THE UNIVERSITY OF STRATHCLYDE

Classical simulations of noisy variational quantum circuits

Enrico Fontana[1,2,3], Manuel S. Rudolph[4], Ivan Rungger[3], Ross Duncan[2], Cristina Cirstoiu[2]

1 University of Strathclyde, Glasgow UK
2 Quantinuum, Cambridge UK
3 National Physical Laboratory, Teddington UK
4 EPFL, Lausanne Switzerland

arXiv:2306.05400

Motivations for VQAs

It can be argued that there are three main motivations for VQAs:

1. **Ease of implementation**

   No fault-tolerance required. Respects native gateset, topology.

2. **Usefulness**

   Can solve interesting problems. Not trivially simulatable.

3. **Noise resilience**

   Conjectured [1] and observed resilience to noise, coherent [2] and incoherent [3].

   Amenable to **quantum error mitigation (QEM)** techniques.

[1] McClean, Jarrod R., et al. "The theory of variational hybrid quantum-classical algorithms." New Journal of Physics 18.2 (2016): 023023.
[2] O'Malley, Peter JJ, et al. "Scalable quantum simulation of molecular energies." Physical Review X 6.3 (2016): 031007.
[3] Sharma, Kunal, et al. "Noise resilience of variational quantum compiling." New Journal of Physics 22.4 (2020): 043006.

Motivations for VQAs

It can be argued that there are three main motivations for VQAs:

1. **Ease of implementation**

   No fault-tolerance required. Respects native gateset, topology.

2. **Usefulness**

   Can solve interesting problems. Not trivially simulatable.

3. **Noise resilience**

   Conjectured [1] and observed resilience to noise, coherent [2] and incoherent [3].

   Amenable to **quantum error mitigation (QEM)** techniques.

[1] McClean, Jarrod R., et al. "The theory of variational hybrid quantum-classical algorithms." New Journal of Physics 18.2 (2016): 023023.
[2] O'Malley, Peter JJ, et al. "Scalable quantum simulation of molecular energies." Physical Review X 6.3 (2016): 031007.
[3] Sharma, Kunal, et al. "Noise resilience of variational quantum compiling." New Journal of Physics 22.4 (2020): 043006.

Questioning VQA noise resilience

Mixed evidence of partial resilience, at best.

- Basic numerical experiments found that there was some resilience, but only for **very low noise rates** [1].
- It also appeared that noise tends to accumulate -> **noise-induced barren plateaus** [2].
- For any QEM technique, overhead scales **worst-case exponentially** in both circuit depth and width [3].

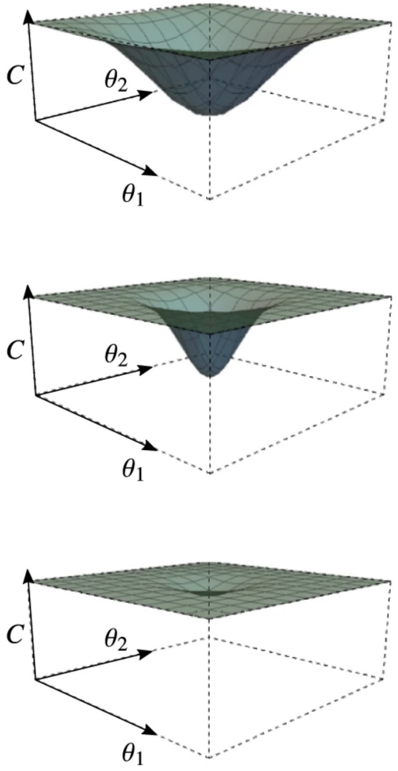

> Even with QEM, asymptotically we cannot get rid of the noise

[1] Fontana, Enrico, et al. "Evaluating the noise resilience of variational quantum algorithms." Physical Review A 104.2 (2021)
[2, image] Wang, Samson, et al. "Noise-induced barren plateaus in variational quantum algorithms." Nature communications 12.1 (2021)
[3] Takagi, Ryuji, et al. "Fundamental limits of quantum error mitigation." npj Quantum Information 8.1 (2022)

THE UNIVERSITY OF STRATHCLYDE

Questioning VQA noise resilience

Does noise prevent quantum advantage?

