



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Quantum singular-value decomposition of nonsparse low-rank matrices

Patrick Rebentrost, Adrian Steffens, Iman Marvian, and Seth Lloyd

Phys. Rev. A **97**, 012327 — Published 24 January 2018

DOI: [10.1103/PhysRevA.97.012327](https://doi.org/10.1103/PhysRevA.97.012327)

Quantum singular value decomposition of non-sparse low-rank matrices

Patrick Rebentrost,^{1,2,*} Adrian Steffens,^{3,1} Iman Marvian,¹ and Seth Lloyd^{1,4,†}

¹Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, MA 02139

²Xanadu, 372 Richmond St W, Toronto, M5V 2L7, Canada

³Dahlem Center for Complex Quantum Systems, Freie Universität Berlin, 14195 Berlin

⁴Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139

In this work, we present a method to exponentiate non-sparse indefinite low-rank matrices on a quantum computer. Given access to the elements of the matrix, our method allows to determine the singular values and their associated singular vectors in time exponentially faster in the dimension of the matrix than known classical algorithms. The method extends to non-Hermitian and non-square matrices via matrix embedding. Moreover, our method preserves the phase relations between the singular spaces allowing for efficient algorithms that require operating on the entire singular value decomposition of a matrix. As an example of such an algorithm, we discuss the Procrustes problem of finding a closest isometry to a given matrix.

Matrix computations are central to many algorithms in optimization and machine learning [1–3]. At the heart of these algorithms regularly lies an eigenvalue or a singular value decomposition of a matrix, or a matrix inversion. Such tasks could be performed efficiently via phase estimation on a universal quantum computer [4], as long as the matrix can be simulated (exponentiated) efficiently and controllably as a Hamiltonian acting on a quantum state. Ref. [5] paved the way for such a simulation of quantum systems by introducing an efficient algorithm for exponentiating Hamiltonians with tensor product structure—enabling applications such as in quantum computing for quantum chemistry [6]. Step by step, more general types of quantum systems were tackled and performance was increased: Aharonov and Ta-Shma [7] showed a method for simulating quantum systems described by sparse Hamiltonians, while Childs *et al.* [8] demonstrated the simulation of a quantum walk on a sparse graph. Berry *et al.* [9] reduced the temporal scaling to approximately linear via higher-order Suzuki integrators and further improvements in the sparsity scaling were presented in Ref. [10]. Beyond sparse Hamiltonians, quantum principal component analysis (qPCA) was shown to be capable of efficiently handling non-sparse positive semidefinite low-rank Hamiltonians [11] given multiple copies (or samples) of the Hamiltonian as a quantum density matrix, as opposed to the requirement of quantum oracles as in the works before. This method offers applications in quantum process tomography and state discrimination [11], as well as in quantum machine learning [12–18], specifically in curve fitting [19] and support vector machines [20]. In an oracular setting, Refs. [10, 21, 22] showed the simulation of non-sparse Hamiltonians via discrete quantum walks. The scaling in terms of the simulated time t is $t^{3/2}$ or even linear in t .

In the spirit of Ref. [11], we provide an alternative method for efficiently simulating non-sparse matrices in an oracular setting that requires only *one*-sparse simulation techniques. Compared to Ref. [11], the matrices are not restricted to be positive semidefinite and the samples required are simple uniform superpositions instead of the actual Hamiltonian. We

achieve a run time in terms of the matrix maximum element and a t^2 scaling, and discuss a class of low-rank matrices which can be efficiently simulated.

In order to effectively treat a general non-Hermitian non-quadratic matrix, we make use of an indefinite “extended Hermitian matrix” that incorporates the original matrix. With such an extended matrix, we are able to efficiently determine the singular value decomposition of dense non-square, low-rank matrices. As one possible application of our method, we discuss the Procrustes problem [1] of finding a closest isometric matrix.

Method.— We have been given an $N \times N$ dense (non-sparse) Hermitian indefinite matrix $A \in \mathbb{C}^{N \times N}$ via efficient oracle access to the elements of A . For the more general case of non-Hermitian matrices, see the section “non-square matrices” below. The oracle either performs an efficient computation of the matrix elements or provides access to a storage medium for the elements such as quantum RAM [23, 24]. Our new method simulates $e^{-i(A/N)t}$ on an arbitrary quantum state for arbitrary times t . Note that the eigenvalues of A/N are bounded by $\pm \|A\|_{\max}$, where $\|A\|_{\max} = \max_{mn} |A_{mn}|$ denotes the maximal absolute value of the matrix elements of A . This means that there exist matrices A for which the unitary $e^{-i(A/N)t}$ can be far from the identity operator for a time of order $\|A\|_{\max}^{-1}$, i.e. an initial state can evolve to a perfectly distinguishable state. For such times, the unitary $e^{-i(A/N)t}$ can be well approximated by a unitary generated by a low-rank matrix.

