control qubit are

$$p_{j\pm} = \left\| \bigotimes_{k \neq j} C_k U_k |+\rangle_k \frac{1 \pm U_j}{2} |\psi\rangle \right\|^2 \quad (\text{B2})$$

$$= \frac{1}{4} \langle \psi | (1 \pm U_j^\dagger) (1 \pm U_j) | \psi \rangle \quad (\text{B3})$$

which coincides with the probabilities of a single Hadamard test with unitary U_j .

When performing echo verification, parallelization is more complicated. The result of verification (the measurement of $\Pi_\psi = |\psi\rangle\langle\psi|$ on the system register) is affected by all the controlled U_k , and thus its result cannot be simply associated to one specific ancilla being in the state $|1\rangle$. To mitigate errors, all the cases in which the register is found in a state orthogonal to $|\psi\rangle$ should be considered as null towards all of the ancilla measurement results. The echo-verified probability of measuring the binary string $\vec{\sigma} = (\sigma_0, \dots, \sigma_k)$, where each σ_k is ± 1 corresponding to the state $|\pm\rangle$ measured on the k th ancilla, is then

$$p_{\vec{\sigma}}^{\text{EV}} = \left| \langle \psi | \prod_k \langle \sigma_k | C_k U_k |+\rangle_k | \psi \rangle \right|^2 \\ = \frac{1}{4^K} \left| \langle \psi | \prod_k (1 + \sigma_k U_k) | \psi \rangle \right|^2. \quad (\text{B4})$$

The product in this equation can be then developed into a linear combination of 2^K expectation values (note that, as all U_k commute, the order does not matter). Under the assumption that all these expectation values are real [granted if $U_k = \text{Re}(U_k)$] Eq. (B4) defines a quadratic system of 2^K equations with $2^K - 1$ unknowns.³ Solving such system we find that the expectation value of a single $\text{Re}(U_j)$ can be

³In the case of a more general $U = \text{Re}(U) + i\text{Im}(U)$, a similar system can be constructed by measuring each U_k and iU_k with $2K$ ancillas. Showing this is out of the scope of our work, and for the sake of simplicity we restrict ourselves to the case of Hermitian $U_k = \text{Re}(U_k)$.

estimated by processing the sampled $p_{\vec{\sigma}}^{\text{EV}}$ as

$$\langle \text{Re}(U_j) \rangle = \underbrace{\left(\sum_{\vec{\sigma}: \sigma_j = +1} \sqrt{p_{\vec{\sigma}}^{\text{EV}}} \right)^2}_{p_{j+}^{\text{EV}}} - \underbrace{\left(\sum_{\vec{\sigma}: \sigma_j = -1} \sqrt{p_{\vec{\sigma}}^{\text{EV}}} \right)^2}_{p_{j-}^{\text{EV}}}, \quad (\text{B5})$$

where we denoted $p_{j\pm}^{\text{EV}}$ the terms that reproduce the probabilities that would be returned by a single, un-parallelized EV experiment

$$p_{j\pm}^{\text{EV}} = \frac{1}{4} |\langle \psi | 1 \pm U_j | \psi \rangle|^2. \quad (\text{B6})$$

We assume $p_{\vec{\sigma}}^{\text{EV}}$ are sampled by averaging M shots of the parallel EV experiment. These are probabilities of mutually exclusive measurements, thus the covariance matrix of the $p_{\vec{\sigma}}^{\text{EV}}$ estimators is defined by

$$\text{Var}[p_{\vec{\sigma}}^{\text{EV}}] = \frac{1}{M} p_{\vec{\sigma}}^{\text{EV}} (1 - p_{\vec{\sigma}}^{\text{EV}}), \quad (\text{B7})$$

$$\text{Cov}[p_{\vec{\sigma}}^{\text{EV}}, p_{\vec{\rho}}^{\text{EV}}] = -\frac{1}{M} p_{\vec{\sigma}}^{\text{EV}} p_{\vec{\rho}}^{\text{EV}} \quad \text{if } \vec{\sigma} \neq \vec{\rho}. \quad (\text{B8})$$

We can then propagate the error through Eq. (B5) to obtain the variance on the parallel-EV (PEV) estimator of $\langle \text{Re}(U_j) \rangle$:

$$M \text{Var}_{\text{PEV}}^*[\langle \text{Re}(U_j) \rangle] \\ = \sum_{\vec{\sigma}} \frac{p_{j\sigma_j}^{\text{EV}}}{p_{\vec{\sigma}}^{\text{EV}}} p_{\vec{\sigma}}^{\text{EV}} (1 - p_{\vec{\sigma}}^{\text{EV}}) - \sum_{\vec{\sigma} \neq \vec{\rho}} \sigma_j \rho_j \frac{\sqrt{p_{j\sigma_j}^{\text{EV}}}}{\sqrt{p_{\vec{\sigma}}^{\text{EV}}}} \frac{\sqrt{p_{j\rho_j}^{\text{EV}}}}{\sqrt{p_{\vec{\rho}}^{\text{EV}}}} p_{\vec{\sigma}}^{\text{EV}} p_{\vec{\rho}}^{\text{EV}} \\ = \sum_{\vec{\sigma}} p_{j\sigma_j}^{\text{EV}} - \langle \text{Re}(U_j) \rangle \\ = 2^{K-1} (p_{j\sigma_+}^{\text{EV}} + p_{j\sigma_-}^{\text{EV}}) - \langle \text{Re}(U_j) \rangle, \quad (\text{B9})$$

which explodes exponentially with the size of the parallelization K .

More generally, we can compute the covariance matrix for all the $p_{j,\sigma_j}^{\text{EV}}$ through error propagation

$$\text{Var}[p_{j,\sigma_j}^{\text{EV}}] = p_{j,\sigma_j}^{\text{EV}} (2^{K-1} - p_{j,\sigma_j}^{\text{EV}}), \quad (\text{B10})$$

$$\text{Cov}[p_{j,\sigma_j}^{\text{EV}}, p_{k,\rho_k}^{\text{EV}}] = \delta_{j,k} \sqrt{p_{j,\sigma_j}^{\text{EV}} p_{k,\rho_k}^{\text{EV}}} - p_{j,\sigma_j}^{\text{EV}} p_{k,\rho_k}^{\text{EV}} \quad (\text{B11})$$

[where the covariance assumes $(j, \sigma_j) \neq (k, \rho_k)$]. This shows that, increasing K , we effectively add to the covariance matrix a positive-semidefinite term with a norm that scales exponentially in K . As all the decompositions Eq. (10) are ultimately to be estimated as linear combinations of the sampled probabilities $p_{j,\sigma_j}^{\text{EV}}$, parallelizing error verification is counterproductive.