- Heuristically: MPS with truncation ≈ noisy QC [1].
- VQE with sufficient noise can be classically simulated via Gibbs states [2].
- This suggests that as noise accumulates, at some point a **transition** occurs and “quantumness” disappears.

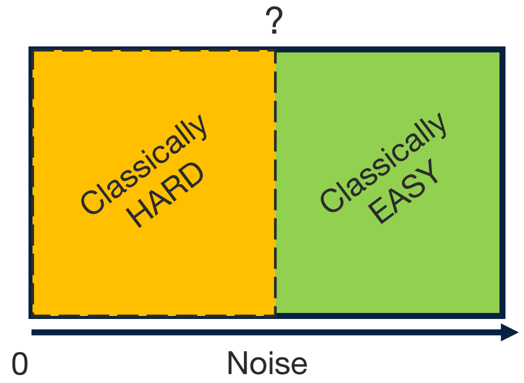

Q: In VQAs, at what point is ‘quantumness’ lost?

[1] Zhou, Yiqing, E. Miles Stoudenmire, and Xavier Waintal. "What limits the simulation of quantum computers?." Physical Review X 10.4 (2020)
[2] Stilck França, Daniel, and Raul Garcia-Patron. "Limitations of optimization algorithms on noisy quantum devices." Nature Physics 17.11 (2021)

Questioning VQA noise resilience

Does noise prevent quantum advantage?

- Heuristically: MPS with truncation ≈ noisy QC [1].
- VQE with sufficient noise can be classically simulated via Gibbs states [2].
- This suggests that as noise accumulates, at some point a **transition** occurs and "quantumness" disappears.

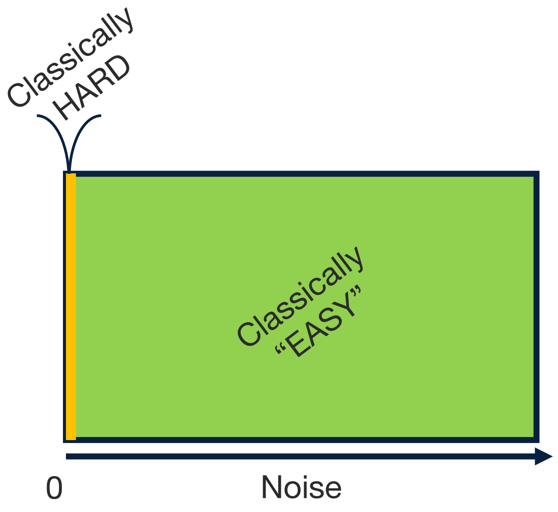

Q: In VQAs, at what point is 'quantumness' lost?

A: In some sense, immediately!

[1] Zhou, Yiqing, E. Miles Stoudenmire, and Xavier Waintal. "What limits the simulation of quantum computers?." Physical Review X 10.4 (2020)
[2] Stilck França, Daniel, and Raul Garcia-Patron. "Limitations of optimization algorithms on noisy quantum devices." Nature Physics 17.11 (2021)

THE UNIVERSITY OF STRATHCLYDE

Questioning VQA noise resilience

Does noise prevent quantum advantage?

- Heuristically: MPS with truncation ≈ noisy QC [1].
- VQE with sufficient noise can be classically simulated via Gibbs states [2].
- This suggests that as noise accumulates, at some point a **transition** occurs and “quantumness” disappears.

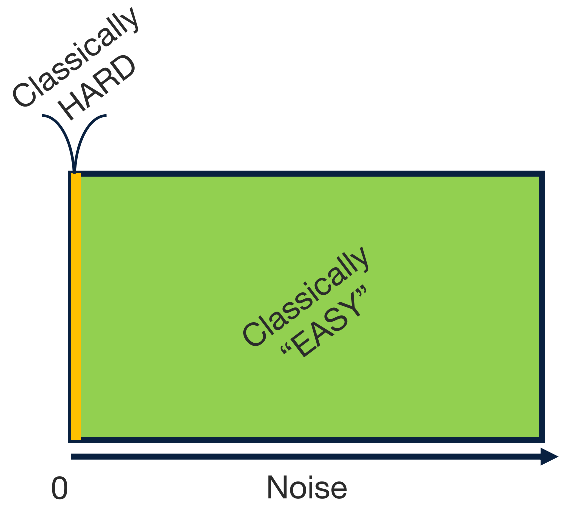

Q: In VQAs, at what point is ‘quantumness’ lost?

