NExt ApplicationS of Quantum Computing

D4.7: Readout Optimization V1.0

Document Properties

Contract Number	951821
Contractual Deadline	31-10-2024
Dissemination Level	Public
Nature	Report
Editors	Giorgio Silvi, HQS Quantum Simulations Jan Reiner, HQS Quantum Simulations
Authors	Giorgio Silvi, HQS Quantum Simulations Jan Reiner, HQS Quantum Simulations
Reviewers	Jan Reiner, HQS Quantum Simulations Arseny Kovyrshin, Data Science and Modeling, Pharmaceutical Sciences, R&D, AstraZeneca, Gothenburg, Sweden
Date	25-10-2024
Keywords	Quantum computing, measurement optimization, shadow spectroscopy, Bayesian statistics, enhanced sampling, n-representability
Status	Final
Release	1.0

This project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 951821

History of Changes

Release	Date	Author, Organisation	Description of Changes
0.1	09-10-2024	Giorgio Silvi, HQS Quantum Simulations	First version for review
0.2	14-10-2024	Giorgio Silvi and Jan Reiner, HQS Quantum Simulations	Implemented changes suggested by Jan Reiner
1.0	25-10-2024	Giorgio Silvi and Jan Reiner, HQS Quantum Simulations	Implemented changes suggested by Arseny Kovyrshin

Table of Contents

1 Executive Summary

This document gives an overview of the software delivery D4.7 “Readout Optimization (RO) V1.0”, comprised of the three readout optimization methods for quantum computing measurements that were researched and implemented by HQS Quantum Simulations within the NEASQC project. The software implementations can be found on the NEASQC GitHub [1].

The three measurement optimization methods studied include one utilizing shadow measurement in the context of spectroscopy, one based on enhanced sampling using Bayesian statistics, and one based on projecting the result to fulfill so-called n-representability constraints which may be violated in a noisy quantum computation. The later two were already reported in the earlier NEASQC delivery D4.5 [2], but are included within this document as well for the sake of completeness and reader convenience. The document is structured as follows:

In Chap. 2, we introduce the technique for utilizing the measurement of classical shadows in the context of spectroscopy, a promising use case for near future quantum computing. We also show numerical results. The method was implemented based on a recent paper [3], and implemented on our NEASQC GitHub repository [1], where the associated code can be found here:
https://github.com/NEASQC/Variationals_algorithms/tree/main/classic_shadows

Next, in Chap. 3 we discuss the method for enhanced sampling and show some key results, which was already reported in a previous NEASQC deliverable [2]. The software implementation we developed is based on what can be found in literature [4], and we will make the code available in the NEASQC GitHub [1]. Currently, it can be found here:
https://github.com/NEASQC/Variationals_algorithms/tree/main/enhanced_sampling

Finally, in Chap. 4 we present briefly the projection method based on n-representability constraints and show selected results (also already reported in the previous deliverable [2]). A detailed analysis can be found in our publication on the subject [5]. The software is available on the NEASQC GitHub[1], with the corresponding code in the following folder:
https://github.com/NEASQC/Variationals_algorithms/tree/main/n-rep_projection

2 Shadow Spectroscopy

Determining energy spectra of quantum systems is a fundamental task in quantum physics and chemistry. Methods like quantum phase estimation require significant quantum resources, making them challenging to implement on near-term quantum devices.

Shadow spectroscopy [3] is a quantum algorithm designed to estimate energy differences (gaps) in a system's Hamiltonian by analyzing the time evolution of quantum states. It requires fewer quantum resources (no ancilla qubits, very few shots) and demonstrates resilience to noise, making it well-suited for noisy intermediate-scale quantum (NISQ) devices. The method exploits the fact that observable quantities evolve over time according to the system's energy differences, which can be extracted through harmonic analysis of the measured signals. It leverages classical shadows [6] – a technique based on randomized measurements – to access a broad set of observables that encode information about the energy spectrum.

In the next two sections, we briefly state the shadow spectroscopy algorithm, and show an example of the algorithm being applied to a quantum chemistry simulation.

2.1 Procedure

The shadow spectroscopy algorithm consists of the following steps: State preparation and time evolution, the measurement of classical shadows with the estimation of expectation values, and a final post-processing step to construct a spectrum. These steps are briefly described in the following three subsections.