Let σ and ρ be N -dimensional density matrices. The state σ is the target state on which the matrix exponential of A/N is applied to, while multiple copies of ρ are used as ancillary states. Our method embeds the N^2 elements of A into a Hermitian sparse matrix $S_A \in \mathbb{C}^{N^2 \times N^2}$, which we call “modified swap matrix” because of its close relation to the usual swap matrix, but with each column of S_A containing a single element of A . The modified swap matrix between the registers for a single copy of ρ and σ is defined as

$$S_A = \sum_{j,k=1}^N A_{jk} |k\rangle\langle j| \otimes |j\rangle\langle k| \in \mathbb{C}^{N^2 \times N^2}. \quad (1)$$

This matrix is one-sparse in a quadratically bigger space and reduces to the usual swap matrix for $A_{jk} = 1$ and

* patrick@xanadu.ai

† slloyd@mit.edu

$j, k = 1, \dots, N$. Given efficient oracle access to the elements, we can simulate a one-sparse matrix such as S_A with a constant number of oracle calls and negligible error [7–9, 25]. We discuss the oracle access below in section “matrix oracle and resource requirements”. The resulting matrix exponential of S_A is applied to a tensor product of a uniform superposition and an arbitrary state. Performing S_A for small Δt leads to a reduced dynamics of σ when expanded to terms of second order in Δt as

$$\text{tr}_1\{e^{-iS_A\Delta t}\rho\otimes\sigma e^{iS_A\Delta t}\} = \sigma - i\text{tr}_1\{S_A\rho\otimes\sigma\}\Delta t + i\text{tr}_1\{\rho\otimes\sigma S_A\}\Delta t + O(\Delta t^2). \quad (2)$$

Here, tr_1 denotes the partial trace over the first register containing ρ . The first $O(\Delta t)$ term can be rewritten as $\text{tr}_1\{S_A\rho\otimes\sigma\} = \sum_{j,k=1}^N A_{jk}\langle j|\rho|k\rangle|j\rangle\langle k|\sigma$. Choosing $\rho = |\bar{1}\rangle\langle\bar{1}|$, with the uniform superposition $|\bar{1}\rangle := \frac{1}{\sqrt{N}}\sum_k|k\rangle$, leads to $\text{tr}_1\{S_A\rho\otimes\sigma\} = \frac{A}{N}\sigma$. This choice for ρ contrasts with the qPCA method, where ρ is proportional to the simulated matrix [11]. Analogously, the second $O(\Delta t)$ term becomes $\text{tr}_1\{\rho\otimes\sigma S_A\} = \sigma\frac{A}{N}$. Thus for small times, evolving with the modified swap matrix S_A on the bigger system is equivalent to evolving with A/N on the σ subsystem,

$$\text{tr}_1\{e^{-iS_A\Delta t}\rho\otimes\sigma e^{iS_A\Delta t}\} = \sigma - i\frac{\Delta t}{N}[A,\sigma] + O(\Delta t^2) \approx e^{-i\frac{A}{N}\Delta t}\sigma e^{i\frac{A}{N}\Delta t}. \quad (3)$$

Let ϵ_0 be the trace norm of the error term $O(\Delta t^2)$. We can bound this error by $\epsilon_0 \leq 2\|A\|_{\max}^2\Delta t^2$ (see Appendix). Note that $\|A\|_{\max}$ coincides with the largest absolute eigenvalue of S_A . The operation in Eq. (3) can be performed multiple times in a forward Euler fashion using multiple copies of ρ . For n steps, the resulting error is $\epsilon = n\epsilon_0$. The simulated time is $t = n\Delta t$. Hence, fixing ϵ and t ,

$$n = O\left(\frac{t^2}{\epsilon}\|A\|_{\max}^2\right) \quad (4)$$

steps are required to simulate $e^{-i(A/N)t}$. The total run time of our method is nT_A : the number steps n is multiplied with the matrix oracle access time T_A (see sec. *Matrix oracle and resource requirements* below).