A: In some sense, immediately!

Fourier

[1] Zhou, Yiqing, E. Miles Stoudenmire, and Xavier Waintal. "What limits the simulation of quantum computers?." Physical Review X 10.4 (2020)
[2] Stilck França, Daniel, and Raul Garcia-Patron. "Limitations of optimization algorithms on noisy quantum devices." Nature Physics 17.11 (2021)

THE UNIVERSITY OF STRATHCLYDE

VQA cost functions are Fourier series

It is well-known [1,2,3,4] that VQA cost functions admit a **Fourier series representation**:

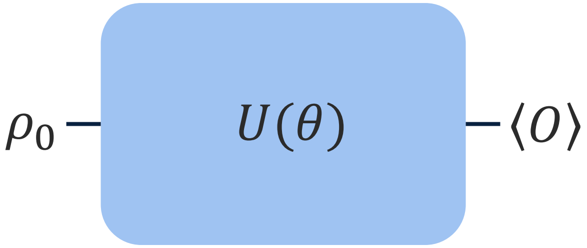

$$\Longrightarrow \quad Tr\{\rho(\boldsymbol{\theta})O\} = \sum_{\boldsymbol{\omega}\in\Omega} a(\boldsymbol{\omega})e^{i\boldsymbol{\omega}\cdot\boldsymbol{\theta}} = \sum_{\omega_0,\omega_1,\dots}$$

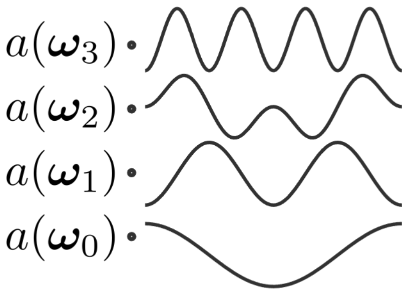

- The frequencies ω live in a discrete, finite (but large) spectrum Ω.
- Spectrum is set of pairwise differences of eigenvalues of rotations:

$$R_z(\theta) = e^{-i\sigma_z\theta/2} \rightarrow \lambda(\sigma_z/2) = \{\pm 1/2\} \rightarrow \Omega = \{-1, 0, +1\}$$

[1] Schuld, Maria, Ryan Sweke, and Johannes Jakob Meyer. "Effect of data encoding on the expressive power of variational quantum-machine-learning models." Physical Review A 103.3 (2021): 032430.
[2] Gil Vidal, Francisco Javier, and Dirk Oliver Theis. "Input redundancy for parameterized quantum circuits." Frontiers in Physics 8 (2020): 297.
[3] Nakanishi, Ken M., Keisuke Fujii, and Synge Todo. "Sequential minimal optimization for quantum-classical hybrid algorithms." Physical Review Research 2.4 (2020): 043158.
[4] Parrish, Robert M., et al. "A Jacobi diagonalization and Anderson acceleration algorithm for variational quantum algorithm parameter optimization." arXiv preprint arXiv:1904.03206 (2019).

VQA cost functions are Fourier series

It is well-known [1,2,3,4] that VQA cost functions admit a **Fourier series representation**:

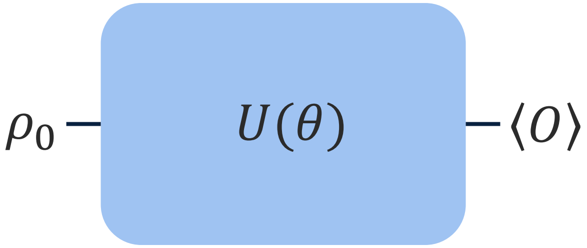

$$\Longrightarrow \quad Tr\{\rho(\boldsymbol{\theta})O\} = \sum_{\boldsymbol{\omega}\in\Omega} a(\boldsymbol{\omega})e^{i\boldsymbol{\omega}\cdot\boldsymbol{\theta}} = \sum_{\omega_0,\omega_1,\dots}$$

