

# Quasiprobability decompositions with reduced sampling overhead

**Journal Article****Author(s):**

Piveteau, Christophe; [Sutter, David](#) ; Woerner, Stefan

**Publication date:**

2022

**Permanent link:**

<https://doi.org/https://doi.org/10.3929/ethz-b-000532053>

**Rights / license:**

[Creative Commons Attribution 4.0 International](#)

**Originally published in:**

npj Quantum Information 8(1), <https://doi.org/10.1038/s41534-022-00517-3>

## ARTICLE OPEN



# Quasiprobability decompositions with reduced sampling overhead

Christophe Piveteau<sup>1,2</sup>, David Sutter<sup>1</sup>✉ and Stefan Woerner<sup>1</sup>

Quantum error-mitigation techniques can reduce noise on current quantum hardware without the need for fault-tolerant quantum error correction. For instance, the quasiprobability method simulates a noise-free quantum computer using a noisy one, with the caveat of only producing the correct expected values of observables. The cost of this error mitigation technique manifests as a sampling overhead which scales exponentially in the number of corrected gates. In this work, we present an algorithm based on mathematical optimization that aims to choose the quasiprobability decomposition in a noise-aware manner. This directly leads to a significantly lower basis of the sampling overhead compared to existing approaches. A key element of the novel algorithm is a robust quasiprobability method that allows for a tradeoff between an approximation error and the sampling overhead via semidefinite programming.

npj Quantum Information (2022)8:12; <https://doi.org/10.1038/s41534-022-00517-3>

## INTRODUCTION

Quantum computing holds the promise for significant advancements in various fields such as quantum chemistry, material science, optimization, and machine learning. Unlike classical computers, quantum computers suffer from a significant amount of noise which accumulates over the execution of a quantum algorithm and thus limits experiments to shallow quantum circuits. The theory of quantum error correction and fault-tolerant quantum computing solves this problem. The celebrated threshold theorem ensures that given the noise level of the physical gates is below some constant threshold value, arbitrarily long calculations are possible at arbitrarily low error rates<sup>1</sup>. The cost for a fault-tolerant implementation of a given circuit is a polylogarithmic overhead that is tame when speaking of asymptotics, but currently available quantum hardware does not fulfil the requirements for quantum error correction<sup>2</sup> and it remains a major challenge to achieve this goal.

It is thus a natural question if it is possible to mitigate noise on quantum hardware without the need of full quantum error correction. Multiple schemes have been proposed<sup>3–6</sup> that aim to reduce the effect of noise while also being significantly easier to implement than quantum error correction. All of these methods come with a certain drawback that prohibits them from achieving large-scale fault-tolerant quantum computation. The term quantum error mitigation (QEM) is often used to group these methods. The hope is to use error mitigation techniques to demonstrate a quantum advantage on a useful task with near-term devices.

One specific method in the family of quantum error mitigation techniques is the quasiprobability method<sup>5</sup> (also known as probabilistic error cancellation). The central idea is to decompose the ideal (noise-free) quantum circuit into a quasiprobabilistic mixture of noisy circuits, called quasiprobability decomposition (QPD), that one can implement on a given hardware. These noisy circuits are randomly sampled and run. By correctly post-processing the measurement outcome, one can obtain an unbiased estimator for the expectation

values of the ideal circuit. The quasiprobability method exhibits a simulation overhead in terms of an additional sampling cost. More precisely, the number of shots required to simulate a circuit scales as  $\mathcal{O}(\gamma^{2|C|})$  where  $\gamma \geq 1$  is the so-called sampling overhead of the QPD and  $|C|$  denotes the number of corrected gates in a given circuit  $C$ . This quantity encapsulates how strongly the method has to compensate for the noise in the quantum system. Notably,  $\gamma$  goes to 1 in the limit where the noise vanishes, making the simulation overhead disappear. This exponential cost restricts the quasiprobability method to shallow quantum circuits.

By the above reasoning, it is evident that one wants to find QPDs that exhibit the smallest possible  $\gamma$ -factor. The arguably most difficult part of finding a suitable QPD is how to choose the noisy quantum circuits, called decomposition set, into which we decompose the ideal quantum circuit. It has been realized<sup>5</sup> that the optimal QPD can be expressed as a linear programme, under the assumption that the decomposition set is already fixed. Furthermore, a concrete decomposition set was proposed in ref. <sup>7</sup> that suffices to decompose any circuit under the constraint of arbitrary noise that is not too strong. However, this decomposition set does not necessarily produce the optimal QPD and indeed we demonstrate that there exist decomposition sets that yield a lower sampling overhead. This motivated the study of improved decomposition sets for the quasiprobability method. In ref. <sup>8</sup>, the authors prove analytical bounds on the best possible sampling overhead under the assumption that all unitary and state preparation operations exhibit the same noise.

In the following, we will introduce the quasiprobability method in detail. Consider a linear operator  $\mathcal{F} \in \mathcal{L}(\mathcal{L}(A), \mathcal{L}(B))$ , where  $\mathcal{L}(A, B)$  denotes the set of linear maps from Hilbert spaces  $A$  to  $B$  and  $\mathcal{L}(A) := \mathcal{L}(A, A)$ . A QPD of  $\mathcal{F}$  is a finite set of tuples  $\{(a_i, \mathcal{E}_i)\}_i$  with  $a_i \in \mathbb{R}$  and completely positive maps  $\mathcal{E}_i$ , such that for all  $\rho \in \mathcal{L}(A)$

$$\mathcal{F}(\rho) = \sum_i a_i \mathcal{E}_i(\rho). \quad (1)$$

<sup>1</sup>IBM Quantum, IBM Research – Zurich, Zurich, Switzerland. <sup>2</sup>Institute for Theoretical Physics, ETH Zurich, Zurich, Switzerland. ✉email: dsu@zurich.ibm.com

We call  $\{\mathcal{E}_i\}_i$  the decomposition set of the QPD. The quasiprobability method requires a QPD  $\{(a_i, \mathcal{E}_i)\}_i$  where the maps  $\mathcal{E}_i$  correspond to operations that can be implemented on the quantum hardware, e.g. they might correspond to the channel representing a noisy quantum gate. In practice, these could be obtained by doing tomography and/or by using prior knowledge of the experimental noise. Note that the  $\mathcal{E}_i$  may not need to be trace-preserving as non-trace-preserving maps can be simulated using measurements and postselection.

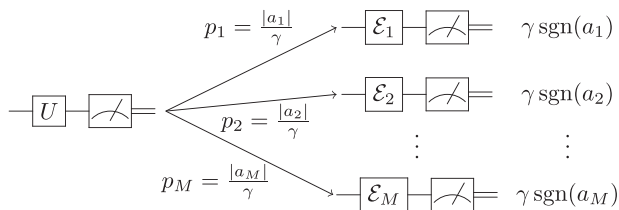
The quasiprobability method allows us to simulate a noiseless execution of  $\mathcal{F}$  using the noisy operations  $\mathcal{E}_i$ . In contrast to quantum error correction, this technique is very hardware-friendly as we do not encode our quantum information in a larger space, so a logical qubit still corresponds to a physical qubit. However, the quasiprobability method is only able to systematically remove the bias from the expectation values and cannot extend the coherence times of the computation. This manifests in a bad asymptotic scaling emphasizing that this technique is most favourable for near-term circuits.

Typically one is interested in the case where  $\mathcal{F}$  is a unitary operation, i.e.  $\mathcal{F} = U$  for some unitary matrix  $U$ , where we use the notation  $\mathcal{U}(\rho) := U\rho U^\dagger$ . This corresponds to simulating an ideal quantum gate  $U$  using noisy operations. But sometimes it can also be interesting to consider  $\mathcal{F}$  which are non-physical, for example  $\mathcal{F}$  could be the inverse of a noise operator that occurred on a certain quantum system. By simulating the inverse, one can revert the effect of the noise on average. Typically the inverse of interesting noise models, such as depolarizing noise, are not completely positive maps. From now on we will restrict ourselves to the case  $\mathcal{F}$  is unitary, though all results also hold in a more generalized setting.