2.1.1 State Preparation and Time Evolution

Prepare an initial quantum state $|\psi(0)\rangle$ and simulate its time evolution under the Hamiltonian H:

$$|\psi(t)\rangle = e^{-iHt}|\psi(0)\rangle. \tag{2.1}$$

The observables O of the system will evolve periodically, containing harmonic components with frequencies corresponding to the energy differences $E_k - E_l$ between the eigenstates $|E_k\rangle$ and $|E_l\rangle$ of H, yielding the signal

$$S(t) = \sum_{k,l} c_k^* c_l \langle\psi_k|O|\psi_l\rangle e^{-i(E_l - E_k)t}, \tag{2.2}$$

2.1.2 Classical Shadows

At various time intervals during the system's evolution, perform measurements in randomly chosen single-qubit bases. This is accomplished by applying random single-qubit Clifford gates $U = \bigotimes_{j=1}^{N} U_j$ across the system, where each U_j is independently selected from the Clifford group. After applying U, measure all qubits in the computational basis, obtaining an N-bit string $\hat{b}_i \in \{0,1\}^N$ for each measurement run i. From these measurement outcomes, construct classical shadows $\hat{\rho}_i$ of the evolved state by applying the inverse of the measurement channel $\mathcal{M}$, which accounts for the randomization over measurement bases. Specifically, the classical shadow for the i-th measurement is given by

$$\hat{\rho}_i = \mathcal{M}^{-1}(U^\dagger|\hat{b}_i\rangle\langle\hat{b}_i|U), \tag{2.3}$$

Note, that we get a single bit string for each measurement. Therefore this process of reversing $\mathcal{M}$ is computationally tractable, as applying Clifford operations on a bit string, which can be represented as a product state of N individual qubits, is computationally efficient.

These classical shadows serve as efficient estimators for the quantum state and enable the estimation of expectation values of observables. For a given observable O, the estimator for its expectation value is calculated as

$$\langle O\rangle \approx \frac{1}{M}\sum_{i=1}^{M} \mathrm{Tr}\left(O\hat{\rho}_i\right), \tag{2.4}$$

where M is the total number of measurements performed.

2.1.3 Post-Processing

After estimating the expectation values of the observables using classical shadows, we organize these values into a data matrix $\mathbf{D} \in \mathbb{R}^{N_o \times N_T}$, where N_o is the number of observables and N_T is the number of time measurements. Each element D_{ij} of the matrix represents the estimated expectation value of the i-th observable at the j-th time point.

To analyze the temporal correlations and identify the dominant dynamical modes of the system, we compute the spectral cross-correlation matrix C defined as

$$C = \frac{1}{N_o}\mathbf{D}^{\mathrm{T}}\mathbf{D}, \tag{2.5}$$

where $\mathbf{D}^{\mathrm{T}}$ denotes the transpose of $\mathbf{D}$. The matrix $C \in \mathbb{R}^{N_T \times N_T}$ captures the correlations between different time measurements aggregated over all observables.

By performing an eigenvalue decomposition of the cross-correlation matrix C, we obtain eigenvalues and corresponding eigenvectors that reveal the principal components of the system's dynamics. The dominant eigenvectors correspond to the most significant harmonic components present in the measured signals, effectively highlighting the primary modes of evolution in the quantum system.

Finally, Fourier-transforming the dominant eigenvectors yields the (shadow) spectra of the system.

2.2 Application to 1,3,5-trimethylenebenzene Molecule

As an example, we provide a jupyter notebook in our repository [1] where we apply shadow spectroscopy to estimate the energy spectrum of a Trimethylenemethane (TMB) molecule, modeled by the spin Hamiltonian:

$$H_S = 3(\epsilon + J - K) + \frac{3J_S}{4} - \frac{J_S}{4}\left(\sigma_1\sigma_2 + \sigma_1\sigma_3 + \sigma_2\sigma_3\right), \tag{2.6}$$

where σ_i are Pauli operators acting on qubit i, and ϵ, J, K, and J_S are system-specific constants.

In our simulation (see Figure 1), we perform shadow measurements with two different numbers of shots per Trotter step: 10 and 1000 shots. The resulting spectra successfully identify a peak at 0.029 $(= 3/2J_S)$ Hartree for both simulations, corresponding to the energy gap of the Hamiltonian.