We now discuss for which Hermitian matrices the algorithm runs efficiently. At a simulation time t , only the eigenvalues of A/N with $|\lambda_j|/N = \Omega(1/t)$ matter. Let the number of these eigenvalues be $r \leq N$. Thus, effectively a matrix A_r/N is simulated, for which the following lower bound holds: $\text{tr}\{A_r^2/N^2\} = \sum_{j=1}^r \lambda_j^2/N^2 = \Omega(r/t^2)$. It also holds that $\text{tr}\{A_r^2/N^2\} \leq \text{tr}\{A^2/N^2\} = \|A\|_F^2/N^2 \leq \|A\|_{\max}^2$, with $\|A\|_F$ the Frobenius norm of A . Combining the upper and the lower bound, we find that the rank of the effectively simulated matrix is $r = O(\|A\|_{\max}^2 t^2)$.

For the algorithm to be efficient in terms of matrix oracle calls, we require that the number of simulation steps n is $O(\text{poly log } N)$. Let the desired error be $1/\epsilon = O(\text{poly log } N)$. Assuming $\|A\|_{\max} = \Theta(1)$, meaning a constant independent of N , we have from Eq. (4) that we can only exponentiate for a time $t = O(\text{poly log } N)$. For such

times, only the large eigenvalues of A/N with $|\lambda_j|/N = \Omega(1/\text{poly log } N)$ matter. Such eigenvalues appear if the matrix is dense enough, for example A/N has $\Theta(N)$ non-zeros of size $\Theta(1/N)$ per row. For the rank of the simulated matrix in this case, we find that $r = O(\text{poly log } N)$, thus effectively a low-rank matrix is simulated. To summarize, we expect the method to work well for low rank matrices A that are dense with relatively small matrix elements.

A large class of matrices satisfies these criteria. Draw a random unitary matrix $U \in \mathbb{C}^{N \times N}$ from the Haar measure and r suitable eigenvalues of size $|\lambda_j| = \Theta(N)$ and multiply them as $U \text{diag}_r(\lambda_j) U^\dagger$ to construct A . Here, $\text{diag}_r(\lambda_j) \in \mathbb{C}^{N \times N}$ is the diagonal matrix with the r eigenvalues on e.g. the first r diagonal entries and zero otherwise. The entries of a typical random normalized vector have absolute size $O(1/\sqrt{N})$, and the entries of the outer product of such a vector with itself have absolute matrix elements of size $O(1/N)$. Each eigenvalue of absolute size $\Theta(N)$ is multiplied with such an outer product and the r terms are summed up. Thus, a typical matrix element of A will be of size $O(\sqrt{r})$ and $\|A\|_{\max} = O(r)$.

Phase estimation. – Phase estimation provides a gateway from unitary simulation to many interesting applications. For the use in phase estimation, we extend our method such that the matrix exponentiation of A/N can be performed conditioned on additional control qubits. With our method, the eigenvalues λ_j/N of A/N can be both positive and negative. The modified swap operator S_A corresponding to a Hermitian matrix A with eigendecomposition $A = \sum_j \lambda_j |u_j\rangle\langle u_j|$ is augmented as $|1\rangle\langle 1| \otimes S_A$, which still is a one-sparse Hermitian operator. The resulting unitary $e^{-i|1\rangle\langle 1| \otimes S_A \Delta t} = |0\rangle\langle 0| \otimes \mathbb{1} + |1\rangle\langle 1| \otimes e^{-iS_A \Delta t}$ is efficiently simulatable. This operator is applied to a state $|c\rangle\langle c| \otimes \rho \otimes \sigma$, where $|c\rangle$ is an arbitrary control qubit state. Sequential application of such controlled operations allows to use phase estimation in order to prepare the state [25]

$$|\phi\rangle = \frac{1}{\sqrt{\sum_j |\beta_j|^2}} \sum_{|\lambda_j|/N \geq \epsilon} \beta_j |u_j\rangle |\lambda_j/N\rangle \quad (5)$$

from an initial state $|\psi\rangle|0\dots 0\rangle$ with $O(\lceil \log(1/\epsilon) \rceil)$ control qubits forming an eigenvalue value register. Here, $\beta_j := \langle u_j | \psi \rangle$ and ϵ is the accuracy for resolving the eigenvalues. To achieve this accuracy, phase estimation is run for a total time $t = O(1/\epsilon)$. Thus, $O(\|A\|_{\max}^2/\epsilon^3)$ queries of the oracle for A are required, which is of order $O(\text{poly log } N)$ under the low-rank assumption for A discussed above.