Consider a gate described by a unitary  $U$  and a QPD  $\{(a_i, \mathcal{E}_i)\}_i$  of the unitary channel  $\mathcal{U}$ . Suppose that we are interested in the expectation value of a projective measurement described by a Hermitian operator  $O$ , which would be performed after the gate  $U$ , i.e. we would like to obtain  $\text{tr}[O \mathcal{U}(\rho)]$ , given a certain input state  $\rho$ . The linearity of the trace together with (1) implies

$$\text{tr}[O \mathcal{U}(\rho)] = \gamma \sum_i \frac{|a_i|}{\gamma} \text{sgn}(a_i) \text{tr}[O \mathcal{E}_i(\rho)], \quad (2)$$

where  $\text{sgn}(\cdot)$  denotes the sign function. Via (2) we have introduced the  $\gamma$ -factor  $\gamma := \sum_i |a_i|$ . The right-hand side of (2) naturally gives us a way to construct an unbiased Monte Carlo estimator for  $\text{tr}[O \mathcal{U}(\rho)]$ , while only having access to the operations  $\mathcal{E}_i$  of the noisy hardware, as depicted in Fig. 1. We choose a random number  $i$  with probability  $|a_i|/\gamma$  and then perform the operation  $\mathcal{E}_i$  followed by the measurement  $O$ . The measurement outcome is weighted by the coefficient  $\gamma \text{sgn}(a_i)$



$$\text{tr}[O \mathcal{U}(\rho)] = \sum_{i=1}^M p_i \text{tr}[O \mathcal{E}_i(\rho)] \gamma \text{sgn}(a_i) \quad \forall \rho$$

**Fig. 1 Graphical representation of the quasiprobability method.** The ideal quantum gate  $U$  is randomly replaced by a quantum channel  $\mathcal{E}_i$  that the hardware can implement. The output of the measurement  $O$  must be weighted according to the sampled operation.

which gives the output of the estimator. By sampling this estimator many times and considering the average value, we can obtain an arbitrarily precise estimate of the true expectation value  $\text{tr}[O \mathcal{U}(\rho)]$ .

In practice the  $\mathcal{E}_i$  are imperfect estimates of the true underlying quantum channels, produced by tomography. Unfortunately, tomography fundamentally cannot give us an arbitrarily precise estimate of the true channels, due to state preparation and measurement (SPAM) errors. An erroneous knowledge of the  $\mathcal{E}_i$  will give erroneous coefficients  $a_i$  in the QPD, which again translate into an error in our estimator of the expectation value. It was shown that this problem can be circumvented by using gate set tomography (GST), as it still allows us to obtain an unbiased Monte Carlo estimator<sup>7</sup>, even though the exact  $\mathcal{E}_i$  are unknown. For simplicity, we will assume from now on that we are able to perform exact tomography of the quantum channel  $\mathcal{E}_i$ .

The quasiprobability estimator constructed above does generally not exhibit the identical statistics as the ideal measurement outcome of  $O$ , it only has the same expectation value. In fact, the variance of the estimator increases with  $\gamma$  and one requires  $\mathcal{O}(\gamma^2)$  more shots to estimate  $\text{tr}[O \mathcal{U}(\rho)]$  to a target accuracy compared to the case where  $\mathcal{U}$  is implemented exactly. This is a direct consequence of Hoeffding's inequality.

The quasiprobability method can be applied to large quantum circuits consisting of many different gates, by obtaining a QPD of each quantum gate individually, and then combining them together into one large QPD of the whole circuit. The sampling of the total quasiprobability estimator can still be done efficiently: Consider a circuit consisting of a sequence of  $m$  unitary gates  $\mathcal{U}_m \circ \dots \circ \mathcal{U}_1$  followed by a measurement described by the observable  $O$ . For each operation  $\mathcal{U}_k$  a QPD  $\{(a_i^{(k)}, \mathcal{E}_i^{(k)})\}$  with  $\gamma$ -factor  $\gamma_k$  is given. Our estimator starts by sampling a random number  $i_1$  according to the probabilities  $\{|a_i^{(1)}|/\gamma_1\}$  and executing the operation  $\mathcal{E}_{i_1}^{(1)}$ . In a second step, we sample a random number  $i_2$  according to the probabilities  $\{|a_i^{(2)}|/\gamma_2\}$  and execute the operation  $\mathcal{E}_{i_2}^{(2)}$ . This procedure is continued for all  $i = 3, \dots, m$  while keeping track of all indices  $i_1, \dots, i_m$  sampled along the way. At the very end we measure the observable  $O$  on the system. The estimator then outputs the outcome of that measurement multiplied by  $\text{sgn}(\prod_{k=1}^m a_{i_k}^{(k)}) \prod_{k=1}^m \gamma_k$ .

We see that the combined  $\gamma$ -factor scales in a multiplicative way as  $\gamma_{\text{total}} = \prod_{k=1}^m \gamma_k$ . Therefore the sampling overhead of the total circuit scales exponentially in the circuit size. The Monte Carlo sampler for multiple error mitigated quantum gates only remains an unbiased estimator under the assumption that the noise is localized and Markovian. In looser terms, this means that the noise on any quantum gate must be uncorrelated with other noise and independent on what operations were performed previously on the circuit. Similarly to previous works we will assume for simplicity that this assumption holds<sup>5,7</sup>. Some more recent research has shown that cross-correlations can be tackled by using variants of the quasiprobability method which do not rely on tomography to find the optimal quasiprobability coefficients<sup>8</sup>.

For a given quantum hardware, can we be certain that there exists a QPD of  $\mathcal{U}$  into operations that the hardware can implement? This question may seem difficult, as its answer depends on the exact details of the capabilities of the hardware, namely what kind of quantum operations it can implement and at what fidelity it does so. Fortunately, the answer to this question is positive, as long as the noise is not too strong<sup>7</sup>.

Consider the simplest case where  $\mathcal{U}$  is a single-qubit gate. Table 1 lists 16 single-qubit operations that can be realized by the successive execution of the Hadamard gate  $H$ , the phase gate  $S$ ,

**Table 1.** Sixteen basis operations constituting the standard basis, which was introduced in ref. <sup>7</sup> using the notation  $[U](\cdot) = U(\cdot)U^\dagger$ .

$[1]$				
$[\sigma^x]$		$=$	$[H][S]^2[H]$	
$[\sigma^y]$		$=$	$[H][S]^2[H][S]^2$	
$[\sigma^z]$		$=$	$[S]^2$	
$[R_x]$	$=$	$\frac{[1+i\sigma^x]}{\sqrt{2}}$	$=$	$[H][S]^3[H]$
$[R_y]$	$=$	$\frac{[1+i\sigma^y]}{\sqrt{2}}$	$=$	$[S][H][S]^3[H][S]^3$
$[R_z]$	$=$	$\frac{[1+i\sigma^z]}{\sqrt{2}}$	$=$	$[S]^3$
$[R_{yz}]$	$=$	$\frac{[\sigma^y+i\sigma^z]}{\sqrt{2}}$	$=$	$[H][S]^3[H][S]^2$
$[R_{zx}]$	$=$	$\frac{[\sigma^z+i\sigma^x]}{\sqrt{2}}$	$=$	$[S]^3[H][S]^3[H][S]^3$
$[R_{xy}]$	$=$	$\frac{[\sigma^x+i\sigma^y]}{\sqrt{2}}$	$=$	$[H][S]^2[H][S]^3$
$[\pi_x]$	$=$	$\frac{[1+\sigma^x]}{2}$	$=$	$[S][H][S][H][P_0][H][S]^3[H][S]^3$
$[\pi_y]$	$=$	$\frac{[1+\sigma^y]}{2}$	$=$	$[H][S]^3[H][P_0][H][S][H]$
$[\pi_z]$	$=$	$\frac{[1+\sigma^z]}{2}$	$=$	$[P_0]$
$[\pi_{yz}]$	$=$	$\frac{[\sigma^y+i\sigma^z]}{2}$	$=$	$[S][H][S][H][P_0][H][S][H][S]^3$
$[\pi_{zx}]$	$=$	$\frac{[\sigma^z+i\sigma^x]}{2}$	$=$	$[H][S]^3[H][P_0][H][S][H][S]^2$
$[\pi_{xy}]$	$=$	$\frac{[\sigma^x+i\sigma^y]}{2}$	$=$	$[P_0][H][S]^2[H]$