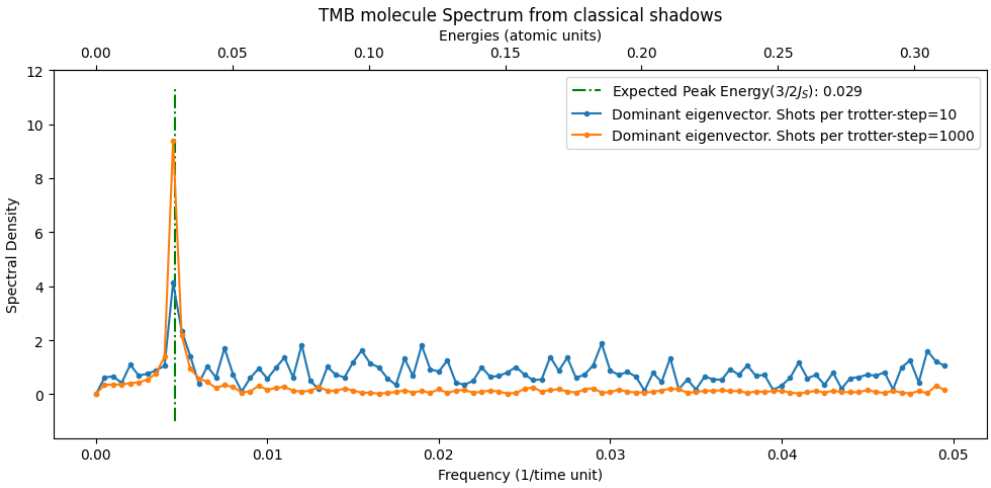

***Figure 1**: Spectral estimation of the TMB molecule using shadow spectroscopy with 10 and 1000 shots per Trotter step. A peak at 0.029 Hartree corresponds to the energy gap of interest.*

In conclusion, shadow spectroscopy provides an efficient and practical approach to estimating energy gaps in quantum systems using NISQ devices. By harnessing classical shadows and time-evolution techniques, it overcomes limitations of traditional methods, requiring fewer quantum resources and offering resilience to noise. This method opens pathways for studying complex quantum systems and advances the capabilities of early quantum hardware.

3 Enhanced Sampling

In this chapter, we give an overview of our investigated method of enhanced sampling based on Bayesian statistics. We first motivate the enhanced sampling method compared to standard sampling, continue with a brief description of the Bayesian sampling procedure, and finally show a handful of results from our numerical analysis.

3.1 Enhanced Sampling vs. Standard Sampling

Practical implementation of quantum algorithms often faces significant challenges. One of these challenges is the high demand for measurements in hybrid quantum-classical algorithms such as the Variational Quantum Eigensolver (VQE). In this section we present a new method of sampling in quantum computing that leverages Bayesian inference: enhanced sampling [4].

3.1.1 Standard Sampling

Standard sampling in quantum computing, particularly in the context of VQE, involves taking a large number of measurements to estimate the expectation value of a Hamiltonian. This process can be computationally expensive and time-consuming, especially for complex problems. The accuracy of the results depends on the number of measurements taken, with more measurements generally leading to more accurate results. However, the high demand for measurements can be a significant barrier for many practical applications.

3.2 Enhanced Sampling

The enhanced sampling method proposed in the paper ”Minimizing estimation runtime on noisy quantum computers” [4] is a technique that maximizes the statistical power of noisy quantum devices. This method is inspired by quantum-enhanced metrology, phase estimation, and the more recent “alpha-VQE” and aims to improve the efficiency of quantum amplitude estimation.

In standard sampling, as used in Variational Quantum Eigensolver (VQE), the estimation process is insensitive to small deviations in the expectation value, leading to low information gain from measurement outcomes. This results in a high runtime cost for interesting practical problems.

Enhanced sampling addresses this issue by engineering likelihood functions that increase the rate of information gain, thereby reducing the runtime of amplitude estimation. The method involves preparing an ansatz state, applying an operator, adding a phase shift about the ansatz state, and then measuring the operator. The phase shift can be achieved by performing the inverse of the ansatz circuit, then a phase shift about the initial state, and then re-applying the ansatz circuit. An example with 1 layer is shown in Fig. 2

The likelihood function in enhanced sampling is a degree-3 Chebyshev polynomial in the expectation value, referred to as a Chebyshev likelihood function (CLF). This is in contrast to standard sampling, where the likelihood function is a degree-1 polynomial. The higher degree polynomial in enhanced sampling leads to a higher rate of information gain, which reduces the runtime of the estimation process.