APPENDIX C: PROOF OF DECOMPOSITION OPTIMALITY HIERARCHY

In this Appendix we build up to the proof that the Ξ decomposition is optimal in terms of cost (15), by proving the lemmas introduced in the main text. We first prove that a norm-preserving subdecomposition has nonincreasing cost with respect to its parent decomposition, for all states $|\psi\rangle$. We then prove that a subdecomposition that does not have the norm-preserving property is always suboptimal (i.e., it has

strictly greater cost than an alternative norm-preserving subdecomposition). The iteration of the norm-preserving subdecomposition procedure leads to one of many alternative improving sequences of decompositions. The end point of each sequence is a norm-preserving linear decomposition of O for which all unitaries are reflection operators. Finally, we prove that one of such decompositions (the Ξ decomposition) achieves the von Neumann measurement variance bound on a certain set of states, and that no unbiased estimator based on single-qubit measurements can achieve this bound on a larger set of states.

1. Proof of Lemma 1, and corollaries

Given a linear decomposition X of an operator O [Eq. (10)], consider a norm-preserving subdecomposition X' where a single term $x \in X$ is split according to Eq. (16). The bound on the total number of shots, Eq. (15), will then change:

$$\begin{aligned} \mathcal{M}_X &\rightarrow \mathcal{M}_{X'} \\ &= \epsilon^{-2} \left[\sum_{y \neq x} c_y \sqrt{1 - \langle \text{Re}(U_y) \rangle^2} + c_{x,0} \sqrt{1 - \langle \text{Re}(U_{x,0}) \rangle^2} \right. \\ &\quad \left. + c_{x,1} \sqrt{1 - \langle \text{Re}(U_{x,1}) \rangle^2} \right]^2 \end{aligned} \quad (\text{C1})$$

[with the change with respect to Eq. (15) being the second row]. This results in a reduction of the cost, as can be seen by calculating

$$\begin{aligned} &c_x^2 [1 - \langle \text{Re}(U_x) \rangle^2] \\ &= (c_{x,0} + c_{x,1})^2 - [c_{x,0} \langle \text{Re}(U_{x,0}) \rangle + c_{x,1} \langle \text{Re}(U_{x,1}) \rangle]^2 \\ &= c_{x,0}^2 [1 - \langle \text{Re}(U_{x,0}) \rangle^2] + c_{x,1}^2 [1 - \langle \text{Re}(U_{x,1}) \rangle^2] \\ &\quad + 2c_{x,0}c_{x,1} [1 - \langle \text{Re}(U_{x,0}) \rangle \langle \text{Re}(U_{x,1}) \rangle] \\ &\geq c_{x,0}^2 [1 - \langle \text{Re}(U_{x,0}) \rangle^2] + c_{x,1}^2 [1 - \langle \text{Re}(U_{x,1}) \rangle^2] \\ &\quad + 2c_{x,0}c_{x,1} \sqrt{[1 - \langle \text{Re}(U_{x,0}) \rangle^2][1 - \langle \text{Re}(U_{x,1}) \rangle^2]} \\ &= [c_{x,0} \sqrt{1 - \langle \text{Re}(U_{x,0}) \rangle^2} + c_{x,1} \sqrt{1 - \langle \text{Re}(U_{x,1}) \rangle^2}]^2, \end{aligned} \quad (\text{C2})$$

where, in the center inequality we have used the fact that for $0 \leq a, b \leq 1$,

$$1 - ab \geq \sqrt{(1 - a^2)(1 - b^2)}. \quad (\text{C3})$$

As a corollary and example, we look at identity shifts of a term $x \in X$. For $\text{Re}(U_x)$ with unit norm, we can assume without loss of generality the largest eigenvalue is $\lambda_{\max} = 1$, and the smallest is λ_{\min} . We can then perform the simple norm-preserving decomposition

$$c_x \text{Re}(U_x) = c_x(1 - \bar{\lambda}) \text{Re}(U_{x'}) + c_x \bar{\lambda} \mathbb{1} \quad (\text{C4})$$

with $\bar{\lambda} = \frac{1}{2}(\lambda_{\min} + \lambda_{\max})$. The resulting $\text{Re}(U_{x'})$ has maximum eigenvalue $+1$ and minimum eigenvalue -1 , thus it does not admit nontrivial identity shift.

A norm-preserving subdecomposition (16) of a term with $|\text{Re}(U_x)| = 1$ will only admit terms with $|\text{Re}(U_{x,i})| = 1$. [This can be checked by taking the expectation value of both sides of Eq. (16) on the eigenstate on which $|\langle \text{Re}(U_x) \rangle| = 1$.] By the same reasoning, terms with $\text{Re}(U_x)$ having maximum eigenvalue $+1$ and minimum eigenvalues -1 [like those obtained by the identity shifts (C4)] only admit subdecompositions whose terms have the same property.

2. Proof of Lemma 2

In this Appendix we compare the costs of two decompositions derived by an original decomposition X : the center \bar{X} where all terms are transformed according to Eq. (17), and the norm-increasing subdecomposition X' where a term $x \in X$ is changed according to Eq. (16) assuming $c_{x,0} + c_{x,1} > c_x$. Remembering that all coefficients are positive $c_y > 0$, the cost of each decomposition (15) is the square of a sum of positive values; the terms in this sum for $y \neq x$ do not change for $X \rightarrow X'$, and have a nonincreasing value for $X \rightarrow \bar{X}$. We thus focus only on the term $x \in X$ and the derived ones, highlighted here:

$$\mathcal{M}_{X'} = \epsilon^{-2} \left[\sum_{j \in \{0,1\}} \overbrace{c_{x,j} \sqrt{1 - \langle \text{Re}(U_{x,j}) \rangle^2}}^{m'} + \dots \right]^2, \quad (\text{C5})$$

$$\mathcal{M}_{\bar{X}} = \epsilon^{-2} \left[\underbrace{c_x(1 - \bar{\lambda}_x) \sqrt{1 - \langle \text{Re}(U_{\bar{x}}) \rangle^2}}_{\bar{m}} + \dots \right]^2. \quad (\text{C6})$$

We now prove there exists a state $|\Psi\rangle$ for which $\bar{m} < m'$, which implies $\mathcal{M}_{\bar{X}} < \mathcal{M}_{X'}$.