Matrix oracle and resource requirements. – To simulate the modified swap matrix S_A , we employ the methods developed in Refs. [8, 9]. First, we assume quantum oracle access to the original matrix A ,

$$|j k\rangle|0\dots 0\rangle \mapsto |j k\rangle|A_{jk}\rangle. \quad (6)$$

This operation can be provided by quantum random access memory (qRAM) [23, 24] using $O(N^2)$ storage space and quantum switches for accessing the data in $T_A = O(\log^2 N)$ operations. Alternatively, matrices whose elements are efficiently computable have by definition $T_A = O(\text{poly log } N)$.

The unitary operation for the simulation of the one-sparse matrix S_A with the sparse methods [8, 9] can be simply constructed from the oracle in Eq. (6) and is given by

$$|(j, k)\rangle|0 \cdots 0\rangle \mapsto |(j, k)\rangle|(k, j), (S_A)_{(k,j),(j,k)}\rangle. \quad (7)$$

Here, we use (j, k) as label for the column/row index of the modified swap matrix.

In the following, we will compare the required resources with those of other methods for sparse and non-sparse matrices: For a general $N \times N$ and s -sparse matrix, $O(sN)$ elements need to be stored. In certain cases, the sparse matrix features more structure and its elements can be computed efficiently [9, 25]. For non-sparse matrices and the qPCA method [11], only multiple copies of the density matrix as opposed to an operation as in Eq. (6) are required for applications such as state tomography. For machine learning via qPCA [11, 20], the density matrix is prepared from a classical source via quantum RAM [23, 24] and requires $O(N^2)$ storage. In comparison, the requirements of the method in this work are in principle not higher than these sparse and non-sparse methods, both in the case of qRAM access and in the case when matrix elements are computed instead of stored.

Non-square matrices.— Our method enables us to also effectively establish properties of general non-square low-rank matrices. To determine the singular value decomposition of a matrix $A = U \Sigma V^\dagger \in \mathbb{C}^{M \times N}$ with rank r , simulating the positive semidefinite matrices AA^\dagger and $A^\dagger A$ via qPCA already yields the correct singular values and vectors. However, essential information is missing, leading to ambiguities in the singular vectors that become evident when inserting diagonal matrices into the singular value decomposition of AA^\dagger that change the relative phases of the singular vectors,

$$AA^\dagger = U \Sigma^2 U^\dagger = U \Sigma D^\dagger V^\dagger V D \Sigma U^\dagger =: \hat{A} \hat{A}^\dagger, \quad (8)$$

with $D := \text{diag}(e^{-i\vartheta_j})$, ϑ_j being arbitrary phases. If $Av_j = \sigma_j u_j$ for each $j = 1, \dots, r$, then

$$\hat{A}v_j = U \Sigma D^\dagger V^\dagger v_j = \sigma_j e^{i\vartheta_j} u_j := \sigma_j \hat{u}_j, \quad (9)$$

which means there are different phase relations between left and right singular vectors in \hat{A} from those in A . Although A and \hat{A} still share the same singular values and even the same singular vectors up to phase factors, $\|A - \hat{A}\|_F$ will in general (with the exception of positive semidefinite matrices, where $U = V$) not be zero or even be small: The matrix A cannot be reproduced this way—a singular value decomposition is more than a set of singular values and normalized singular vectors. This affects all kinds of algorithms that require the appropriate phase relations between each left singular vector u_j and the according right singular vector v_j . Such applications are determining the best low-rank approximation of a matrix, signal processing algorithms discussed in Ref. [26], or determining the nearest isometric matrix, related to the unitary Procrustes problem, of a non-Hermitian matrix.

In order to overcome this issue, consider the “extended matrix”

$$\tilde{A} := \begin{bmatrix} 0 & A \\ A^\dagger & 0 \end{bmatrix}, \quad (10)$$