3.2.1 Enhanced-Circuit Choice vs ELF

In the paper ”Minimizing estimation runtime on noisy quantum computers” [4], the Engineered Likelihood Function (ELF) is used to estimate the parameters of a quantum circuit. The ELF maps the parameters of the quantum circuit to a likelihood value, which measures how well the circuit fits the data. It is constructed by defining a prior distribution over the parameters, and then using Bayes’ rule to compute the posterior distribution given the data. The likelihood function is then defined as the marginal distribution of the data given the parameters. The ELF is used to perform maximum likelihood estimation of the parameters, which involves finding the parameters that maximize the likelihood function. This is done using a classical optimization algorithm, such as gradient descent.

In our work [1], we have chosen to implement an alternative approach where, for a given ansatz and some pre-sampling, the algorithm selects between two enhanced-circuit with different number of layers, instead of using the Engineered Likelihood Function (ELF). This method involves selecting the enhanced circuit with the highest Fisher

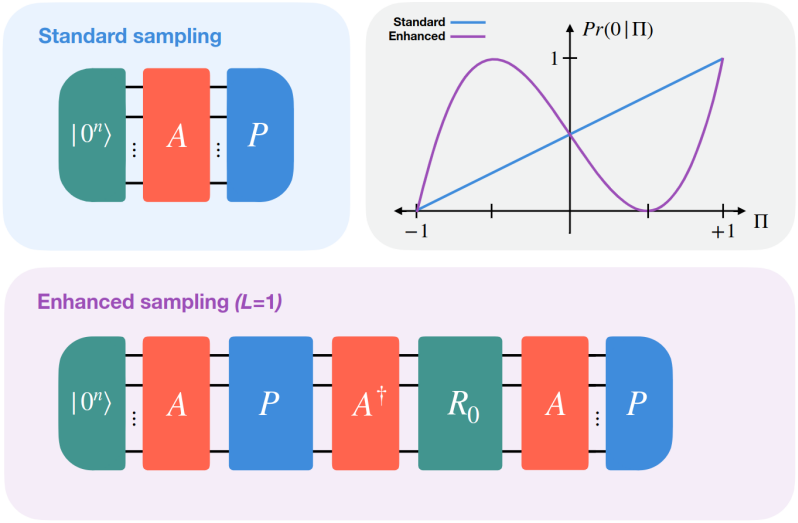

***Figure 2**: Circuit for standard sampling vs enhanced sampling. Source [1].*

information to be used for sampling. Fisher information is a measure of the amount of information that an observable random variable carries about an unknown parameter, and maximizing it can lead to more efficient sampling.

By selecting the circuit with the highest Fisher information, we aim to maximize the information gain from each measurement, thereby reducing the total number of measurements needed. This approach is more practical for our purposes than implementing the ELF, as it does not require the time-consuming adaptive scheme to be applied at every shot.

3.2.2 Results

In our repository, we offer a demonstrative notebook that showcases the capabilities of the algorithm. Specifically, given a Hamiltonian and an ansatz, we evaluate the root mean square error of the measured energy relative to the exact energy. As illustrated in Fig. 3, following an initial shared presampling phase (standard), we compare the performance of the enhanced sampling technique to that of standard sampling as additional measurements are conducted in both scenarios. In a noiseless environment, it is evident that the enhanced sampling technique outperforms standard sampling. However, this improvement is accompanied by a rise in complexity and, more significantly, an increase in circuit depth. The latter is particularly disadvantageous in noisy environments such as those experienced by Near-Term Intermediate-Scale Quantum (NISQ) devices. This drawback is evident in the case of the noisy 3-qubit system depicted in Fig. 3d. Here, the performance of enhanced sampling deteriorates to a point below that of standard sampling.

3.2.3 Conclusion

While the ELF formalism offers a promising method for enhancing the power of sampling on quantum devices, its practical implementation can be challenging due to the time required for the adaptive scheme. The two enhanced-circuit choice offers a practical alternative that can reduce the number of measurements and runtime compared to standard sampling, making it a promising approach for the implementation of quantum algorithms in practical applications.