Let $|\psi_+\rangle$ and $|\psi_-\rangle$ be eigenvectors of $\text{Re}(U_{\bar{x}})$ with eigenvalues $+1$ and -1 , respectively. We consider three cases:

(1) $|\langle \psi_\sigma | \text{Re}(U_{x,j}) | \psi_\sigma \rangle| < 1$ for at least one combination of $\sigma \in \{+, -\}$ and $j \in \{0, 1\}$. In this case, on the state $|\Psi\rangle = |\psi_\sigma\rangle$ we get $\bar{m} = 0 < m' \neq 0$.

(2) $\langle \psi_\sigma | \text{Re}(U_{x,j}) | \psi_\sigma \rangle = \sigma$ for all combinations of $\sigma \in \{+, -\}$ and $j \in \{0, 1\}$. By combining Eqs. (16) and (17) and taking the expectation value on $|\psi_\sigma\rangle$ we obtain $\sigma[c_{x,0} + c_{x,1} - c_x(1 - \bar{\lambda}_x)] = c_x \bar{\lambda}_x$, which implies $c_{x,0} + c_{x,1} = c_x$, violating one of the hypotheses of the lemma.

(3) $\langle \psi_\sigma | \text{Re}(U_{x,j}) | \psi_\sigma \rangle = (-1)^j \sigma$ for all combinations of $\sigma \in \{+, -\}$ and $j \in \{0, 1\}$. We define the state $|\Psi\rangle = \frac{|\psi_+\rangle + |\psi_-\rangle}{\sqrt{2}}$, on which $\langle \text{Re}(U_{\bar{x}}) \rangle = \langle \text{Re}(U_{x,0}) \rangle = \langle \text{Re}(U_{x,1}) \rangle = 0$. On this state, the costs are $\mathcal{M}_{\bar{X}} = \epsilon^2 c_x^2 (1 - \bar{\lambda}_x)^2$ and $\mathcal{M}_{X'} = \epsilon^2 (c_{x,0} + c_{x,1})^2$. As $\bar{\lambda}_x \geq 0$ and $c_{x,0} + c_{x,1} > c_x$, $\mathcal{M}_{X'} < \mathcal{M}_{\bar{X}}$.

3. Proof of Lemma 3

In this Appendix, we prove that the end point of norm-preserving decomposition sequences are reflection operators. In other terms, if $\text{Re}(U_x)$ is a reflection operator, it only admits a norm-preserving subdecomposition [Eq. (16)] if $\text{Re}(U_{x,0}) = \text{Re}(U_{x,1}) = \text{Re}(U_x)$.

To prove this, consider a state $|\psi\rangle$ in the $+1$ eigenspace of $\text{Re}(U_x)$. For a norm-preserving decomposition, we

must have

$$\begin{aligned} c_{x,0} + c_{x,1} &= c_x = c_x \langle \psi | \text{Re}(U_x) | \psi \rangle \\ &= c_{x,0} \langle \psi | \text{Re}(U_{x,0}) | \psi \rangle + c_{x,1} \langle \psi | \text{Re}(U_{x,1}) | \psi \rangle. \end{aligned} \quad (\text{C7})$$

As $\|\text{Re}(U_{x,0})\|, \|\text{Re}(U_{x,1})\| \leq 1$, this equality can only be satisfied if $|\psi\rangle$ is also a $+1$ eigenstate of both $U_{x,0}$ and $U_{x,1}$. A similar argument holds for all -1 eigenstates of U_x , and so $U_{x,0}$, $U_{x,1}$, and U_x share the same eigenstates and eigenvalues and must be equal. Taking such a subdecomposition has no effect on the estimator of $\langle O \rangle$, as the same HT are performed and the total number of shots does not change, i.e., $\mathcal{M}_{X'} = \mathcal{M}$ in Eq. (C1).

4. Examples of reflection decompositions

The simplest example of a reflection-based decomposition is a decomposition in terms of Pauli operators

$$O = \sum_j^J c_j Z_j, \quad (\text{C8})$$

with $c_j \geq 0$. We could be tempted to measure $\langle O \rangle$ with a single HT circuit (assuming access to a block encoding of $\frac{O}{\|O\|}$, which is optimal). In this case, as $O = \|O\| \text{Re}(U)$, the bound (15) is

$$M \geq \epsilon^{-2} \|O\|^2 \left[1 - \frac{\langle O \rangle^2}{\|O\|^2} \right]. \quad (\text{C9})$$

To improve on this, we can estimate each $\langle Z_j \rangle$ separately, each with a Hadamard test with $U_j = \langle Z_j \rangle$ (a binary operator). As the spectral norm of O is equal to the induced 1-norm $\|O\|_1 = \sum_j^J c_j$, Eq. (C8) is a norm-preserving decomposition. The bound (15) then becomes

$$M \geq \epsilon^{-2} \left[\sum_j c_j \sqrt{1 - \langle Z_j \rangle^2} \right]^2, \quad (\text{C10})$$

which is always smaller or equal than Eq. (C9) [easily proven through Eq. (C3)]. This inequality is only saturated when the considered state ρ has support only on the $\|O\|^2$ -eigenvalue subspace of O^2 ; the operator O projected on this subspace is effectively a binary operator.