All these operations can be realized using Hadamard gates  $H$ , phase gates  $S$  and measurement-postselection operations  $P_0$ .

and/or a measurement-postselection operation  $P_0 = |0\rangle\langle 0|$ . The measurement-postselection operation can be simulated in the context of the quasiprobability sampling estimator by performing a measurement in the computational basis and aborting if the outcome is 1, where aborting means that the estimator outputs 0. Therefore any quantum computer able to implement Clifford gates and measurements in the computational basis can also implement these operations, at least approximately. It can be verified that the mentioned set of 16 operations forms a basis of the (real) space of one-qubit Hermitian-preserving operations and hence we call it standard basis. The space spanned by the 16 operations contains only Hermitian-preserving maps, as the 16 basis elements are themselves Hermitian-preserving. Every Hermitian-preserving map has a real  $16 \times 16$  Pauli transfer matrix, which can be linearly decomposed into the Pauli transfer matrices of the 16 basis operations. It can be shown that the standard basis is robust to small errors in the sense that it remains a basis of the space of Hermitian-preserving operations if the individual elements suffer from a small amount of noise<sup>7</sup>.

Consider the setting where we have a fixed decomposition set  $\{\mathcal{E}_i\}_i$  and we would like to find the optimal (in terms of lowest possible  $\gamma$ -factor) quasiprobability coefficients  $a_i$  such that (1) is fulfilled. If our decomposition set consists of linearly independent operations, which is e.g. the case for the standard basis, then this problem corresponds to solving a set of linear equations. But in practice one is often confronted with the more general case where there is an infinity of possible decompositions, so we require some method to pick out the best one of them. This optimization problem can be written in terms of a linear programme (LP)

$$\min_{a_i \in \mathbb{R}} \left\{ \sum_i |a_i| : \mathcal{F} = \sum_i a_i \mathcal{E}_i \right\}, \quad (3)$$

and hence can be solved efficiently<sup>5</sup>. We note that the number of equality constraints in (3) is exponential in the number of qubits involved in the unitary  $\mathcal{U}$ . This exponential cost implies that the LP

can only be solved for few-qubit gates, typically one- and two-qubit gates.

## RESULTS

In this work, we introduce a novel method, which we call Stinespring algorithm, for finding efficient decomposition sets for single-qubit and two-qubit gates. Note that any quantum circuit can be decomposed into single-qubit and two-qubit gates, therefore it is sufficient to obtain quasiprobability decompositions of these gates in order to error mitigate an arbitrary quantum circuit. Based on mathematical optimization techniques for convex and non-convex problems this new iterative algorithm takes into account the hardware noise. As a building block for the Stinespring algorithm, we introduce a robust quasiprobability method called approximate quasiprobability decomposition, which is interesting on its own. Instead of perfectly simulating a certain circuit, we allow for a small approximation error. This enables us to make a tradeoff between the approximation quality and the sampling overhead of the quasiprobability method. We will illustrate our results with simulations showing that the new methods significantly reduce the  $\gamma$ -factor of the QPD compared to existing approaches.

### Approximate quasiprobability decomposition

Our first result is a relaxation of the QPD where we allow for an error in the decomposition. In mathematical terms, we require the condition (1) to hold approximately, i.e.,

$$\mathcal{U}(\rho) \approx \sum_i a_i \mathcal{E}_i(\rho). \quad (4)$$

An erroneous QPD leads to an error in the Monte Carlo estimator. Giving up the exactness of the quasiprobability method allows us to find a better decomposition in terms of the  $\gamma$ -factor. More precisely the approximate QPD results in a tradeoff between the sampling overhead and the error in the method.

In a first step we have to quantify the error in the approximate QPD given by (4). A natural candidate is to use the diamond norm as it has a strong operational interpretation. In addition, the diamond norm fits very naturally in our mathematical optimization setting, as it has been shown to be expressible as a semidefinite programme (SDP)<sup>10</sup>.

**Theorem 2.1** (SDP for diamond norm<sup>10</sup>). *Let  $\mathcal{G} \in \mathcal{L}(\mathcal{L}(A), \mathcal{L}(B))$  and denote its Choi matrix by  $\Lambda_{\mathcal{G}}$ . Then*

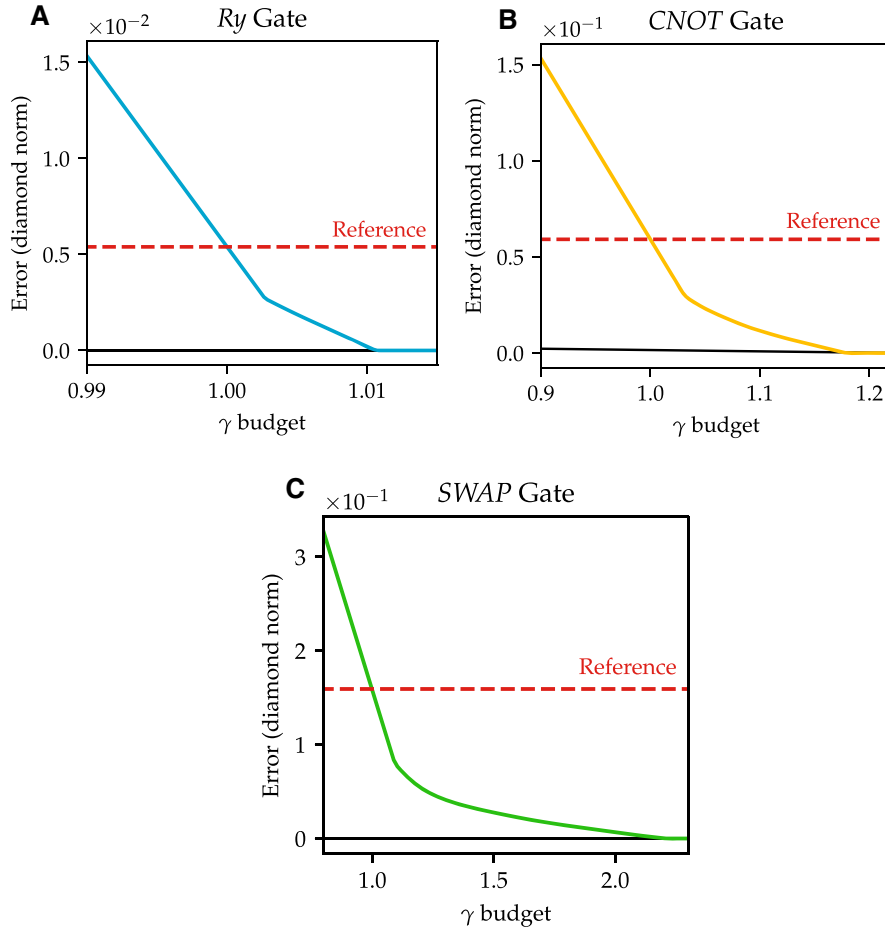
$$\|\mathcal{G}\|_{\diamond} = \begin{cases} \max_{\substack{\rho_0, \rho_1 \in \mathcal{L}(A) \\ X \in \mathcal{L}(B \otimes A)}} & \frac{1}{2} \langle \Lambda_{\mathcal{G}}, X \rangle + \frac{1}{2} \langle \Lambda_{\mathcal{G}}^*, X^* \rangle \\ \text{s.t.} & \begin{pmatrix} \mathbb{1}_B \otimes \rho_0 & X \\ X^* & \mathbb{1}_B \otimes \rho_1 \end{pmatrix} \geq 0 \\ & \rho_0 \geq 0, \rho_1 \geq 0, \rho_0^\dagger = \rho_0, \rho_1^\dagger = \rho_1, \end{cases} \quad (5)$$

which is a SDP. (Note that  $X^*$  and  $\Lambda_{\mathcal{G}}^*$  denote the adjoint operators of  $X$  and  $\Lambda_{\mathcal{G}}$ , respectively.)