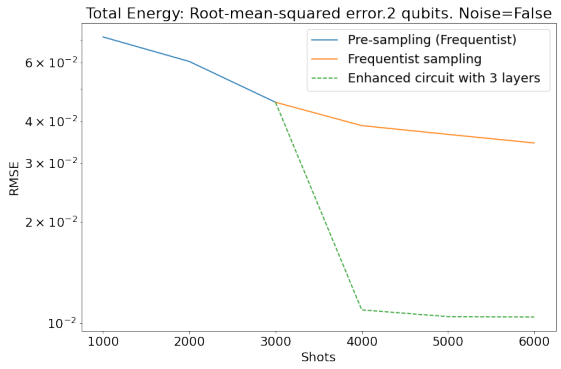

(a) 5 Runs on a 2-qubit EfficientSU2 ansatz on a noiseless simulator.

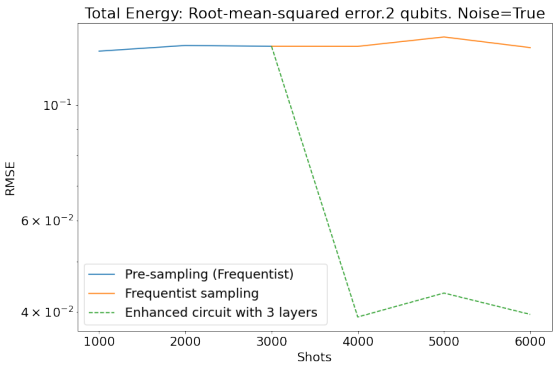

(b) 5 Runs on a 2-qubit EfficientSU2 ansatz on a noisy simulator.

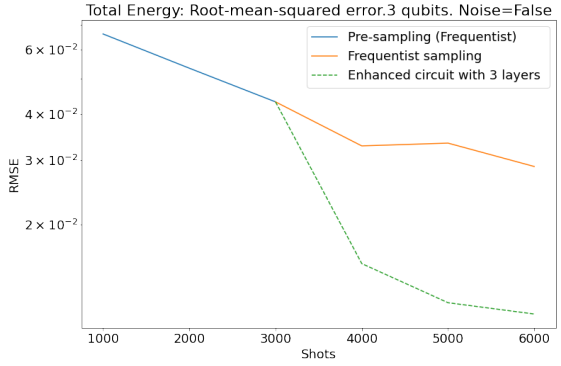

(c) 5 Runs on a 3-qubit EfficientSU2 ansatz on a noiseless simulator.

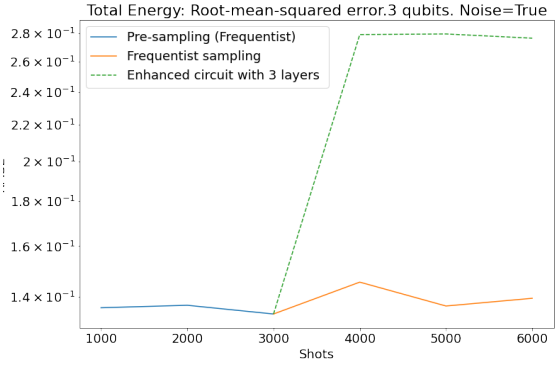

(d) 5 Runs on a 3-qubit EfficientSU2 ansatz on a noisy simulator.

***Figure 3**: Root mean square errors for various experimental setups and configurations.*

4 Post-processing noisy quantum computations utilizing N-representability constraints

In this chapter, we consider a method for improving the results of a noisy quantum computation, e.g., obtained from a NISQ device. We will describe here briefly how the method works in the following theory subsection and show some selected results in the subsequent subsection. A thorough analysis of the approach, including a detailed explanation of the method and the motivation behind it, as well extensive numerical data and a discussion of its caveats are given in our publication on this subject [5]:

Tomislav Piskor, Florian G. Eich, Michael Marthaler, Frank K. Wilhelm, and Jan-Michael Reiner, *Post-processing noisy quantum computations utilizing N-representability constraints*, arXiv:2304.13401 [quant-ph] (2023), https://arxiv.org/abs/2304.13401