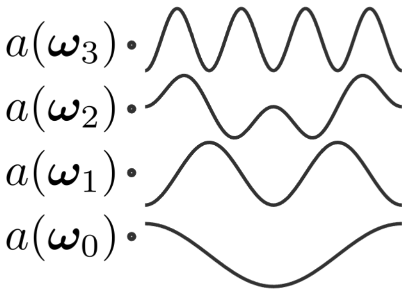

- The frequencies ω live in a discrete, finite (but large) spectrum Ω.
- Spectrum is set of pairwise differences of eigenvalues of rotations:

$$R_z(\theta) = e^{-i\sigma_z\theta/2} \rightarrow \lambda(\sigma_z/2) = \{\pm 1/2\} \rightarrow \Omega = \{-1, 0, +1\}$$

Q: What happens when noise is present?

[1] Schuld, Maria, Ryan Sweke, and Johannes Jakob Meyer. "Effect of data encoding on the expressive power of variational quantum-machine-learning models." Physical Review A 103.3 (2021): 032430.
[2] Gil Vidal, Francisco Javier, and Dirk Oliver Theis. "Input redundancy for parameterized quantum circuits." Frontiers in Physics 8 (2020): 297.
[3] Nakanishi, Ken M., Keisuke Fujii, and Synge Todo. "Sequential minimal optimization for quantum-classical hybrid algorithms." Physical Review Research 2.4 (2020): 043158.
[4] Parrish, Robert M., et al. "A Jacobi diagonalization and Anderson acceleration algorithm for variational quantum algorithm parameter optimization." arXiv preprint arXiv:1904.03206 (2019).

Channel mode decomposition

- To include noise, move to channels.

$$\mathbf{Rz}(\theta) = \mathbf{C}_0 + e^{i\theta}\mathbf{C}_1 + e^{-i\theta}\mathbf{C}_{-1}$$

- Now go from Fourier to **trigonometric decomposition** for single-Pauli rotation gate channel:

$$\mathbf{Rz}(\theta) = \mathbf{D}_0 + \cos(\theta)\mathbf{D}_1 + \sin(\theta)\mathbf{D}_{-1}$$

- With noise, the channel modes $\{\mathbf{D}_0, \mathbf{D}_{\pm 1}\}$ will depend on the error parameters.
- Single qubit Pauli noise following rotation give PTM:

$$\mathbf{D}_0 = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & q_Z \end{pmatrix}, \quad \mathbf{D}_1 = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & q_X & 0 & 0 \\ 0 & 0 & q_Y & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad \mathbf{D}_{-1} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & -q_X & 0 \\ 0 & q_Y & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

Decomposition gives PTMs that map 1 Pauli into ≤1 Paulis

Noisy Heisenberg-picture simulation

- Evolving a Pauli measurement in the Heisenberg picture, we obtain a trigonometric decomposition of the **cost function** that includes noise factors.

$$U(\boldsymbol{\theta}) = \prod_{i=1}^{m} (C_i \cdot Rz(\theta_i)) \cdot C_0$$

$$\mathbf{Rz}(\theta) = \mathbf{D}_0 + \cos(\theta)\mathbf{D}_1 + \sin(\theta)\mathbf{D}_{-1}$$

- The cost function (expectation value of Pauli operator) is

$$f(\boldsymbol{\theta}) = \sum_{\boldsymbol{\omega}} \Phi_{\boldsymbol{\omega}}(\boldsymbol{\theta}) \langle\langle 0 | \mathbf{U}_{\boldsymbol{\omega}}^{\dagger} | O \rangle\rangle$$