The dual formulation of the SDP (5) is given by

$$\|\mathcal{G}\|_{\diamond} = \begin{cases} \min_{Y_0, Y_1 \in \mathcal{L}(B \otimes A)} & \frac{1}{2} \|\text{tr}_B[Y_0]\|_{\infty} + \frac{1}{2} \|\text{tr}_B[Y_1]\|_{\infty} \\ \text{s.t.} & \begin{pmatrix} Y_0 & -\Lambda_{\mathcal{G}} \\ -\Lambda_{\mathcal{G}}^* & Y_1 \end{pmatrix} \geq 0 \\ & Y_0 \geq 0, Y_1 \geq 0, \end{cases} \quad (6)$$

where  $\|\cdot\|_{\infty}$  denotes the spectral norm. Suppose we have a fixed decomposition set  $\{\mathcal{E}_i\}_i$  and we would like to find the best possible approximate QPD of an operation  $\mathcal{U}$  into that decomposition set. More precisely, we give a certain  $\gamma$ -factor budget, denoted  $\gamma_{\text{budget}}$  such that the QPD may not have a  $\gamma$ -factor higher



**Fig. 2 Simulation of Tradeoff curves.** Panels **A**, **B** and **C** depict the tradeoff curves for three quantum gates Ry (with an arbitrary angle since the noise model is the same for all rotation angles), CNOT and SWAP respectively using the standard basis as decomposition set. The noisy channels of the three gates and the standard basis are extracted from a noise model in Qiskit<sup>21</sup> which approximates the noise on the `ibmq_melbourne` device. More precisely, the noise model consists of a combination of depolarizing and thermal relaxation errors where the noise parameters are inferred from physical quantities measured on the hardware (e.g.  $T_1$ ,  $T_2$ , gate times, qubit frequency, etc.). The CNOT gates exhibit an error rate of around 2% and the single-qubit gates exhibit an error rate of around 0.1%. The red dashed line represents the error (in diamond norm) of the reference noisy channel, i.e. when the gate is implemented as-is without QEM. We use the SDP solver in the MOSEK<sup>22</sup> software package through the CVXPY<sup>23,24</sup> modelling language.

than that. The corresponding optimization problem is

$$\begin{cases} \min_{a_i \in \mathbb{R}} & \|\mathcal{U} - \sum_i a_i \mathcal{E}_i\|_\diamond \\ \text{s.t.} & \sum_i |a_i| \leq \gamma_{\text{budget}} \end{cases} \quad (7)$$

By inserting the SDP from (6) into (7) the problem can be rewritten as

$$\begin{cases} \min_{a_i \in \mathbb{R}, Y_0, Y_1 \in \mathcal{L}(A \otimes B)} & \frac{1}{2} \|\text{tr}_B[Y_0]\|_\infty + \frac{1}{2} \|\text{tr}_B[Y_1]\|_\infty \\ \text{s.t.} & \left( Y_0 \sum_i a_i \Lambda_{\mathcal{E}_i} - \Lambda_{\mathcal{U}} \sum_i a_i \Lambda_{\mathcal{E}_i}^* - \Lambda_{\mathcal{U}}^* Y_1 \right) \geq 0 \\ & \sum_i |a_i| \leq \gamma_{\text{budget}} \\ & Y_0 \geq 0, Y_1 \geq 0, \end{cases} \quad (8)$$

where  $\Lambda_{\mathcal{U}}$  and  $\Lambda_{\mathcal{E}_i}$  are the Choi matrices of  $\mathcal{U}$  and  $\mathcal{E}_i$ , respectively. Luckily (8) is still a SDP, which allows us to efficiently solve our optimization problem. The approximation  $\mathcal{U}_{\text{approx}} := \sum_i a_i \mathcal{E}_i$  of  $\mathcal{U}$  obtained from the optimization problem in (7) is not guaranteed to be a physical map. In practice, this implies that we might simulate the execution of a non-physical map using the

quasiprobability method. It is possible to enforce complete positivity and/or trace-preservingness of  $\mathcal{U}_{\text{approx}}$  into the optimization problem by adding a positivity/partial-trace constraint on the Choi matrix  $\Lambda_{\mathcal{U}_{\text{approx}}}$  of  $\mathcal{U}_{\text{approx}}$ .

We note that the idea of considering approximate QPDs has been considered in some form or another in other works. For instance, in ref. <sup>9</sup>, the quasiprobability coefficients of a complete circuit are optimized all at once in order to minimize the error of the computation. This optimisation problem is very difficult and does not come with the same properties and guarantees as a SDP. In ref. <sup>11</sup>, the authors utilize the quasiprobability method to simulate noise with reduced strength in order to be used in conjunction with the error extrapolation method. The used QPD is exact and not approximate, but the target channel  $\mathcal{F}$  is chosen to be noisy instead of an ideal unitary channel.

One can solve the SDP (7) for different values of  $\gamma_{\text{budget}}$  to obtain a relation between the diamond norm error and the  $\gamma$ -factor. This function, which we call tradeoff curve, encapsulates the tradeoff between the approximation error and the sampling overhead. Example tradeoff curves for three gates are computed using a specific noise model and depicted in Fig. 2.

As expected, if the  $\gamma$ -factor budget is larger than the optimal  $\gamma$ -factor of the non-approximate QPD, the error becomes zero. Similarly, when the  $\gamma$ -factor budget is exactly 1, then one does not gain any advantage over implementing the gate as-is without QEM. The more interesting regime is in between these two values of  $\gamma_{\text{budget}}$ : One clearly sees that we can still significantly reduce the error of a gate, without having to pay the full  $\gamma$ -factor necessary for a non-approximate QPD. For the SWAP gate, the exact QPD requires a  $\gamma$ -factor of 2.21 to completely correct the gate. However, if we only pay a  $\gamma$ -factor of 1.21, we can still reduce the error by 67%. The saved costs in terms of the sampling overhead are substantial, since the number of shots scales as  $\gamma^{2|C|}$ , where  $|C|$  is the number of gates.

If one applies the approximate quasiprobability method to a circuit with multiple gates, a new degree of freedom emerges, which is not present in the original formulation of the quasiprobability method: How much  $\gamma$ -factor budget do we give to every individual gate? Assume we have a budget  $\gamma_{\text{total}}$  for the whole circuit, how do we distribute that budget optimally across the whole circuit? More concretely, given  $N$  gates we have to find individual budgets  $\gamma_{\text{budget},i} \geq 0$  for the  $i$ -th gate where  $i = 1, \dots, N$  such that  $\prod_{i=1}^N \gamma_{\text{budget},i} = \gamma_{\text{total}}$ . This is discussed in more detail in the Supplementary Information.

### Stinespring algorithm

For a fixed decomposition set an optimal QPD can be found by solving the LP (3). To further reduce the  $\gamma$ -factor it would be advantageous to also find the optimal decomposition set for a given hardware noise. This task is non-trivial for at least two reasons:

1. It is difficult to optimize over the space of all possible operations that a given quantum hardware can perform. In general, this space is large and might vary significantly from one machine to another.
2. Performing tomography to characterize the operations  $\mathcal{E}_i$  is expensive. Hence, we want to limit the number of required uses of tomography.

In this section, we will introduce a novel method, called the Stinespring algorithm, that is able to deal with both issues. While the Stinespring algorithm is generally not able to produce the optimal decomposition set, it can be seen as an empirical approximation thereof. We demonstrate with simulations that the obtained decomposition set exhibits a significantly reduced  $\gamma$ -factor compared to the standard basis. The algorithm relies on a fundamental result in quantum information theory, stating that any quantum channel can be expressed as a unitary evolution on some extended Hilbert space.