4.1 Theoretical description

The internal energy is given by the expectation value of the Hamiltonian,

$$\langle H \rangle = E_0 + \sum_{ij} t_{ij} ({}^1D_{ij}) + \sum_{ijkl} V_{ijkl} ({}^2D_{ijkl}), \tag{4.1}$$

with the annihilation (creation) operators $c_i^{(\dagger)}$ of an electron in orbital i, the energy offset E_0, the one- and two-electron integrals t_{ij} and V_{ijkl}, as well as the one-particle RDM,

$${}^1D_{ij} = \langle c_i^\dagger c_j \rangle, \tag{4.2}$$

and the two-particle RDM

$${}^2D_{ijkl} = \langle c_i^\dagger c_j^\dagger c_l c_k \rangle. \tag{4.3}$$

If ${}^1D$ and ${}^2D$ of the ground state are obtained from a quantum computation on a NISQ device, they are obscured by decoherence and shot noise. We assume, that the dominating noise source is decoherence, and in this case, the calculated energy would be higher then the actual ground state energy.

One can mitigate this error and the statistical variance from shot noise by imposing constraints that the RDMs need to fulfill:

From the anti-commutation relations of the fermionic operators, on can derive that they are Hermitian and obey a certain set of anti-symmetry relations. These constraints are usually fulfilled in a quantum computation by simply calculating only a minimal necessary set of matrix elements and reconstructing the rest through the Hermiticity and anti-symmetry constraints.

Less trivial are the constraints that the RDMs need to be positive semi-definite and the trace is given be the particle number n:

$$\begin{aligned} \mathrm{tr}({}^1D) &= n, \\ \mathrm{tr}({}^2D) &= n(n-1). \end{aligned}$$

We impose these by projecting; we find the RDM that fulfills these conditions and is, in some norm, closest to the RDM calculated on the quantum computer.

Furthermore, one can transform the one- and two-particle RDMs into the hole and particle-hole sector. This means we can transform to the one- and two-hole RDMs,

$${}^1Q_{ij} = \langle c_i c_j^\dagger \rangle, \tag{4.4}$$

$${}^2Q_{ijkl} = \langle c_i c_j c_l^\dagger c_k^\dagger \rangle, \tag{4.5}$$

and the particle-hole RDM,

$$^2G_{ijkl} = \langle c_i^\dagger c_j c_l^\dagger c_k \rangle. \tag{4.6}$$

In these sectors, the same constraints hold (where the trace now depends not only on the number of electrons in the system, but also the number of holes). Instead of measuring, projecting, and calculating the energy in the particle sector, one can, after measurement, transform into a different sector, project there, transform back to the particle RDMs and calculate the energy.

We perform the projection of the measured RDMs in all three sectors, and return the energetically best of these results, as we assume the noise to be dominated by decoherence.

Again, more detailed information and reasoning can be found in our paper [5].

4.2 Selected Results

To showcase the performance of the method, we give here examples of a ground state calculation of H_2 under the influence of damping noise in Fig. 4, and under the influence of damping noise and shot noise in Fig. 5. For damping noise only we look at the energy difference of the calculation to the actual ground state energy, as well as the final state fidelity. When including shot noise, we looked at the energy difference to the ground state and the measurement variance.