which was introduced for singular value computations in Ref. [27] and recently in sparse quantum matrix inversion in [25]. The non-zero eigenvalues of \tilde{A} correspond to $\{\pm\sigma_j\}$ with $\{\sigma_j\}$ being the singular values of A for $j = 1, \dots, r$. The corresponding eigenvectors are proportional to $(u_j, \pm v_j) \in \mathbb{C}^{M+N}$ (see Appendix). The left and right singular vectors of A can be extracted from the first M and last N entries, respectively. Since \tilde{A} is Hermitian, its eigenvectors can be assumed to be orthonormal: $\|(u_j, v_j)\|^2 = \|u_j\|^2 + \|v_j\|^2 = 1$, and $(u_j, v_j) \cdot (u_j, -v_j)^\dagger = \|u_j\|^2 - \|v_j\|^2 = 0$, from which follows that the norm of each of the subvectors u_j and v_j is $1/\sqrt{2}$, independent of their respective lengths M and N . The important point is that the eigenvectors of the extended matrix preserve the correct phase relations between the left and right singular vectors since $(e^{i\vartheta_j} u_j, v_j)$ is only an eigenvector of \tilde{A} for the correct phase $e^{i\vartheta_j} = 1$.

The requirements for our quantum algorithm can be satisfied also for the extended matrix. For randomly sampled left and right singular vectors, the matrix elements have maximal size of $O(\sum_{j=1}^r \sigma_j / \sqrt{MN})$, thus $\sigma_j = O(\sqrt{MN})$. In addition, an $1/(M+N)$ factor arises in the simulation of the extended matrix from the ancillary state $\rho = |\vec{1}\rangle\langle\vec{1}|$ as before, which leads to the requirement $\sigma_j = \Theta(M+N)$. These two conditions for σ_j can be satisfied if the matrix A is not too skewed, i.e. $M = \Theta(N)$. In summary, by simulating the corresponding Hermitian extended matrices, general complex matrices of low rank can be simulated efficiently, yielding the correct singular value decomposition.

Procrustes problem.— The unitary Procrustes problem is to find the unitary matrix that most accurately transforms one matrix into another. It appears in many fields, such as in shape/factor/image analysis and statistics [1]. As an application for our method in the case of non-square matrices, we will discuss the more general Procrustes problem of finding the *isometry* W that most accurately transforms a matrix B into a matrix C : Formally, minimize $\|WB - C\|_F$ among all isometries $W \in \mathbb{C}^{M \times N}$, $W^\dagger W = \mathbb{1}$, with $B \in \mathbb{C}^{N \times K}$ and $C \in \mathbb{C}^{M \times K}$, where $M > N$. The problem is equivalent to the problem of finding the nearest isometric matrix $W \in \mathbb{C}^{M \times N}$ to a matrix $A \in \mathbb{C}^{M \times N}$ by taking $A = CB^\dagger$. Since our quantum algorithm is restricted to low rank matrices, let $A = CB^\dagger$ be low-rank with rank r and singular value decomposition $A = U \Sigma V^\dagger$ with $U \in \mathbb{C}^{M \times r}$, $\Sigma \in \mathbb{R}^{r \times r}$, and $V \in \mathbb{C}^{N \times r}$. The optimal solution to the Procrustes problem is $W = U V^\dagger$ [1], setting all singular values to one, in both the low-rank and the full-rank situation. Since A is assumed to be low rank, we find a *partial* isometry with $W^\dagger W = \mathbb{P}_{\text{col}(V)}$, with $\mathbb{P}_{\text{col}(V)}$ being the projector into the subspace spanned by the columns of V . Thus, W acts as an isometry for vectors in that subspace (see Appendix).

In a quantum algorithm, intrinsically consisting of quantum operations acting on quantum states, let the task be to apply the nearest low-rank isometry to a quantum state $|\psi\rangle$. The state $|\psi\rangle$ is assumed to be in or close to the subspace spanned by the columns of V . We assume that the extended matrix for A in Eq. (10) is given in oracular form and that A is not too skewed such that $\sigma_j/(M+N) = \Theta(1)$ and $\|A\|_{\max} = \Theta(1)$. We perform phase estimation on the input state $|0, \psi\rangle|0 \dots 0\rangle$

and, analogous to Eq. (5), obtain a state proportional to

$$\sum_{\frac{\sigma_j}{M+N} \geq \epsilon} \beta_j^\pm |u_j, \pm v_j\rangle \left| \pm \frac{\sigma_j}{M+N} \right\rangle \quad (11)$$

with $\beta_j^\pm = \langle u_j, \pm v_j | 0, \psi \rangle = \pm \langle v_j | \psi \rangle / \sqrt{2}$. The sum has $2r$ terms corresponding to the eigenvalues of the extended matrix with absolute value greater than $(M+N)\epsilon$. Performing a σ_z operation on the qubit encoding the sign of the respective eigenvalue and uncomputing the eigenvalue register yields a state proportional to $\sum_j \beta_j |u_j, \pm v_j\rangle$. Projecting onto the u_j part (with success probability $1/2$) results in a state proportional to

$$\sum_{\frac{\sigma_j}{M+N} \geq \epsilon} |u_j\rangle \langle v_j | \psi \rangle \propto U V^\dagger | \psi \rangle. \quad (12)$$

This procedure leads to the preparation of the desired state for the non-square low-rank Procrustes problem with accuracy ϵ in runtime $O(\|A\|_{\max}^2 \log^2(N+M)/\epsilon^3)$. In contrast, performing the singular value decomposition of a low-rank A classically requires in general without further structural assumptions a runtime $O(N^3)$.