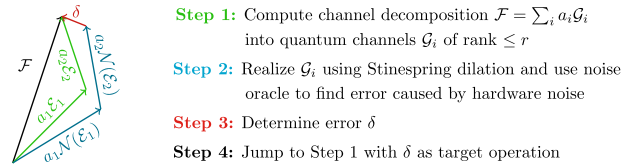
**Theorem 2.2** (Stinespring dilation). *Consider a trace-preserving completely positive (TPCP) map  $\mathcal{E} \in \text{TPCP}(A, A)$ . There exists a Hilbert space  $R$  and an isometry  $V \in L(A, A \otimes R)$  such that for all density matrices  $\rho$*

$$\mathcal{E}(\rho) = \text{tr}_R[V\rho V^\dagger].$$

Furthermore, there exists an isometry with  $\dim(R) \leq r$ , where  $r$  is the rank of the quantum channel defined by  $r := \text{rank}(\Lambda_{\mathcal{E}})$  for  $\Lambda_{\mathcal{E}}$  the Choi matrix of  $\mathcal{E}$ .

Any isometry  $V$  can be extended (generally non-uniquely) to a unitary  $U \in U(A \otimes R, A \otimes R)$  such that  $U$  acts equivalently to  $V$  on the space of states of the form  $\rho_A \otimes |0\rangle\langle 0|_R$ .

The core idea of the Stinespring algorithm is to find the set of general quantum channels  $\{\mathcal{E}_i\}$  that allows the realization of a QPD with lowest possible  $\gamma$ -factors. These quantum channels are then approximately implemented on the quantum hardware by use of a Stinespring dilation. (We can straightforwardly use the Stinespring dilation to approximate an arbitrary quantum channel on a quantum computer by making use of clean ancilla qubits,



**Fig. 3 Stinespring algorithm.** Conceptual visualization of the iterative process involved in the Stinespring algorithm.

performing a circuit corresponding to the unitary  $U$  on the extended space, and finally discarding the ancilla qubits.) However, due to the significant hardware noise, the realized operation will only be an approximation of the desired quantum channel. The approximation error of the dilation is then accounted for in an iterative manner.

In practice, we will restrict the Stinespring algorithm to generate decomposition sets for error-mitigation of one- and two-qubit gates, as the involved computation and tomography quickly become expensive for larger gates. This does not restrict the applicability of the method, as any multi-qubit unitary can be decomposed in single and two-qubit gates.

The result in the section ‘Channel decomposition’ allows us to quasiprobabilistically decompose any TPCP map  $\mathcal{F}$ , which could for example be our desired unitary  $\mathcal{U}$ , into rank  $r$  quantum channels that can be approximated with  $\lceil \log_2 r \rceil$  ancilla qubits. By choosing  $r$  small enough and by making use of a technique called variational unitary approximation introduced in the section ‘Variational unitary approximation’, we can ensure that this approximation is reasonably good. Still, there is an important step missing before we can practically use this result. The quasiprobability method requires a QPD where the elements of the decomposition set correspond to channels describing noisy operations which the actual quantum hardware can execute. However, due to the noisy nature of the hardware, we can only execute an approximation of the channels  $\mathcal{G}_i$  obtained from the channel decomposition. We have to take into account the inaccuracy of the implemented Stinespring dilation when constructing our decomposition set. This will be achieved by the use of an iterative algorithm.

Assume that we have access to a noise oracle  $\mathcal{G} \mapsto \mathcal{N}(\mathcal{G})$  which gives us the actual quantum channel executed by the hardware when we implement the channel  $\mathcal{G}$  using a Stinespring dilation. In general, since we are considering single and two-qubit gates, this oracle can be realized by simply performing tomography of the noisy realisations of the  $\mathcal{G}_i$  on the hardware. Assume that we found a channel decomposition

$$\mathcal{F} = \sum_{i=1}^{n_{\text{pos}}} a_i^+ \mathcal{G}_i^+ - \sum_{i=1}^{n_{\text{neg}}} a_i^- \mathcal{G}_i^-$$

using the rank-constrained optimization (10). If we were to implement  $\mathcal{F}$  with the quasiprobability method by using the Stinespring approximation of the involved channels  $\mathcal{G}_i^\pm$ , an error  $\delta$  would occur where

$$\delta := \mathcal{F} - \sum_{i=1}^{n_{\text{pos}}} a_i^+ \mathcal{N}(\mathcal{G}_i^+) - \sum_{i=1}^{n_{\text{neg}}} a_i^- \mathcal{N}(\mathcal{G}_i^-).$$

We can iteratively repeat our procedure, but this time decomposing  $\delta$  instead of  $\mathcal{F}$ . During each point of the iteration we store the  $\mathcal{N}(\mathcal{G}_i^\pm)$  into our decomposition set. The  $\delta$  can be obtained by performing an approximate QPD of the target unitary channel  $\mathcal{U}$  into the decomposition set. At some point our decomposition set will be large enough such that the resulting approximation error is smaller than a desired threshold. A conceptual visualization of the procedure is depicted in Fig. 3 and the exact algorithm can be found in Algorithm 1.

**Algorithm 1:** Stinespring algorithm.

---

**Given:** a target unitary operation  $[U]$  for  $U \in U(A)$ , a noise oracle  $\mathcal{N}$ , a threshold  $\Delta_{\text{threshold}}$  for the diamond norm error;

**Result:** A decomposition set for a QPD with low sampling overhead

$\mathcal{D} := \{\mathcal{N}([U])\}$ ;  
 $\mathcal{F} := [U]$ ;

**repeat**

$\{(a_j, D_j)\}_j :=$  compute approximate QPD coefficients of  $\mathcal{F}$  using decomposition set  $\mathcal{D}$ ;

$\Delta :=$  get diamond norm error of approximate QPD  $\{(a_j, D_j)\}_j$ ;

**if**  $\Delta < \Delta_{\text{threshold}}$  **then**

| break;

**end**

$\delta :=$  get the remaining error of the approximate QPD  $\{(a_j, D_j)\}_j$ ;

$\{\mathcal{G}_i\}_i :=$  perform rank-constrained channel decomposition of  $\delta$  into a set of channels;

$\{V_i\}_i :=$  get Stinespring dilation isometries of the channels  $\{\mathcal{G}_i\}_i$ ;

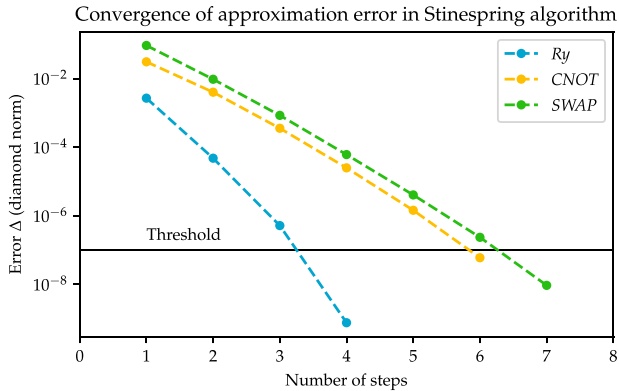
$\{C_i\}_i :=$  get variational unitary approximation circuits of the isometries  $\{V_i\}_i$ ;

$\{\mathcal{E}_i\}_i :=$  apply noise oracle  $\mathcal{N}$  on the circuits  $\{C_i\}_i$ ;

$\mathcal{D} := \mathcal{D} \cup \{\mathcal{E}_i\}_i$ ;

**return**  $\mathcal{D}$

---



**Fig. 4** Evolution of the approximation error in each step of the Stinespring algorithm. Three different runs of the algorithm on three different gates Ry, CNOT and SWAP are depicted. The hardware noise is estimated using a noise model approximating the `ibmq_melbourne` device. The horizontal grey line denotes the threshold of  $10^{-7}$  at which the algorithm stops.

It is crucial that the Stinespring dilation allows for a reasonably accurate approximation of a desired quantum channel. In other words, the blue and green arrows in Fig. 3 have to be sufficiently close. In the simulations presented next, we observed that this was only the case with the inclusion of the rank constraint in the channel decomposition. The variational unitary approximation technique significantly improves the approximation further.