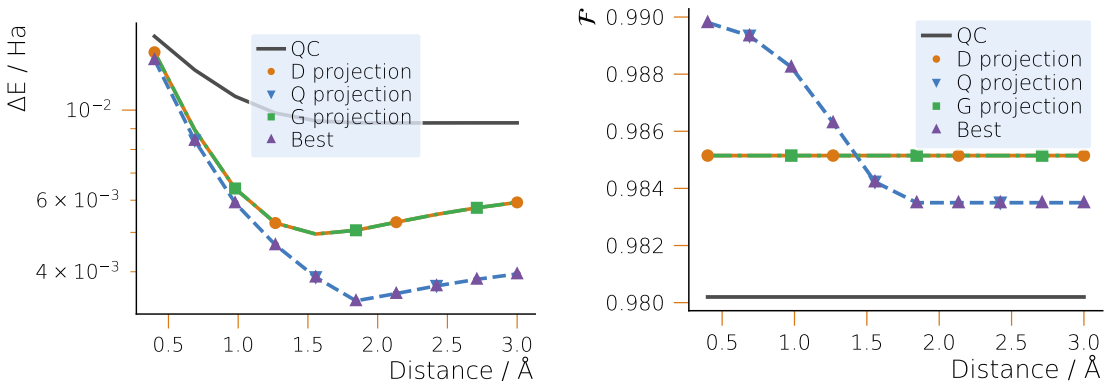

***Figure 4**: Energy difference to the ground state and fidelity of the final state towards the exact ground state for H_2; directly from the quantum calculation (QC), the three individual projection in the particle, hole, or particle-hole sector (D, Q, or G) and the energetically best result (Best).*

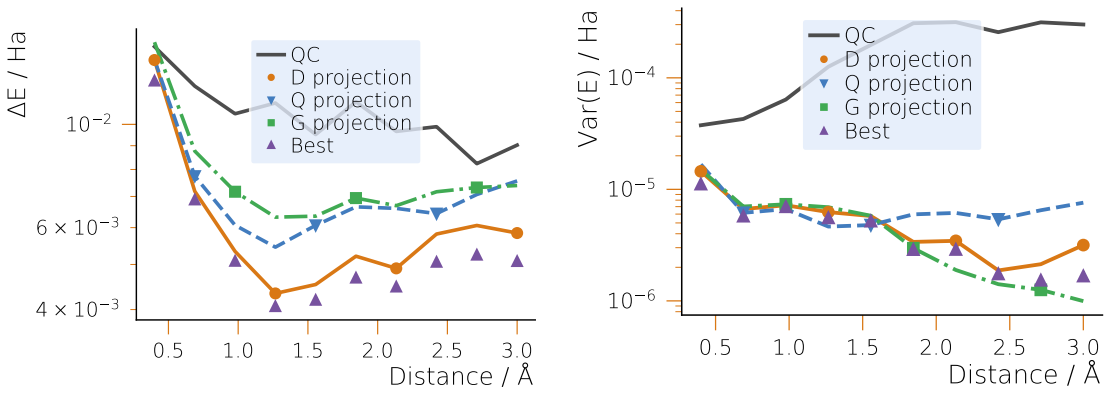

***Figure 5**: Energy difference to the ground state and measurement variance for H_2; directly from the quantum calculation (QC), the three individual projection in the particle, hole, or particle-hole sector (D, Q, or G) and the energetically best result (Best).*

In Fig. 4, note the improvement of the energy difference of almost an order of magnitude. Note also the general improvement in fidelity, but that the energetically best result does not guarantee the best fidelity.

In Fig. 5, we should mention that *Best* is strictly lower than other projections in this case, as here we average over 100 repetitions of the same measurement with different shot noise, to obtain a variance, and in each repetition different projections might lead to the energetically best result. The plot shows significant improvement of the energy difference, but also – and arguably more importantly – the reduction of the measurement variance up to two orders of magnitude. Note also that there is general reduction in the variance, but that the energetically best result does not guarantee the smallest variance.

Once more, for a detailed numerical analysis see our manuscript [5].

List of Figures

Bibliography

[1]NEASQC GitHub repository, *Variational Algorithms and Measurement Optimization*, https://github.com/NEASQC/Variationals_algorithms, 2024.

[2]NEASQC Deliverable D4.5, *RO Beta and QCCC Beta*, https://www.neasqc.eu/wp-content/uploads/2023/07/NEASQC_D4.5_RO-Beta-and-QCCC-Beta-v1.0.pdf, 2023.

[3]H. H. S. Chan, R. Meister, M. L. Goh, and B. Koczor, *Algorithmic shadow spectroscopy*, arXiv:2212.11036 (2022).

[4]G. Wang, D. E. Koh, P. D. Johnson, and Y. Cao, "Minimizing estimation runtime on noisy quantum computers", PRX Quantum **2**, 010346 (2021).

[5]T. Piskor, F. G. Eich, M. Marthaler, F. K. Wilhelm, and J.-M. Reiner, *Post-processing noisy quantum computations utilizing n-representability constraints*, arxiv:2304.13401 (2023).

[6]H.-Y. Huang, R. Kueng, and J. Preskill, "Predicting many properties of a quantum system from very few measurements", en, Nature Physics **16**, 1050–1057 (2020).