Conclusion.— The method presented here allows non-sparse low-rank non-positive Hermitian $N \times N$ matrices A/N to be exponentiated for a time t with accuracy ϵ in run time

$O\left(\frac{t^2}{\epsilon} \|A\|_{\max}^2 T_A\right)$, where $\|A\|_{\max}$ is the maximal absolute element of A and T_A is the data access time. If the matrix elements are accessed via quantum RAM or computed efficiently and the significant eigenvalues of A are $\Theta(N)$, our method can achieve a run time of $O(\text{poly log } N)$ for a large class of matrices. Our method allows non-Hermitian and non-square matrices to be exponentiated via extended Hermitian matrices.

We have shown how to compute the singular value decomposition of a non-Hermitian non-sparse matrix on a quantum computer directly while keeping the relative phase information. The numerous potential applications of the quantum singular value decomposition include determining the pseudoinverse of a matrix or its closest isometry exponentially faster than any known classical algorithm. In addition, the present method has been modified to a continuous variable setting as a subroutine for Gaussian process regression [28]. In addition, by using a (possibly unknown) ancillary state different from the uniform superposition, the oracular setting of the present work and the tomography setting of Ref. [11] can be combined.

ACKNOWLEDGMENTS

We acknowledge support from DARPA, NSF, and AFOSR. AS thanks the German National Academic Foundation (Studienstiftung des deutschen Volkes) and the Fritz Haber Institute of the Max Planck Society for support.

-
- [1] G. H. Golub and C. F. Van Loan, *Matrix computations*, Vol. 3 (JHU Press, 2012).
 - [2] S. Boyd and L. Vandenberghe, *Convex Optimization* (Cambridge University Press, 2004).
 - [3] K. P. Murphy, *Machine Learning: A Probabilistic Perspective* (MIT Press, 2012).
 - [4] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information* (Cambridge university press, 2010).
 - [5] S. Lloyd, *Science* **273**, 1073 (1996).
 - [6] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, *Science* **309**, 1704 (2005).
 - [7] D. Aharonov and A. Ta-Shma, in *Proceedings of the Thirty-fifth Annual ACM Symposium on Theory of Computing*, STOC '03 (ACM, New York, NY, USA, 2003) pp. 20–29.
 - [8] A. M. Childs, R. Cleve, E. Deotto, E. Farhi, S. Gutmann, and D. A. Spielman, in *Proceedings of the Thirty-fifth Annual ACM Symposium on Theory of Computing*, STOC '03 (ACM, New York, NY, USA, 2003) pp. 59–68.
 - [9] D. W. Berry, G. Ahokas, R. Cleve, and B. C. Sanders, *Comm. Math. Phys.* **270**, 359 (2007).
 - [10] D. W. Berry and A. M. Childs, *Quantum Info. Comput.* **12**, 29 (2012).
 - [11] S. Lloyd, M. Mohseni, and P. Rebentrost, *Nature Physics* **10**, 631 (2014).
 - [12] N. Wiebe, A. Kapoor, and K. M. Svore, arXiv preprint arXiv:1401.2142 (2014).
 - [13] N. Wiebe, A. Kapoor, and K. M. Svore, arXiv preprint arXiv:1412.3489 (2014).
 - [14] M. Benedetti, J. Realpe-Gómez, R. Biswas, and A. Perdomo-Ortiz, arXiv preprint arXiv:1510.07611 (2015).
 - [15] M. Schuld, I. Sinayskiy, and F. Petruccione, *Physics Letters A* **379**, 660 (2015).
 - [16] M. Schuld, I. Sinayskiy, and F. Petruccione, arXiv preprint arXiv:1601.07823 (2016).
 - [17] I. Kerenidis and A. Prakash, arXiv preprint arXiv:1603.08675 (2016).
 - [18] H.-K. Lau, R. Pooser, G. Siopsis, and C. Weedbrook, arXiv preprint arXiv:1603.06222 (2016).
 - [19] G. Wang, arXiv preprint arXiv:1402.0660 (2014).
 - [20] P. Rebentrost, M. Mohseni, and S. Lloyd, *Physical Review Letters* **113**, 130503 (2014).
 - [21] A. Childs, *Comm. Math. Phys.* **294**, 581 (2010).
 - [22] A. Childs and R. Kothari, in *Theory of Quantum Computation, Communication, and Cryptography*, Lecture Notes in Computer Science Vol. 6519, edited by W. van Dam, V. Kendon, and S. Severini (Springer, Berlin, Heidelberg, 2011), p. 94, ISBN 978-3-642-18072-9 (2011).
 - [23] V. Giovannetti, S. Lloyd, and L. Maccone, *Phys. Rev. Lett.* **100**, 160501 (2008).
 - [24] V. Giovannetti, S. Lloyd, and L. Maccone, *Phys. Rev. A* **78**, 052310 (2008).
 - [25] A. W. Harrow, A. Hassidim, and S. Lloyd, *Phys. Rev. Lett.* **103**, 150502 (2009).
 - [26] A. Steffens, P. Rebentrost, I. Marvian, J. Eisert, and S. Lloyd,