Norm-preserving decompositions do not need to involve only mutually commuting Pauli operators. As a practical example, we consider the two-qubit operator $O = XX + YY$, which appears commonly in quantum Hamiltonians. As $O = 2 \text{Im}[i\text{SWAP}]$, this operator can be measured with a single Hadamard test circuit. Furthermore, in the context of electronic structure Hamiltonians, O preserves particle number, so in general a control-free scheme using the vacuum as reference state can be employed for the measurement. This operator has three eigenvalues $\{0, \pm 1\}$, which means we can improve its measurement by decomposing it in binary operators. We propose three decompositions $O = \text{Re}U_0 + \text{Re}U_1$. The obvious Pauli decomposition $U_0 = XX$, $U_1 = YY$ has the downside of not conserving particle number. To fix this, we

can take

$$U_j = \frac{1}{2}[(XX + YY) + (-1)^j(Z\mathbb{1} + \mathbb{1}Z)]. \quad (\text{C11})$$

These are particle-number preserving, reflection operators and can be easily implemented by combining $i\text{SWAP}$ with single qubit $e^{\pm iZ\pi/4}$ rotations on both qubits. The last decomposition,

$$U_j = \frac{1}{2}[(XX + YY) + (-1)^j(ZZ + \mathbb{1}\mathbb{1})], \quad (\text{C12})$$

uses particle-preserving reflection operators with different eigenvalue multiplicities: unlike Pauli operators, the ± 1 -eigenvalue subspaces of U_j have unequal dimension 1 and 3. For any state in the 0-eigenvalue subspace, spanned by $\{|00\rangle, |11\rangle\}$, the estimate variance $\text{Var}^*[\langle \text{Re}(U_j) \rangle] = 0$ for decomposition Eq. (C12). This is not true for the other two decompositions, which indicates that not all decompositions in binary operators are born equal. We will deal with this in the next section. Another example of a few-qubit reflection operator that is a sum of noncommuting Pauli operators is the three-spin all-to-all Heisenberg coupling

$$O = \frac{1}{3} \sum_{l=1}^2 \sum_{m=0}^{l-1} X_m X_l + Y_m Y_l + Z_m Z_l, \quad (\text{C13})$$

which appears, e.g., in the kagome-Heisenberg Hamiltonian.

5. Proof of Lemma 4

In this Appendix we prove Lemma 4, which formally states the optimality and uniqueness of the Ξ decomposition. To do this, we first define a variance bound for a class of estimators of $\langle O \rangle$ on a state $|\psi\rangle$. We prove that the bound is achieved on all eigenstates of O if all the sampled operators $\text{Re}(U_x)$ are diagonal in the eigenbasis of O . We then construct the Ξ decomposition, and prove that the related estimator saturates the bound on the set S_2 of all states with support on at most two eigenstates of O . Finally, we prove no other decomposition satisfies this requirement (i.e., the Ξ decomposition is unique), and no decomposition satisfies the bound on a superset $S \supset S_2$.

A decomposition X [Eq. (10)] of an operator O is optimal on a state $|\psi\rangle$ if no other decomposition produces an estimator with lower cost [Eq. (15)] for that state. Optimality can be defined for a set S of states: X is optimal on S if, for each $|\psi\rangle \in S$, no decomposition X' has lower cost $M_{X'} < M_X$. (Note that this can be readily generalized to mixed state, without changing any of our next results.) Lemmas 1–3 imply a necessary condition for optimality on the whole Hilbert space: X can only be optimal on all states if it has the form

$$\begin{aligned} O &= \bar{\lambda}_O \mathbb{1} + \sum_{x \in X} c_x \text{Re}(U_x), \\ c_x &> 0, \quad |\bar{\lambda}_O| + \sum_{x \in X} c_x = \|O\|, \quad \text{Re}(U_x)^2 = \mathbb{1}, \end{aligned} \quad (\text{C14})$$

where $\bar{\lambda}_O$ is the average of the largest and smallest eigenvalues of O . In other words, X is a norm-preserving decomposition of the center of O where all sampled terms are reflection operators. This condition is not sufficient, as many nonequivalent instances of such decompositions exist, as exemplified in Appendix C4.

We now construct a bound on the variance of the estimator of $\langle O \rangle$ based on the decomposition X : saturating this bound on all $|\psi\rangle \in S$ implies optimality of X on S [the cost of the decomposition (15) is defined as the minimum value of M required to achieve target variance ϵ^2 , so minimum variance at fixed M implies minimum cost at fixed ϵ]:

$$\text{Var}_X^*[\langle O \rangle] = \frac{1}{M} \left[\sum_x c_x \sqrt{1 - \langle \text{Re}(U_x) \rangle^2} \right]^2 \geq \frac{\text{Var}[O]}{M}. \quad (\text{C15})$$

This bound is implied by Eqs. (9) and (12), with the choice of optimal shot allocation (14). Its physical interpretation is rooted in the following observation: a von Neumann measurement of O is the lowest-variance unbiased estimator of $\langle O \rangle$ when given access to a single state preparation. Thus, given M independent experiments each with a single state preparation, the mean of von Neumann measurements is the lowest-variance unbiased estimator.

We first consider the set S_1 of all eigenstates of O . For any $|\phi\rangle \in S_1$, the value of the bound in Eq. (C15) becomes $\text{Var}[O] = 0$. The bound is thus saturated only if we choose all reflection operators $\text{Re}(U_x)$ diagonal in any eigenvector basis of O , i.e., $[U_x, O] = 0$ and $U_x|\phi\rangle = \pm|\phi\rangle$ for any $|\phi\rangle \in S_1$. For any decomposition of this form, we can write all U_x in terms of the eigenspace projectors of O :

$$U_x = \sum_{j=0}^{J-1} \xi_{x,j} \Pi_j, \quad \xi_{x,j} \in \{\pm 1\}, \quad (\text{C16})$$

where Π_j is the projector on the (eventually degenerate) λ_j eigenspace of O , J is the number of distinct eigenvalues $\{\lambda_j\}$ of O , and without loss of generality we assume $\lambda_j > \lambda_{j-1}$. The coefficients will then have to satisfy the relation $\lambda_j = \sum_x c_x \xi_{x,j}$.

We define the Ξ decomposition based on Eq. (C16), by choosing $\xi_{x,j} = -1$ if $j < x$, and $+1$ otherwise. The resulting decomposition is presented in Lemma 4, Eq. (18). The operators Ξ_x are reflections by definition, and it is easy to check that the decomposition satisfied the necessary condition (C14). Note that $c_0 = (\lambda_0 + \lambda_j)/2$ defines the optimal identity shift (producing the center of O) and the $c_x = (\lambda_x - \lambda_{x-1})/2$ complete the decomposition.