We next show the performance of the Stinespring algorithm via simulation results on the gates Ry, CNOT and SWAP, that we already analysed in Fig. 2. For all three gates we enforce a rank constraint  $r=2$  during the channel decomposition. This corresponds to allowing for at most one ancilla qubit during the Stinespring dilation. For the two-qubit gates, we additionally make use of the variational unitary approximation with a three-qubit RyRz variational form of depth 6 with linear qubit connectivity. The noise oracle is obtained from the noise model that was already used in Fig. 2, which aims to approximate the noise on the `ibmq_melbourne` device. Using a noise model instead of a full tomography significantly speeds up the simulation. We use a threshold of  $\Delta_{\text{threshold}} = 10^{-7}$  to stop the iterative procedure. During each iteration of the Stinespring algorithm, we store the diamond norm error of the current approximate QPD (denoted  $\Delta$  in Algorithm 1). Figure 4 shows how this error decreases during

**Table 2.** Simulation results of the Stinespring algorithm applied on the three quantum gates Ry, CNOT and SWAP.

	$\gamma$ -factor	
	Standard basis + noisy gate	Stinespring algorithm
Ry	1.0106	1.0056
CNOT	1.1789	1.0815
SWAP	2.2095	1.2294

The obtained  $\gamma$ -factor is compared to the value obtained when the respective gates are decomposed into the standard basis.

the Stinespring algorithm. It is clearly visible that this decrease is exponential.

As discussed in detail in the Supplementary Information, each iteration of the Stinespring algorithm applied on the two-qubit gates extends the decomposition set by 16 new operations. By considering that we need around 6–7 steps to reach the desired threshold, this implies that the produced decomposition set is significantly smaller than the standard basis (around 70–80 elements instead of 256). This result is remarkable and indicates that the Stinespring algorithm really does find a decomposition set that is well adapted to the hardware noise. As a reminder, the 256 elements in the standard basis were needed to completely span the space of Hermitian-preserving operators. The decomposition set produced by the Stinespring algorithm spans a significantly smaller space, as it is tailored to only represent one specific operation.

Table 2 denotes the  $\gamma$ -factors obtained by finding the optimal QPD using the decomposition set produced by the Stinespring algorithm for the three gates in question. We see that a significant improvement can be observed compared to using the decomposition set consisting of the standard basis and the respective noisy gate. As previously discussed, the sampling overhead of the quasiprobability method scales exponentially in the number of gates, where the  $\gamma$ -factor forms the basis of that exponential cost. Therefore any reduction in the  $\gamma$ -factor is significant and allows us to implement considerably larger quantum circuits.

Consider for instance that we have a circuit consisting of gates with equal sampling overhead  $\gamma = 1 + \kappa$  and we wish to fix the total sampling overhead  $\Gamma = (1 + \kappa)^{2n_{\text{gates}}}$ . Then the number of gates  $n_{\text{gates}}$  that one can error mitigate without exceeding the total sampling overhead  $\Gamma$  is given by  $\frac{1}{2} \frac{\ln \Gamma}{\ln(1 + \kappa)}$  which is  $\approx \frac{1}{2} \frac{\ln \Gamma}{\kappa}$  for  $\kappa \approx$

**Table 3.** Simulation results of the Stinespring algorithm applied on the CNOT gate using noise models from different IBMQ hardware backends.

Hardware Backend	CNOT		$\gamma$ -factor
	Error rate	Standard basis + noisy CNOT	
ibmq_melbourne	2.03%	1.1789	1.0815
ibmq_mumbai	1.12%	1.0968	1.0486
ibmq_sydney	0.96%	1.0667	1.0415

The obtained  $\gamma$ -factor is compared to the value obtained when the respective gates are decomposed into the standard basis.

0. So a reduction of  $\kappa$  by a factor 2, which is roughly what we achieved for the Ry and CNOT gates, corresponds to doubling the number of gates that we can error mitigate.

Finally, we repeat the experiment on the CNOT gate for two additional noise models. We again consider noise models from Qiskit which approximate the noise on the `ibmq_sydney` and `ibmq_mumbai` hardware platforms. These two chips are newer and more advanced than `ibmq_melbourne` and therefore exhibit lower noise. The results are depicted in Table 3 and one sees that the Stinespring algorithm again provides a decomposition set with a significantly reduced  $\gamma$ -factor.

We do not have a guarantee that the Stinespring algorithm always converges. Such a proof would require a formalization of the heuristics that were made during the construction of the algorithm. In the Supplementary Information we summarize these heuristics in detail.

## DISCUSSION

The quasiprobability method is an error mitigation technique that allows us to suppress errors on small-scale devices without the need of implementing universal fault-tolerant quantum computing. The cost is a sampling overhead scaling as  $\mathcal{O}(\gamma^{2|C|})$ , where  $\gamma \geq 1$  is the  $\gamma$ -factor and  $|C|$  denotes the number of gates in the circuit that need to be error mitigated. We presented two novel methods to reduce the  $\gamma$ -factor. Since the sampling overhead scales exponentially in the number of gates with the  $\gamma$ -factor being the basis of the exponent, reducing  $\gamma$  is crucial.

In this work, we presented two improvements of the QPD method that both allow to reduce the overhead by minimizing the  $\gamma$ -parameter. The first contribution, i.e., the approximate QPD, allows us to elegantly generalize the quasiprobability method in an approximate setting, where one can make a tradeoff between an approximation error and the sampling overhead. Since in many practical applications it is very difficult to obtain exact estimates of the hardware noise, the quasiprobability method will anyways incur a significant error. In this setting it might be possible to reduce the sampling overhead without significantly increasing the error by making use of our approximate quasiprobability method.

Whereas the approximate QPD approach is fully analytical and mathematically rigorous, the second improvement (the Stinespring algorithm) relies on some numerical heuristics that we tested on various practical problems and verified its usefulness via simulation results. Nevertheless, it would be important to prove the convergence or even the convergence rate of the Stinespring algorithm. We leave this for future work. Our results indicate that this approach could have the potential to significantly reduce the sampling overhead compared to existing methods, as can be seen from a simple calculation. Suppose we have a quantum circuit

where we want to correct as many SWAP gates as possible with a total sampling overhead that is not larger than  $10^3$ . (Note that due to limited connectivity on near-term devices SWAP gates will occur very often.) Table 2 shows that via the previous QPD approach we can correct 4 SWAP gates, whereas our improved QPD method allows us to correct up to 16 SWAP gates. (Recall that the overhead scales as  $\mathcal{O}(\gamma^{2|C|})$ , where  $|C|$  is the number of mitigated gates.) From an experimental point of view this technique also has the advantage that, in contrary to the standard basis, no measurements have to be executed on the data qubits.

Due to its inherent exponential scaling, neither the QPD nor other error mitigation techniques are believed to be able to replace the need for error correction for complicated quantum computations. However, error mitigation techniques may be able to smoothen the transition into the fault-tolerant quantum computing era. Hence as a next step for future work, it would be interesting to see how the improved QPD techniques presented in this work can be combined with error correction ideas.

## METHODS

### Variational unitary approximation

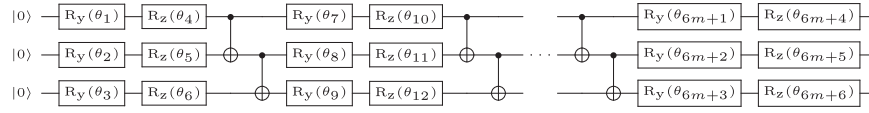
To optimize the approximate execution of a quantum channel via the Stinespring dilation, we make use of two central observations:

1. Because the choice of the unitary  $U$  extending the isometry  $V$  is not unique we can try to choose the  $U$  which requires the least amount of two-qubit gates to implement.
2. Since the hardware is noisy, it may not be advantageous to implement a circuit that represents  $U$  exactly. It might be better to use a circuit that approximates  $U$  and requires less two-qubit gates and thus suffers less from noise on the hardware.