Contains sin/cos and noise

Constant (±1 or 0)
Classically simulatable

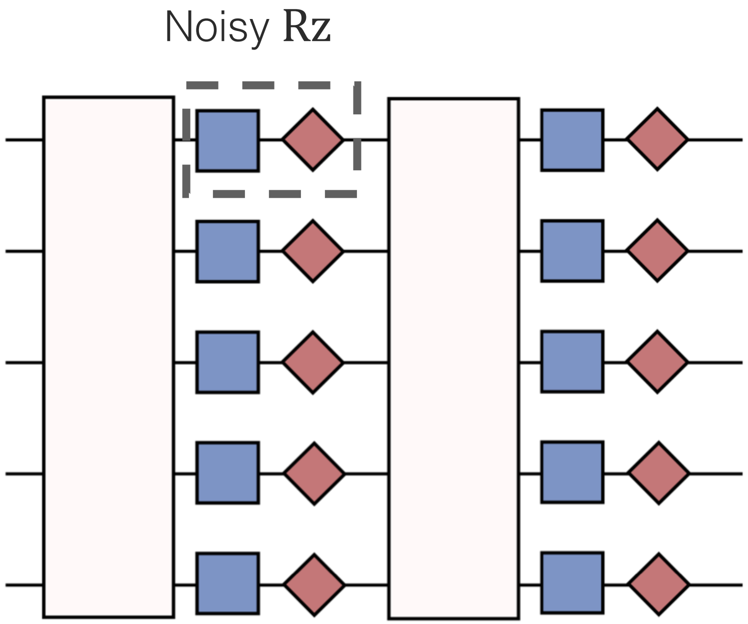

Low-weight efficient simulation algorithm (LOWESA)

Weight of a path w
= number of sin/cos

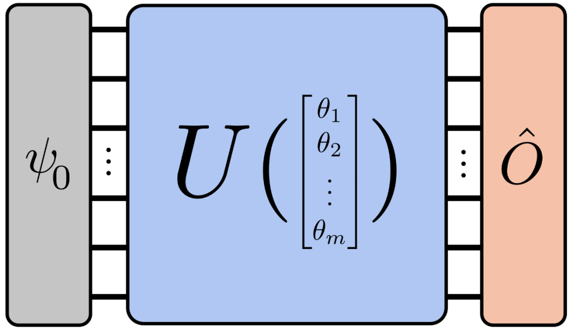

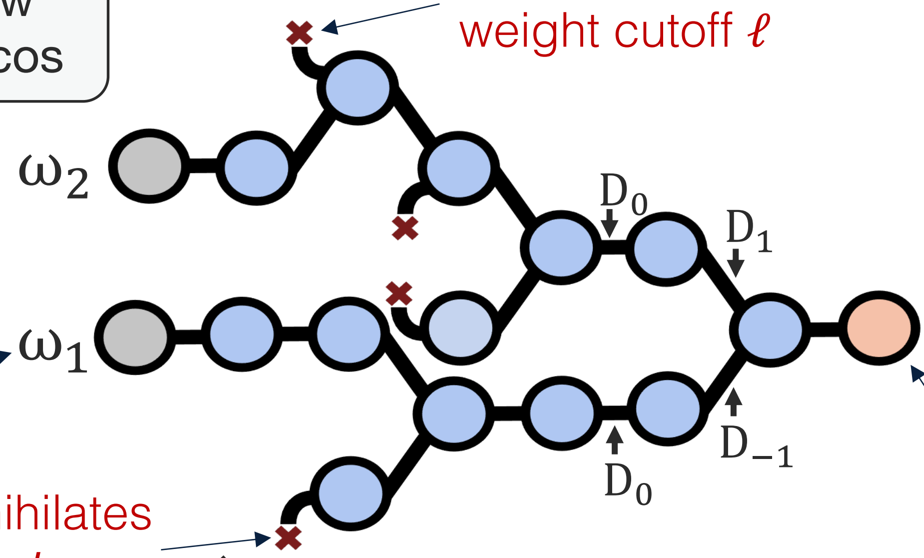

Trigonometric polynomial

$$f(\boldsymbol{\theta}) \approx \sum_{\omega_0, \omega_1, \dots} a(\omega_3) \cdot \; a(\omega_2) \cdot \; a(\omega_1) \cdot \; a(\omega_0) \cdot$$

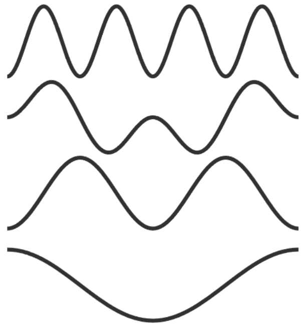

Each path contributes one term (trigonometric monomial) to the sum

THE UNIVERSITY OF STRATHCLYDE

Why this algorithm?