We now prove that the Ξ decomposition is optimal on the set S_2 of states with support on two eigenstates of O :

$$S_2 = \left\{ \frac{\alpha|\lambda_m\rangle + \beta|\lambda_n\rangle}{\sqrt{\alpha^2 + \beta^2}} : |\lambda_m\rangle, |\lambda_n\rangle \in S_1 \right\}. \quad (\text{C17})$$

On a general state $|\psi\rangle$ with eigenspace occupations $a_j = \langle \psi | \Pi_j | \psi \rangle$, the estimator based on the Ξ decomposition has variance

$$\text{Var}_\Xi^*[\langle O \rangle] = \frac{1}{M} \left[\sum_j \frac{\delta\lambda_j}{2} \sqrt{4 \left(\sum_{i<j} a_i \right) \left(\sum_{i\geq j} a_i \right)} \right]^2. \quad (\text{C18})$$

For a state $|\phi\rangle \in S_2$, only two occupations are nonzero $a_m, a_n \neq 0$ (we assume without loss of generality $m < n$), thus the term under square root is reduced to $4a_m a_n$ if $m < j \leq n$ and 0 otherwise. The resulting variance

$$\begin{aligned} \text{Var}_\Xi^*[\langle \phi | O | \phi \rangle] &= \frac{1}{M} \left[\frac{\lambda_n - \lambda_m}{2} \sqrt{4a_m a_n} \right]^2 \\ &= \frac{1}{M} a_n a_m (\lambda_n - \lambda_m) = \text{Var}[\langle O \rangle], \end{aligned} \quad (\text{C19})$$

thus saturating the bound (12).

We now prove that the only optimal decomposition on S_2 is the Ξ decomposition (or equivalent up to relabeling and trivial subdecompositions). First of all, $S_1 \subset S_2$, so the terms of the decomposition need to be of the form of Eq. (C16). Consider a family of states $\sqrt{a_m}|\lambda_m\rangle + \sqrt{a_n}|\lambda_n\rangle$ for any $n > m$, with only two nonzero eigenstate occupations $a_m + a_n = 1$. On such a state,

$$\begin{aligned} \text{Var}_X^*[O] &= \frac{1}{M} \left[\sum_x c_x \sqrt{1 - [a_m \xi_{x,m} + a_n \xi_{x,n}]^2} \right]^2 \\ &= \frac{a_m a_n}{M} \left[\sum_x 2c_x \frac{1 - \xi_{x,m} \xi_{x,n}}{2} \right]^2. \end{aligned} \quad (\text{C20})$$

The bound (C15) is then saturated when

$$\left[\sum_x 2c_x \frac{1 - \xi_{x,m} \xi_{x,n}}{2} \right]^2 = \lambda_n - \lambda_m, \quad (\text{C21})$$

where we simplified out the free parameter $\frac{a_m a_n}{M}$. This can be rewritten as

$$\sum_x c_x \xi_{x,n} (\xi_{x,n} - \xi_{x,m}) = \sum_x c_x (\xi_{x,n} - \xi_{x,m}) \quad (\text{C22})$$

using the condition on the decomposition coefficients $\lambda_j = \sum_x c_x \xi_{x,j}$. This implies that if $\xi_{x,n} = -1$, then $(\xi_{x,m} - \xi_{x,n}) = 0$ (recall that $c_x > 0$), i.e., $\xi_{x,m} = -1$. Thus, the only U_x that can appear in this decomposition are of the same form as the operators in the Ξ decomposition ($\xi_{j,m} = -1, \xi_{j,n} = +1$ for $m < j \leq n$), and thus X is either Ξ or a trivial subdecomposition of it.

We now show that the Ξ decomposition does not saturate the bound (C15) for a state $|\psi\rangle$ with three nonzero occupations $a_m, a_n, a_p \neq 0$ ($m < n < p$). On this state we can write

$$\text{Var}_\Xi^*[\langle O \rangle] = \frac{1}{M} [(\lambda_n - \lambda_m) \sqrt{a_m(a_n + a_p)} \quad (\text{C23})$$

$$+ (\lambda_p - \lambda_n) \sqrt{a_p(a_m + a_n)}]^2. \quad (\text{C24})$$

Subtracting from this $\text{Var}[\langle O \rangle]$, expanding and then collecting terms we get

$$\begin{aligned} \text{Var}_\Xi^*[\langle O \rangle] - \text{Var}[\langle O \rangle] &= [(\lambda_n - \lambda_p)(\lambda_n - \lambda_m)] \\ &\quad \times [a_m a_p - \sqrt{a_m a_p (a_m + a_n)(a_p + a_n)}] > 0, \end{aligned} \quad (\text{C25})$$

as both the terms in square brackets are strictly smaller than zero. This (along with the uniqueness of Ξ as the optimal estimator on S_2) implies that no HT-based estimator can saturate the bound (C15) for arbitrary states.


In fact, the bound can only be saturated on states in S_2 : on these states the von Neumann measurement has only two possible outcomes (λ_m and λ_n) with nonzero probability. The adaptive shot allocation scheme then ensures (for a large enough M) that most of the measurements we take (Ξ_x with $m \leq x < n$) reproduce the statistics of the von Neumann measurement, with the single bit we sample in every experiment always distinguishing between λ_m and λ_n . On any state $|\psi\rangle \in \overline{S_2}$, the von Neumann measurement has three or more outcomes with nonzero probability, and we cannot reproduce its statistics by sampling a single qubit per experiment. This, along with the uniqueness of Ξ , implies that no decomposition can satisfy the sufficient condition for optimality on a superset $S \supset S_2$. The numerical results presented in this paper quantify the increase in variance with respect to the bound, along with

confirming the Ξ decomposition outperforms other decompositions on all states.

APPENDIX D: IMPLEMENTATION OF THE Ξ DECOMPOSITION VIA QUANTUM SIGNAL PROCESSING

Verifiable sampling of QSP polynomials. To measure the operators in the Ξ decomposition (20), we implement a Hadamard test (or EV) on trigonometric polynomials of $(H - \mu_x)t$ generated by the quantum signal processing. We tune the QSP coefficients such that the polynomials approximate the sign function in a suitable range. In this Appendix we display and analyze this technique.

The full circuit we use to achieve this is



The first control qubit (labeled HT) takes care of the Hadamard test. The second ancilla (labeled QSP) manages the quantum signal processing subroutine, extended through the sign-controlled evolution $e^{-iZ \otimes (O - \mu_x)t}$ to implement a quantum signal processing (QSP) on the operator $e^{(O - \mu_x)t}$. We now describe how the measurement scheme works, and how to select the ϕ parameters to approximate a measurement of $\text{sgn}[(O - \mu_x)t]$ in the interval $[-\pi, \pi]$.