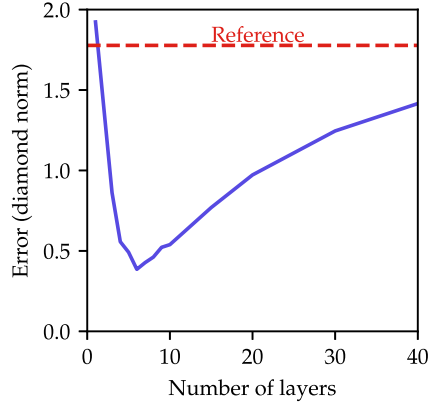
In this section, we propose a method that makes use of both of these observations, which we call variational unitary approximation. We use a variational circuit to implement the dilation unitary and fit its parameters in order to approximate the Stinespring isometry  $V$  as well as possible. Denote by  $\theta$  the tuple of all variational parameters and  $U_{\text{var}}(\theta)$  the unitary represented by the variational form. Furthermore, we denote by  $V_{\text{var}}(\theta)$  the submatrix of  $U_{\text{var}}(\theta)$  restricted on the subspace where the input ancilla qubits are fixed to the zero state. We want to choose our parameters  $\theta$  such that we minimize the difference between  $V_{\text{var}}(\theta)$  and  $V$ , where this optimization is done with a classical computer. The variational unitary approximation allows us to freely choose the expressiveness of our variational circuit. More concretely, we can tune how many two-qubit gates it shall contain and therefore influence the tradeoff between the approximation error and the hardware noise error. Variational unitary approximations are studied in various settings such as<sup>12</sup>, where the interested reader can find more information.

In the following, we are going to look at a concrete example to illustrate the gains from the variational unitary approximation. Consider an arbitrary two-qubit quantum channel with rank 2 that we want to approximate with a Stinespring dilation. Let  $U$  be a three-qubit unitary realizing a Stinespring dilation of the given channel. The exact representation of  $U$  typically requires around 100 CNOT gates (assuming linear qubit connectivity), which is significantly more than what current hardware can reasonably handle. We replace this ideal circuit with a RyRz variational circuit depicted in Fig. 5, where  $m$  denotes the depth of the variational form, the total number of parameters is  $6(m+1)$  and the total number of CNOT gates is  $2m$ . We use the gradient-based BFGS algorithm implemented in SciPy<sup>13</sup> to minimize the objective  $\|V - V_{\text{var}}(\theta)\|_2$ . In order to avoid having to estimate the gradient with the finite differences method, we make use of the Autograd<sup>14</sup> package for automatic differentiation. As initial guess for the parameters, we use uniformly random numbers. The objective is non-convex and in practice the optimization algorithm finds a different local minimum depending on the provided initial guess. To obtain a good result in practice, we observed that it is enough to repeat the optimization 5 times for different initial values and then take the best result.

For purpose of numerical demonstration, the above procedure is performed for different depths of the variational form on a Haar-random three-qubit unitary  $U$ . By using the same noise model as in Fig. 2, we can estimate how well our method allows us to approximate the two-qubit quantum channel  $\rho \mapsto \text{tr}_3[U\rho U^\dagger]$  where  $\text{tr}_3$  stands for tracing out the third



**Fig. 5 RyRz variational circuit.** Variational circuit used to approximate the Stinespring unitary. It is parametrized by the angles  $\theta_1, \dots, \theta_{6(m+1)}$ .



**Fig. 6 Numerical results from the variational unitary approximation of a Haar-random three-qubit unitary  $U$  using a RyRz variational form.** This plot depicts the diamond norm error of the induced two-qubit quantum channel for different depths of the variational form. The depicted error contains the effects of the hardware noise and of the approximation error of the variational unitary. The noise model used is the same as in Fig. 2. The dashed red reference line corresponds to the error obtained when  $U$  is (exactly) decomposed in single and two-qubit gates (without the variational unitary approximation) and then executed on the hardware.

qubit. More precisely, we can compute the diamond norm error between the obtained two-qubit quantum channel and the ideal two-qubit quantum channel. This error encapsulates the approximation error of the variational form combined with the hardware noise. The results can be seen in Fig. 6. The variational unitary approximation technique allows us to significantly reduce the error to roughly one-quarter of its reference value. One can also clearly see that there is a sweet spot in the variational circuit depth. This makes sense intuitively when considering the tradeoff mentioned previously: if the depth is too short, then  $U$  is not well represented by the variational form and if the depth is too long, then the hardware noise takes over and we start to mostly sample noise. The optimum is reached at a variational circuit depth of 6.

### Channel decomposition

Consider a trace-preserving linear operator  $\mathcal{F} \in \mathcal{L}(\mathcal{L}(A), \mathcal{L}(B))$  of which we want to find an optimal QPD. In this section, we will generalize the LP (3) by not only optimizing over the quasiprobability coefficients  $a_i$ , but also over the decomposition set  $\{\mathcal{E}_i\}_i$ . The resulting QPD will not be directly usable for the quasiprobability method, as the obtained decomposition will generally not be implementable by a quantum computer in practice. To highlight this, we will denote the decomposition set obtained by this optimisation procedure by  $\{\mathcal{G}_i\}_i$  instead of  $\{\mathcal{E}_i\}_i$  throughout this section. We call the resulting QPD the channel decomposition of  $\mathcal{F}$ . We note that this decomposition was independently derived in refs. 15,16. The channel decomposition will be an important step towards the realization of the Stinespring algorithm.

We denote the Choi matrix of  $\mathcal{F}$  by  $\Lambda_{\mathcal{F}}$  and our goal is now to construct a finite set of Choi matrices  $\{\Lambda_{\mathcal{G}_i}\}_i$  which correspond to a decomposition set  $\{\mathcal{G}_i\}_i$ . The Choi representation is very convenient, as it allows us to formulate the TPCP-condition on the decomposition basis in a straightforward way:  $\Lambda_{\mathcal{G}_i} \geq 0$  and  $\text{tr}_2[\Lambda_{\mathcal{G}_i}] = \frac{1}{2^n} \mathbb{1}$  for all  $i$  where  $\text{tr}_2$  stands for the partial trace over the ancillary Hilbert space of the Choi-Jamiolkowski isomorphism and  $n$  is the number of qubits involved in  $\mathcal{F}$ . The

optimization problem at hand can be expressed as

$$\left\{ \begin{array}{ll} \min_{a_i \in \mathbb{R}, \Lambda_{\mathcal{G}_i} \in \mathbb{C}^{4^n \times 4^n}} & \sum_{j=1}^N |a_j| \\ \text{s.t.} & \Lambda_{\mathcal{F}} = \sum_{j=1}^N a_j \Lambda_{\mathcal{G}_j} \\ & \Lambda_{\mathcal{G}_i} \geq 0 \\ & \text{tr}_2[\Lambda_{\mathcal{G}_i}] = \frac{1}{2^n} \mathbb{1}. \end{array} \right. \quad (9)$$

By grouping all positive and negative quasiprobability coefficients together

$$a^+ := \sum_i a_i^+, \quad a^- := \sum_i a_i^-, \quad \mathcal{G}^+ := \frac{1}{a^+} \sum_i a_i^+ \mathcal{G}_i^+, \quad \text{and} \quad \mathcal{G}^- := \frac{1}{a^-} \sum_i a_i^- \mathcal{G}_i^-,$$

one can see that choosing  $N=2$  is enough to reach the optimum. By performing the substitution  $\tilde{\Lambda}^{\pm} := a^{\pm} \Lambda_{\mathcal{G}^{\pm}}$  we can transform the problem into a SDP which can be solved efficiently

$$\left\{ \begin{array}{ll} \min_{a_i \in \mathbb{R}, \Lambda_{\mathcal{G}_i} \in \mathbb{C}^{4^n \times 4^n}} & a^+ + a^- \\ \text{s.t.} & \Lambda_{\mathcal{F}} = \tilde{\Lambda}_{\mathcal{G}^+} - \tilde{\Lambda}_{\mathcal{G}^-} \\ & \tilde{\Lambda}_{\mathcal{G}^{\pm}} \geq 0 \\ & \text{tr}_2[\tilde{\Lambda}_{\mathcal{G}^{\pm}}] = a^{\pm} \frac{1}{2^n} \mathbb{1}. \end{array} \right.$$

The channel decomposition asserts that we can find TPCP maps  $\mathcal{G}^+, \mathcal{G}^-$  and  $a^+, a^- \geq 0$  such that  $\mathcal{F} = a^+ \mathcal{G}^+ - a^- \mathcal{G}^-$  with optimal  $\gamma$ -factor  $\gamma = a^+ + a^-$ .