- Similar algorithms had been developed before [1, 2].
- Same algorithm developed in parallel by Nemkov *et al.* [3] and Begusic *et al.* [4] (different truncation scheme).
- However, these versions had **no noise**.

What is the effect of noise on the cost function?

$$f(\boldsymbol{\theta}) = \sum_{\omega} d_\omega \Phi_\omega(\boldsymbol{\theta}) \xrightarrow{\text{Dephasing}} \tilde{f}(\boldsymbol{\theta}) = \sum_{\omega} (1-p)^{|\omega|} d_\omega \Phi_\omega(\boldsymbol{\theta})$$

With noise, paths with large weights are **suppressed exponentially**

[1] Cerf, N. J., and S. E. Koonin. "Monte Carlo simulation of quantum computation." Mathematics and Computers in Simulation 47.2-5 (1998): 143-152.
[2] Gao, Xun, and Luming Duan. "Efficient classical simulation of noisy quantum computation." 1810.03176
[3] Nemkov, Nikita, et al. "Fourier expansion in variational quantum algorithms." arXiv:2304.03787
[4] Begušić, Tomislav, Kasra Hejazi, and Garnet Kin Chan. "Simulating quantum circuit expectation values by Clifford perturbation theory." arXiv:2306.04797

Theoretical performance

- Truncation reduces the paths from exponential in the number of rotations (2^m for m rotations) to constant (2^ℓ for cutoff ℓ).
- The overall runtime is now **polynomial** in the number of qubits and gates: $\mathcal{O}(n^2 m 2^\ell)$.
- Since discarded paths are the ones most affected by noise, the **average error** over landscape is **bounded**:

$$\Delta(\tilde{f}, \tilde{g}) \leq (1 - 2p - 2p_Z)^{\ell+1} \leq e^{-2(p+p_Z)\ell}$$

- Given any nonzero, constant amount of noise, a VQA is approximately **simulatable** in **polynomial** time with guaranteed average error.

Theoretical performance

- Truncation reduces the paths from exponential in the number of rotations (2^m for m rotations) to constant (2^ℓ for cutoff ℓ).
- The overall runtime is now **polynomial** in the number of qubits and gates: $\mathcal{O}(n^2 m 2^\ell)$.
- Since discarded paths are the ones most affected by noise, the **average error** over landscape is **bounded**:

$$\Delta(\tilde{f}, \tilde{g}) \leq (1 - 2p - 2p_Z)^{\ell+1} \leq e^{-2(p+p_Z)\ell}$$

- Given any nonzero, constant amount of noise, a VQA is approximately **simulatable** in **polynomial** time with guaranteed average error.

Within a small average error, there is **no exponential separation** between classical and quantum with noise

Experiments

- Simulated large random noisy circuits (up to 50 qubits) in 2D topology, in under ~10s on laptop.