to be submitted (2016).

- [27] G. Golub and W. Kahan, *Journal of the Society for Industrial and Applied Mathematics, Series B: Numerical Analysis* **2**, 205 (1965).
 [28] S. Das, G. Siopsis, and C. Weedbrook, arXiv:1707.00360 (2017).

APPENDIX

Norms.— Denote the maximum absolute element of a matrix $A \in \mathbb{C}^{N \times N}$ with $\|A\|_{\max} = \max_{j,k} |A_{jk}|$. The Frobenius or Hilbert-Schmidt norm is given by $\|A\|_F = \sqrt{\sum_{j,k} |A_{jk}|^2}$ and its nuclear norm by $\|A\|_* = \sum_{i=1}^r \sigma_i$, where r is the rank and σ_j are the singular values.

Modified swap matrix.— The modified swap matrix is defined as

$$S_A = \sum_{j,k=1}^N A_{jk} |k\rangle\langle j| \otimes |j\rangle\langle k| \in \mathbb{C}^{N^2 \times N^2}. \quad (13)$$

Taking $A_{jk} = 1$ leads to the original swap matrix $S = \sum_{j,k=1}^N |k\rangle\langle j| \otimes |j\rangle\langle k| \in \mathbb{C}^{N^2 \times N^2}$. The N^2 eigenvalues of S_A are

$$A_{11}, A_{22}, \dots, A_{NN}, A_{12}, -A_{12}, \dots, A_{j,k>j}, -A_{j,k>j}, \dots, \quad (14)$$

where $k > j$ denotes an index k greater than j . The maximal absolute eigenvalue of S_A is thus $\max_{j,k} |A_{jk}| \equiv \|A\|_{\max}$, corresponding to the maximal absolute matrix element of A . The square of the modified swap matrix is

$$(S_A)^2 = \sum_{j,k=1}^N |A_{jk}|^2 |k\rangle\langle k| \otimes |j\rangle\langle j| \leq \|A\|_{\max}^2 \mathbb{1}. \quad (15)$$

Its eigenvalues are $|A_{jk}|^2$ and the maximal eigenvalue is $\|A\|_{\max}^2$. This already points to the result that the second order error of our method naturally scales with $\|A\|_{\max}^2$, which we will now derive.

Error analysis.— In the following, we will estimate the error from the second-order term in Δt in Eq. (2). The nuclear norm of the operator part of the second order error is

$$\begin{aligned} \epsilon_{\rho,\sigma} &= \|\text{tr}_1\{S_A \rho \otimes \sigma S_A\} - \frac{1}{2} \text{tr}_1\{(S_A)^2 \rho \otimes \sigma\} \\ &\quad - \frac{1}{2} \text{tr}_1\{\rho \otimes \sigma (S_A)^2\}\|_*. \end{aligned} \quad (16)$$