First, we analyze the QSP routine. Let us assume $|\psi\rangle$ to be an eigenstate of $(O - \mu_x)t$ with eigenvalue $\omega \in (-\pi, \pi)$, and only consider the effect of the controlled gates (removing the HT qubit). Then, we can reduce the circuit to an effective single qubit gate on the QSP qubit, with action

$$Q_\phi(\omega) = e^{-i\frac{Y}{2}\pi} e^{-i\frac{X}{2}\phi_R} \left[\prod_{r=1}^R e^{-i\frac{Z}{2}2\omega} e^{-i\frac{X}{2}\phi_{R-r}} \right] = \begin{pmatrix} S_\phi(\omega) & \cdots \\ \cdots & \cdots \end{pmatrix} \quad (\text{D1})$$

which is a block encoding of $S_\phi(\omega)$, a degree- R trigonometric polynomial of ω . For the sake of simplicity we inserted the final gate $e^{-i\frac{Y}{2}\pi} = -iY$, shifting the polynomial of interest S from the block $\langle 1|Q|0\rangle$ to $\langle 0|Q|0\rangle$. We ensure $S_\phi(\omega)$ is real and odd by constraining

$$\phi_r = -\phi_{R-r} \Rightarrow S(\omega) = -S(-\omega) \in \mathbb{R}. \quad (\text{D2})$$

Reintroducing the system register, i.e., taking a general $|\psi\rangle_S$, can be done by linearity taking $Q(\omega) \mapsto Q[(O - \mu_x)t]$ and recovering the circuit above.

The result of the verified Hadamard test (or EV) is obtained by measuring on the output state of the circuit the expectation value of Z_{HT} (or $Z_{\text{HT}} \otimes |0\rangle\langle 0|_{\text{QSP}} \otimes |\psi\rangle\langle \psi|_S$). (In the

absence of noise these two expectation values are equal. In the presence of noise, an additional measurement at $t = 0$ can be taken to mitigate errors. For more details on the technique we refer the reader to the original work on EV [12].)

Approximating the sign function. To approximate the operators (20) that make up the Ξ decomposition, we need to choose the QSP parameter ϕ such that $S_\phi(\omega)$ in Eq. (D1) approximates $\text{sgn}[\omega]$. The polynomial $S_\phi(\omega)$ is odd, real, and 2π periodic, thus having nodes $S_\phi(0) = S_\phi(\pm\pi) = 0$. To account for the approximation error in the neighborhood of these nodes, we introduce a resolution parameter $\delta \geq 0$, and request the approximation to be effective only in the $[\delta, \pi - \delta]$ interval. Choosing $\delta > 0$ implies accepting a larger error in approximating the sign function close to zero. For example, we know the eigenvalues of $(O - \mu_x)t$ closest to zero have absolute value $\frac{\delta\lambda_x}{2}t$, and we can use this knowledge to choose δ .

We define a loss function to characterize the quality of the approximation: the average error

$$\mathcal{L}_\delta(\phi) = \frac{1}{\pi - 2\delta} \int_\delta^{\pi-\delta} d\omega [\text{sgn}(\omega) - \text{Im}[S_\phi(\omega)]]. \quad (\text{D3})$$

To choose the optimal parameters ϕ , we minimize this loss under the constraints (D2). Although an analytical approach to this problem is possible building on the techniques described in [18], we take the numerical route to this approximation (which is efficient, scalable, and easy to implement). The integral is thus substituted with a sum on a grid with a number of points much larger than the degree of the trigonometric polynomial. We plot in Fig. 2 the minimized cost function, as a function of the approximation's order R and of the resolution parameter δ . We find that the loss always decays




FIG. 2. Loss equation (D3) for the optimal choice of QSP parameters ϕ , as a function of the order R (number of QSP layers) and resolution parameter δ . The dotted lines are log-lin fits for $R > 10$. The dependence of the fit parameter β on the resolution δ is shown in the inset.

exponentially with an increasing order R , with a decay rate depending on δ .

APPENDIX E: THE GENERALIZED PARAMETER-SHIFT KERNEL DECOMPOSITION OF A DIAGONAL OPERATOR WITH LADDER SPECTRUM

In [19] the authors propose techniques to estimate derivatives $\langle \frac{d}{dt} U(t) \rangle$ of a unitary $U(t) = e^{iOt}$ generated by O , by sampling $\langle \sin(Ot_l) \rangle = \langle \text{Re}[U(t_l)] \rangle$ at a discrete set of points $\{t_l\}$. This technique can be used to estimate expectation values of O , as $\langle O \rangle = \langle [-i \frac{d}{dt} e^{iOt}]_{t=0} \rangle$, and it is clearly compatible with Hadamard test or EV measurements (as it only requires sampling $\langle \text{Re}U(t_l) \rangle$).

For an operator O with equispaced eigenvalues $\Omega, 2\Omega, \dots, R\Omega$ (commonly referred to as a “ladder spectrum”), the authors give a choice of $\{t_l\}$ and explicit coefficients c_l for the linear combination $\langle \frac{d}{dt} U(t) \rangle = \sum_l c_l \langle \text{Re}[-iU(t_l)] \rangle$. Assuming $\Omega = 1$ (which can be considered a choice of units for the energy), the time points are chosen as $\{t_l = \frac{2l}{2R+1}\pi\}$. We can then define a modified version of the Dirichelet kernel,

$$\tilde{D}_l(t) = \frac{1}{R} \cos(t_l) \left[\frac{1}{2} \sin(Rt) + \sum_{j=1}^{R-1} \sin(jt) \right], \quad (\text{E1})$$

which satisfies $\tilde{D}_l(t_l) = \delta_{ll}$. This is a linear combination of the R basis functions $\{\sin(jt)\}_{j=1, \dots, R}$, like $\langle \sin(Ot) \rangle$. Thus, as the equality

$$\langle \sin(Ot) \rangle = \sum_{l=1}^R \langle \sin(Ot_l) \rangle \tilde{D}_l(t) \quad (\text{E2})$$

holds for all $\{t_l\}_{l=1, \dots, R}$, it must to hold for all t . We can then differentiate the kernel rather than the expectation value itself. Evaluating $[\frac{\partial}{\partial t} \tilde{D}_l(t)]_{t=0}$ and combining the equations above