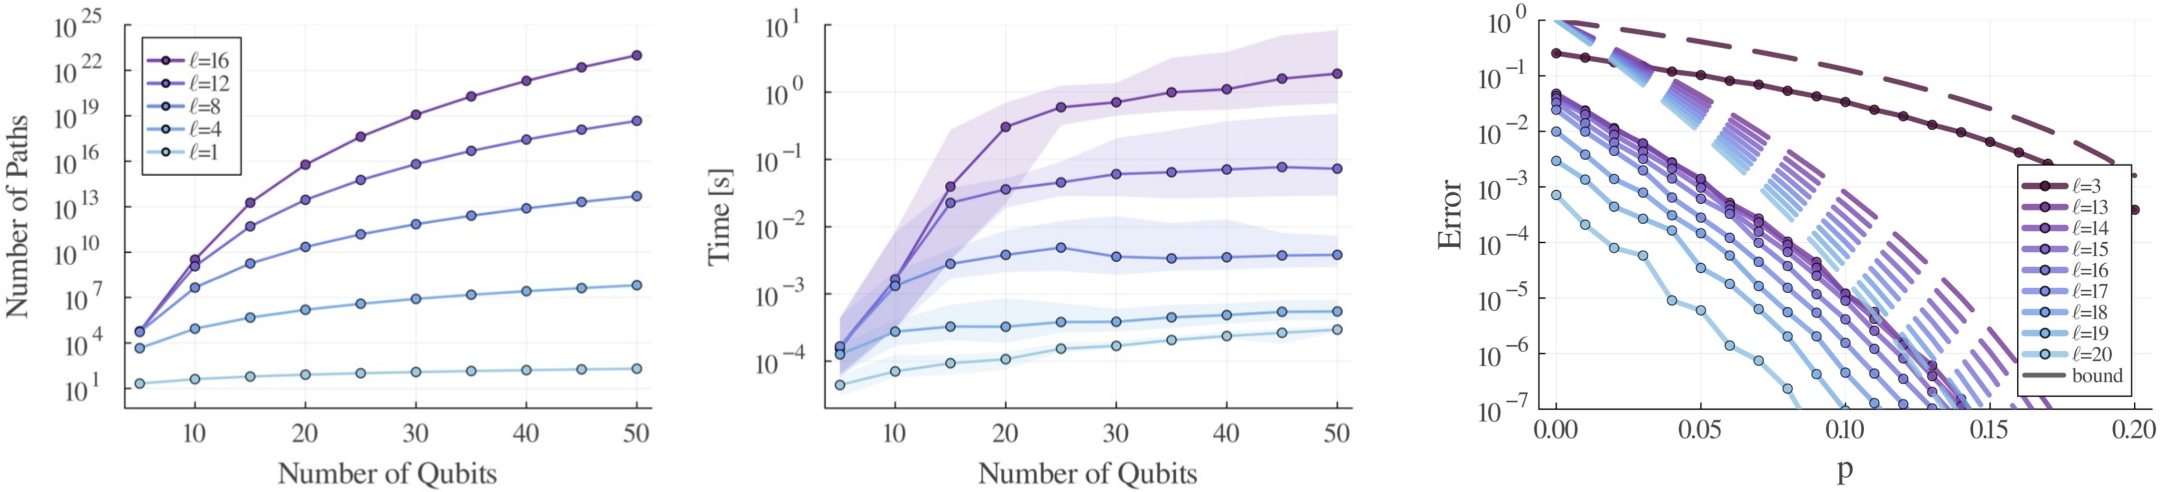

- With high ℓ, worst-case number of paths is **huge**.
- However, vast majority **annihilate early** and so simulation is still fast.
- Bounds on simulation error are loose by 1-3 orders of magnitude.

Caveats

- In theory, there is a strong dependency of optimal cutoff on the noise rate: $\ell \approx p^{-1} \log \epsilon^{-1}$

  With 1% phase noise, 1% error tolerance, worst-case runtime has a prefactor of 10^{100} = 1Google

  -> Large **gap** between empirical and worst-case performance.
- Bound holds for **uncorrelated**, single qubit rotation gates only.
- Expectation values only.

Open questions

- Can we improve the bounds?
- Can we extend to correlated parameter VQAs?
- Can we do sampling?

Outlook

- Presented LOWESA, a Fourier-based simulation algorithm that can handle noise.
- Under a reasonable noise model, uncorrelated parameter VQAs are classically simulatable in polynomial time.

Noise is a **fundamental barrier** to **exponential** quantum advantage for VQAs.

Is it a practical algorithm, though?

Outlook

- Presented LOWESA, a Fourier-based simulation algorithm that can handle noise.
- Under a reasonable noise model, uncorrelated parameter VQAs are classically simulatable in polynomial time.

> Noise is a **fundamental barrier** to **exponential** quantum advantage for VQAs.