For the sake of the Stinespring algorithm, we would like to approximate the  $\mathcal{G}^{\pm}$  using a Stinespring dilation. This will only work reasonably well when the number of required ancilla qubits is not too large. We can enforce this by adding an additional constraint  $\text{rank}(\tilde{\Lambda}_{\mathcal{G}_i}^{\pm}) \leq r$  for some  $r \in \mathbb{N}$  into the SDP (9), i.e.,

$$\left\{ \begin{array}{ll} \min_{a_i^{\pm} \in \mathbb{R}^+, \tilde{\Lambda}_{\mathcal{G}_i}^{\pm} \in \mathbb{C}^{4^n \times 4^n}} & \sum_{i=1}^{n_{\text{pos}}} a_i^+ + \sum_{i=1}^{n_{\text{neg}}} a_i^- \\ \text{s.t.} & \Lambda_{\mathcal{F}} = \sum_{i=1}^{n_{\text{pos}}} \tilde{\Lambda}_{\mathcal{G}_i}^+ - \sum_{i=1}^{n_{\text{neg}}} \tilde{\Lambda}_{\mathcal{G}_i}^- \\ & \tilde{\Lambda}_{\mathcal{G}_i}^{\pm} \geq 0 \\ & \text{tr}_2[\tilde{\Lambda}_{\mathcal{G}_i}^{\pm}] = a_i^{\pm} \frac{1}{2^n} \mathbb{1} \\ & \text{rank}(\tilde{\Lambda}_{\mathcal{G}_i}^{\pm}) \leq r, \end{array} \right. \quad (10)$$

where  $\tilde{\Lambda}_{\mathcal{G}_i}^{\pm} := a_i^{\pm} \Lambda_{\mathcal{G}_i}^{\pm}$ . Note that with the rank constraint we do not have the guarantee anymore that we can decompose  $\mathcal{F}$  into just two channels, so generally it is not evident how small we can choose the number of positive and negative channels,  $n_{\text{pos}}$  and  $n_{\text{neg}}$ , respectively, while still finding the optimum. Because SDPs with rank constraints are NP-hard, we have to resort to heuristics to find a solution. One common approach is due to Burer-Monteiro<sup>17,18</sup>. Suppose that we want to solve a given SDP with a rank constraint  $\text{rank}(C) \leq r$  for some positive-semidefinite  $n \times n$  matrix  $C$ . The main idea is to parametrize  $C = X^T X$  for some  $r \times n$  complex matrix  $X$  and then optimize over the matrix elements of  $X$ . The positive-semidefiniteness and rank constraint of  $C$  are automatically enforced from the construction. Unfortunately, the objective and the constraints generally contain quadratic terms of  $X$ , so the problem becomes a quadratically-constrained quadratic programme. Still, recent research has demonstrated that the objective landscape of these problems tends to behave nicely and that using local optimization methods can provably lead to the global optimum under some assumptions<sup>19,20</sup>. We defer further technical details to the Supplementary Information.

### DATA AVAILABILITY

The code to reproduce all numerical experiments and plots in this paper is available at <https://github.com/ChriPiv/stinespring-algo-paper>.

Received: 19 April 2021; Accepted: 15 December 2021;  
Published online: 03 February 2022

## REFERENCES

- Aharonov, D. & Ben-Or, M. Fault-tolerant quantum computation with constant error rate. *SIAM J. Comput.* **8**, 1207–1282 (2008).
- Preskill, J. Quantum computing in the NISQ era and beyond. *Quantum* **2**, 79 (2018).
- McClean, J. R., Kimchi-Schwartz, M. E., Carter, J. & de Jong, W. A. Hybrid quantum-classical hierarchy for mitigation of decoherence and determination of excited states. *Phys. Rev. A* **95**, 042308 (2017).
- Li, Y. & Benjamin, S. C. Efficient variational quantum simulator incorporating active error minimization. *Phys. Rev. X* **7**, 021050 (2017).
- Temme, K., Bravyi, S. & Gambetta, J. M. Error mitigation for short-depth quantum circuits. *Phys. Rev. Lett.* **119**, <https://doi.org/10.1103/physrevlett.119.180509> (2017).
- M. Otten and S. Gray. Accounting for errors in quantum algorithms via individual error reduction. *npj Quantum Inf.* **5**, <https://doi.org/10.1038/s41534-019-0125-3> (2018).
- Endo, S., Benjamin, S. C. & Li, Y. Practical quantum error mitigation for near-future applications. *Phys. Rev. X* **8**, 031027 (2018).
- Takagi, R. Optimal resource cost for error mitigation. *Phys. Rev. Research* **3**, 033178 (2021).
- Strikis, A., Qin, D., Chen, Y., Benjamin, S. C. & Li, Y. Learning-based quantum error mitigation. *PRX Quantum* **2**, 040330 (2021).
- Watrous, J. Simpler semidefinite programs for completely bounded norms. *Chicago J. Theor. Comput. Sci.* <https://doi.org/10.4086/cjcs.2013.008> (2013).
- Cai, Z. Multi-exponential error extrapolation and combining error mitigation techniques for NISQ applications. *npj Quantum Inf.* **7**, 80 (2021).
- Khatri, S. et al. Quantum-assisted quantum compiling. *Quantum* **3**, 140 (2019).
- Jones, E. et al. SciPy: open source scientific tools for Python. <http://www.scipy.org/> (2001).
- Maclaurin, D., Duvenaud, D. & Adams, R. P. Autograd: effortless gradients in NumPy. *ICML 2015 AutoML Workshop*. <https://indico.lal.in2p3.fr/event/2914/session/1/contribution/6/3/material/paper/0.pdf> (2015).
- Piveteau, C. Advanced methods for quasiprobabilistic quantum error mitigation. <https://doi.org/10.3929/ethz-b-000504508>. Master thesis, ETH Zurich (2020).
- Jiang, J., Wang, K. & Wang, X. Physical implementability of quantum maps and its application in error mitigation. *Quantum* **5**, 600 (2021).
- Burer, S. & Monteiro, R. A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Math. Program., Ser. B* **95**, 329–357 (2003).
- Burer, S. & Monteiro, R. Local minima and convergence in low-rank semidefinite programming. *Math. Program.* **103**, 427–444 (2005).
- Cifuentes, D. On the Burer–Monteiro method for general semidefinite programs. *Optim. Lett.* <https://doi.org/10.1007/s11590-021-01705-4> (2021).
- Boumal, N., Voroninski, V. & Bandeira, A. In *Advances in Neural Information Processing Systems*, Vol. 29. <https://proceedings.neurips.cc/paper/2016/file/3de2334a314a7a72721f1f74a6cb4cee-Paper.pdf> (2016).
- Abraham, H. et al. Qiskit: an open-source framework for quantum computing. <https://doi.org/10.5281/zenodo.2562110> (2019).
- MOSEK ApS. *MOSEK Optimizer API for Python 9.2.8*. <https://docs.mosek.com/9.2/pythonapi/index.html> (2020).
- Diamond, S. & Boyd, S. CVXPY: a Python-embedded modeling language for convex optimization. *J. Mach. Learn. Res.* **17**, 1–5 (2016).
- Agrawal, A., Verschuere, R., Diamond, S. & Boyd, S. A rewriting system for convex optimization problems. *J. Control Decis.* **5**, 42–60 (2018).

## ACKNOWLEDGEMENTS

We thank Sergey Bravyi, Jay Gambetta and Kristan Temme for discussions on quasiprobability decompositions. IBM, the IBM logo, and ibm.com are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. The current list of IBM trademarks is available at <https://www.ibm.com/legal/copytrade>.

## AUTHOR CONTRIBUTIONS

All authors contributed equally to this work.

## COMPETING INTERESTS

The authors declare no competing interests.

## ADDITIONAL INFORMATION

**Supplementary information** The online version contains supplementary material available at <https://doi.org/10.1038/s41534-022-00517-3>.

**Correspondence** and requests for materials should be addressed to David Sutter.

**Reprints and permission information** is available at <http://www.nature.com/reprints>

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2022