FIG. 3. Comparison study of variances of different decompositions on random states generated by a hardware-efficient Ansatz (see text for details). Different colors correspond to different decompositions [Eq. (10)] of the target operator O (see text for the description of all decompositions). Dashed lines are exponential fits $[a \exp(mN + b)]$ to the data (the parameter m is given in legend).

we obtain

$$\langle O \rangle = \sum_{l=1}^R \frac{(-1)^{l-1}}{2R \sin^2(\frac{1}{2}t_l)} \langle \sin(Ot_l) \rangle \quad (\text{E3})$$

$$= \sum_{l=1}^R c_l \langle \text{Re}[-iU(t_l)] \rangle. \quad (\text{E4})$$

This matches the form of decompositions (10). We call this the generalized parameter-shift kernel (GPSK) decomposition. Under the optimal shot allocation choice [Eq. (14)], the shot variance of the estimator based on this decomposition is

$$M \text{Var}_{\text{GPSK}}^* = \left[\sum_{l=1}^R \frac{\sqrt{1 - \langle \sin(Ot_l) \rangle^2}}{|2R \sin^2(\frac{1}{2}t_l)|} \right]^2. \quad (\text{E5})$$

APPENDIX F: DETAILS ON NUMERICAL SIMULATIONS AND FURTHER NUMERICAL RESULTS

We measure the variances on random states generated by hardware-efficient Ansatzes using PENNYLANE [37]. For each value of N , 100 random set of parameters (and therefore 100 random states) are generated and measured for all decompositions. For each decomposition X , we first use 10^5 shots (allocated proportionally to the weight of each term) to obtain a rough estimate of the expectation value of each term $\langle \text{Re}(U_x) \rangle$ for $x \in X$. These values are plugged in Eq. (14) to get an estimate of the optimal shot allocation ratios $r_x = \frac{m_x}{M_x}$. The variance of each term $\text{Var}^*[\langle \text{Re}(U_x) \rangle]$ is obtained by Eq. (8) (or by sampling in the case of the QSP-approximation decomposition SGN). With these we compute the final shot variance $M_X \text{Var}_X^*[\langle O \rangle] = \sum_{x \in X} r_x^{-1} \text{Var}^*[\langle \text{Re}(U_x) \rangle]$. Finally, we average the values of $M_X \text{Var}_X^*[\langle O \rangle]$ obtained for each random state. This average is the quantity reported in Figs. 1, 3, and 4.

The terms Ξ_x are constructed as per Eq. (18) using the known eigenvectors of O , and projectively measured on the

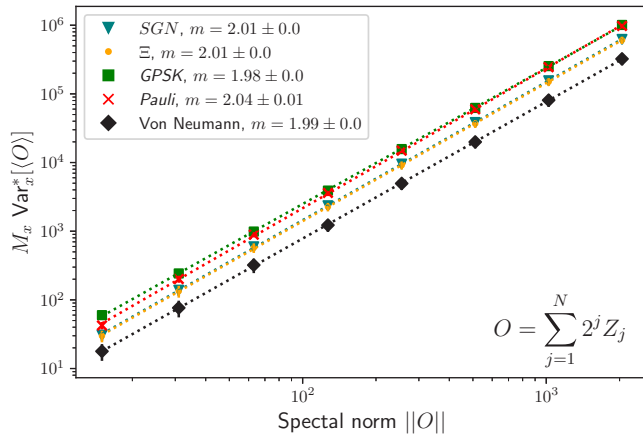


FIG. 4. Comparison study of variances of different decompositions on random states generated by a hardware-efficient *Ansatz* (see text for details). Different colors correspond to different decompositions [Eq. (10)] of the target operator O (see text for the description of all decompositions). Dashed lines are exponential-power-law fits $[\exp(aN^2 + bN + c)]$ to the data (the dominant scaling parameter a is given in the legend).

prepared state (as these are reflection operators, Hadamard test samples match projective measurement samples). The terms in the Pauli decomposition are also directly measured on the prepared state. The GPSK decomposition is constructed as described in Appendix E and measured through a Hadamard test. The von Neumann variance $\text{Var}[O]$ is computed analytically.

The QSP approximation of Ξ (denoted SGN from the sign term approximation) is implemented as described in Appendix D for $R = 20$ and $\delta = 0$. For fair comparison with the other methods, echo verification is not used. The comparison between the Ξ and SGN decomposition shows how the approximation increased the final variance. (The approximation also introduces a bias, see Appendix D.)

All the simulations assume Hadamard-test-based measurement in an ideal circuit simulation: no circuit-level noise is considered and EV is not implemented.

We additionally report scaling results for the shot variances of two other observables, $O = \sum_j jZ_j$ and $O = \sum_j 2^j Z_j$. The overall scaling of all decompositions matches the scaling of the operator norm $\|O\|$. Similarly to the case of Fig. 1, the Ξ decomposition performs best, the SGN approximation has a relatively small effect on the shot variance, and the Pauli decomposition shows the worst scaling.