In Ref. [11], this error was equal to $\epsilon_{\rho,\sigma}^{\text{qPCA}} = \|\rho - \sigma\|_* \leq 2$, which is achieved in the present algorithm by choosing A such that $A_{jk} = 1$ for each j, k . Here, our algorithm coincides with the qPCA method for ρ chosen as the uniform superposition. For general low-rank A , we bound Eq. (16) via the triangle inequality. Taking the nuclear norm of the first term results in

$$\begin{aligned} \|\text{tr}_1\{S_A \rho \otimes \sigma S_A\}\|_* &\leq \|S_A \rho \otimes \sigma S_A\|_* \\ &\leq \|\rho \otimes \sigma\|_* \|S_A^2\|_* \leq \|A\|_{\max}^2. \end{aligned} \quad (17)$$

The second and third term can be treated similarly. We obtain $\|\text{tr}_1\{(S_A)^2 \rho \otimes \sigma\}\|_* \leq \|A\|_{\max}^2$. Combining all terms yields the bound

$$\epsilon_{\rho,\sigma} \leq 2\|A\|_{\max}^2. \quad (18)$$

Extended matrices.— We define the Hermitian extended matrix \tilde{A} of a general complex-valued, not necessarily square matrix $A \in \mathbb{C}^{M \times N}$ as

$$\tilde{A} = \begin{bmatrix} 0 & A \\ A^\dagger & 0 \end{bmatrix} \in \mathbb{C}^{(M+N) \times (M+N)}. \quad (19)$$

Using block matrix identities for the determinant, we obtain its characteristic polynomial

$$\chi_{\tilde{A}}(\lambda) = \lambda^{|M-N|} \det(\lambda \mathbb{1} + \sqrt{AA^\dagger})(\lambda \mathbb{1} - \sqrt{AA^\dagger}). \quad (20)$$

The eigenvalues of \tilde{A} are either zero or correspond to $\{\pm\sigma_j\}$, the singular values of A for $j = 1, \dots, r$ with an additional sign. Hence, if A has low rank r , then \tilde{A} has low rank $2r$. The corresponding eigenvectors are proportional to $(u_j, \pm v_j) \in \mathbb{C}^{M+N}$ since

$$\begin{bmatrix} \mp\sigma_j \mathbb{1} & A \\ A^\dagger & \mp\sigma_j \mathbb{1} \end{bmatrix} \cdot \begin{bmatrix} u_j \\ \pm v_j \end{bmatrix} = 0, \quad (21)$$

where u_j and v_j are the j th left and right singular vector of A , respectively. The important point is that the eigenvectors of the extended matrix preserve the correct phase relations between the left and right singular vectors since $(e^{i\vartheta_j} u_j, \pm v_j)$ is only an eigenvector of \tilde{A} for the correct phase $e^{i\vartheta_j} = 1$,

$$\begin{aligned} \begin{bmatrix} \mp\sigma_j \mathbb{1} & A \\ A^\dagger & \mp\sigma_j \mathbb{1} \end{bmatrix} \cdot \begin{bmatrix} e^{i\vartheta_j} u_j \\ \pm v_j \end{bmatrix} &= \begin{bmatrix} \mp\sigma_j e^{i\vartheta_j} u_j \pm A v_j \\ e^{i\vartheta_j} A^\dagger u_j - \sigma_j v_j \end{bmatrix} \\ &= (e^{i\vartheta_j} - 1) \sigma_j \begin{bmatrix} \mp u_j \\ v_j \end{bmatrix}. \end{aligned} \quad (22)$$

The right hand side is only equal to zero for the correct phase $e^{i\vartheta_j} = 1$.

Low-rank Procrustes.— Let the isometry be $W = UV^\dagger$ with $U \in \mathbb{C}^{M \times r}$ and $V \in \mathbb{C}^{N \times r}$. Assume that $M > N$, leading to orthogonal columns in the full-rank Procrustes problem ($r = N$). We find for the low-rank (partial) isometry that

$$W^\dagger W = VU^\dagger UV^\dagger = VV^\dagger = \sum_{j=1}^r \vec{v}_j \vec{v}_j^\dagger. \quad (23)$$

Pick an arbitrary vector $\vec{x} = \sum_{j=1}^r \alpha_j \vec{v}_j + \vec{x}^\perp = \vec{x}^\parallel + \vec{x}^\perp$, where \vec{x}^\perp denotes the part orthogonal to the orthonormal vectors \vec{v}_j . Then,

$$W^\dagger W \vec{x} = \sum_{j=1}^r \alpha_j \vec{v}_j = \vec{x}^\parallel. \quad (24)$$

Thus, $W^\dagger W$ acts as the identity operator in the low-rank subspace, and projects out the space perpendicular to that subspace.