- [1] S. L. Braunstein and C. M. Caves, Statistical Distance and the Geometry of Quantum States, *Phys. Rev. Lett.* **72**, 3439 (1994).
- [2] V. Giovannetti, S. Lloyd, and L. Maccone, Quantum-enhanced measurements: Beating the standard quantum limit, *Science* **306**, 1330 (2004).
- [3] B. L. Higgins, D. W. Berry, S. D. Bartlett, M. W. Mitchell, H. M. Wiseman, and G. J. Pryde, Demonstrating Heisenberg-limited unambiguous phase estimation without adaptive measurements *New Journal of Physics* Demonstrating Heisenberg-limited unambiguous phase estimation without adaptive measurements, *New J. Phys.* **11**, 073023 (2009).
- [4] E. Knill, G. Ortiz, and R. D. Somma, Optimal quantum measurements of expectation values of observables, *Phys. Rev. A* **75**, 012328 (2007).
- [5] W. J. Huggins, K. Wan, J. McClean, T. E. O'Brien, N. Wiebe, and R. Babbush, Nearly Optimal Quantum Algorithm for Estimating Multiple Expectation Values, *Phys. Rev. Lett.* **129**, 240501 (2022).
- [6] J. Preskill, Quantum Computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
- [7] X. Bonet-Monroig, R. Sagastizabal, M. Singh, and T. E. O'Brien, Low-cost error mitigation by symmetry verification, *Phys. Rev. A* **98**, 062339 (2018).
- [8] S. McArdle, X. Yuan, and S. Benjamin, Error-Mitigated Digital Quantum Simulation, *Phys. Rev. Lett.* **122**, 180501 (2019).
- [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 Inf.* **7**, 23 (2021).
- [10] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, A variational eigenvalue solver on a photonic quantum processor, *Nat. Commun.* **5**, 4213 (2014).
- [11] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, The theory of variational hybrid quantum-classical algorithms, *New J. Phys.* **18**, 023023 (2016).
- [12] T. E. O'Brien, S. Polla, N. C. Rubin, W. J. Huggins, S. McArdle, S. Boixo, J. R. McClean, and R. Babbush, Error mitigation via verified phase estimation, *PRX Quantum* **2**, 020317 (2021).
- [13] Z. Cai, Resource-efficient purification-based quantum error mitigation, [arXiv:2107.07279](https://arxiv.org/abs/2107.07279).
- [14] M. Huo and Y. Li, Dual-state purification for practical quantum error mitigation, *Phys. Rev. A* **105**, 022427 (2022).
- [15] T. E. O'Brien, G. Anselmetti, F. Gkritis, V. E. Elfving, S. Polla, W. J. Huggins, O. Oumarou, K. Kechedzhi, D. Abanin, R. Acharya, *et al.*, Purification-based quantum error mitigation of pair-correlated electron simulations, [arXiv:2210.10799](https://arxiv.org/abs/2210.10799).
- [16] G. H. Low and I. L. Chuang, Optimal Hamiltonian Simulation by Quantum Signal Processing, *Phys. Rev. Lett.* **118**, 010501 (2017).
- [17] G. H. Low and I. L. Chuang, Hamiltonian simulation by qubitization, *Quantum* **3**, 163 (2019).
- [18] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe, Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics, in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing* (ACM, New York, 2019), pp. 193–204.
- [19] D. Wierichs, J. Izaac, C. Wang, and C. Y.-Yu. Lin, General parameter-shift rules for quantum gradients, *Quantum* **6**, 677 (2022).
- [20] G. G. Guerreschi and M. Smelyanskiy, Practical optimization for hybrid quantum-classical algorithms, [arXiv:1701.01450](https://arxiv.org/abs/1701.01450).
- [21] Y. Li and S. C. Benjamin, Efficient Variational Quantum Simulator Incorporating Active Error Minimization, *Phys. Rev. X* **7**, 021050 (2017).

- [22] V. Giovannetti, S. Lloyd, and L. Maccone, Quantum Metrology, *Phys. Rev. Lett.* **96**, 010401 (2006).
- [23] J. Saunders and J.-F. Van Huele, Qubit quantum metrology with limited measurement resources, [arXiv:2108.02876](https://arxiv.org/abs/2108.02876).
- [24] C. Di Franco, M. Paternostro, and M. S. Kim, Hamiltonian Tomography in an Access-Limited Setting without State Initialization, *Phys. Rev. Lett.* **102**, 187203 (2009).
- [25] D. Burgarth, K. Maruyama, and F. Nori, Indirect quantum tomography of quadratic hamiltonians, *New J. Phys.* **13**, 013019 (2011).
- [26] E. Knill and R. Laflamme, Power of One Bit of Quantum Information, *Phys. Rev. Lett.* **81**, 5672 (1998).
- [27] Z. Cai, R. Babbush, S. C. Benjamin, S. Endo, W. J. Huggins, Y. Li, J. R. McClean, and T. E. O'Brien, Quantum error mitigation, [arXiv:2210.00921](https://arxiv.org/abs/2210.00921).
- [28] Y. Gu, Y. Ma, N. Forcellini, and D. E. Liu, Noise-Resilient Phase Estimation with Randomized Compiling, *Phys. Rev. Lett.* **130**, 250601 (2023).
- [29] K. Mitarai and K. Fujii, Methodology for replacing indirect measurements with direct measurements, *Phys. Rev. Res.* **1**, 013006 (2019).
- [30] A. W. Harrow and J. C. Napp, Low-Depth Gradient Measurements Can Improve Convergence in Variational Hybrid Quantum-Classical Algorithms, *Phys. Rev. Lett.* **126**, 140502 (2021).
- [31] G. Ortiz, J. E. Gubernatis, E. Knill, and R. Laflamme, Quantum algorithms for fermionic simulations, *Phys. Rev. A* **64**, 022319 (2001).
- [32] D. Wecker, M. B. Hastings, and M. Troyer, Progress towards practical quantum variational algorithms, *Phys. Rev. A* **92**, 042303 (2015).
- [33] N. C. Rubin, R. Babbush, and J. R. McClean, Application of fermionic marginal constraints to hybrid quantum algorithms, *New J. Phys.* **20**, 053020 (2018).
- [34] V. Verteletskyi, T.-C. Yen, and A. F. Izmaylov, Measurement optimization in the variational quantum eigensolver using a minimum clique cover, *J. Chem. Phys.* **152**, 124114 (2020).
- [35] L. Lin and Y. Tong, Near-optimal ground state preparation, *Quantum* **4**, 372 (2020).
- [36] 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 (London)* **549**, 242 (2017).
- [37] V. Bergholm, J. Izaac, M. Schuld, C. Gogolin, M. S. Alam, S. Ahmed, J. M. Arrazola, C. Blank, A. Delgado, S. Jahangiri, K. McKiernan, J. J. Meyer, Z. Niu, A. Száva, and N. Killoran, PENNYLANE: Automatic differentiation of hybrid quantum-classical computations.
- [38] A. F. Izmaylov, T.-C. Yen, and I. G. Ryabinkin, Revising measurement process in the variational quantum eigensolver: Is it possible to reduce the number of separately measured operators? *Chem. Sci.* **10**, 3746 (2019).
- [39] O. Crawford, B. van Straaten, D. Wang, T. Parks, E. Campbell, and S. Brierley, Efficient quantum measurement of pauli operators in the presence of finite sampling error, *Quantum* **5**, 385 (2021).
- [40] T.-C. Yen, A. Ganeshram, and A. F. Izmaylov, Deterministic improvements of quantum measurements with grouping of compatible operators, non-local transformations, and covariance estimates, *npj Quantum Inf.* **9**, 14 (2023).