Is it a practical algorithm, though?

- If hardware noise is large, LOWESA may be more accurate and efficient than a QC.
- There's indication that the bounds are potentially very loose for certain circuits.
- Appears that LOWESA has a role as a standalone (noiseless) simulation algorithm [1].

[1] Rudolph, Manuel S., et al. "Classical surrogate simulation of quantum systems with LOWESA." arXiv:2308.09109

THE UNIVERSITY OF STRATHCLYDE

Thank you!

University of Strathclyde Glasgow

NPL National Physical Laboratory

QUANTINUUM

Bonus: Simulating 127 **clean** qubits

Transverse Field Ising Hamiltonian

$$H = \sum_{\langle i,j \rangle} J^{(i,j)} Z_i Z_j + \sum_i h^{(i)} X i$$

Trotter Time Evolution Circuit

$$U(\boldsymbol{\theta}) = \prod_{l=1}^{L} \left[\prod_{\langle i,j \rangle} \left(R_{ZZ}(\theta_J^{(i,j)}) \right) \prod_i \left(R_X(\theta_h^{(i)}) \right) \right]$$

L = 6

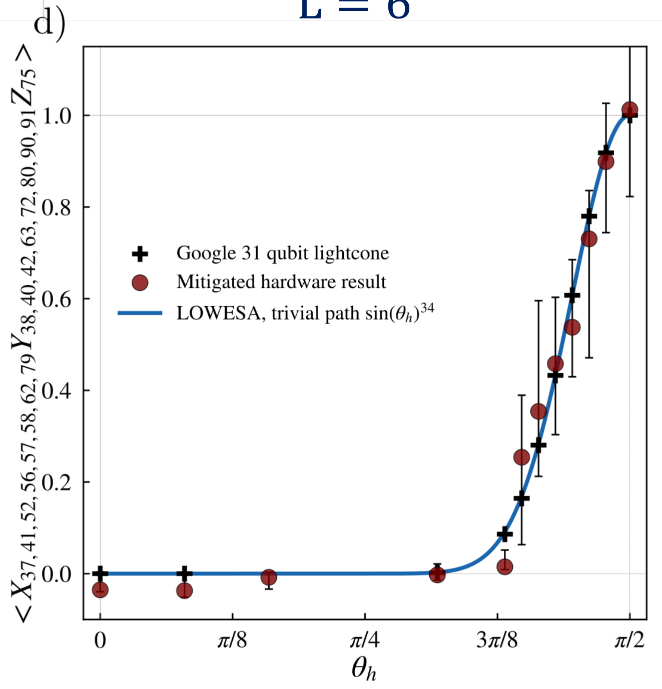

L = 20

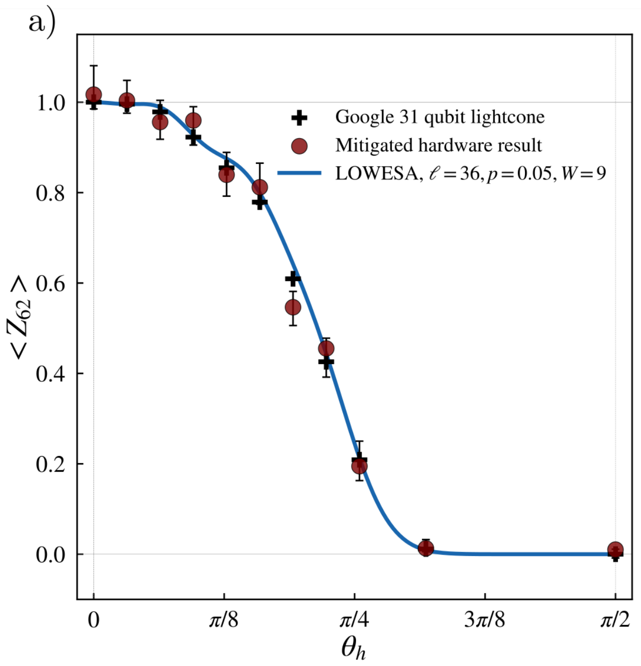

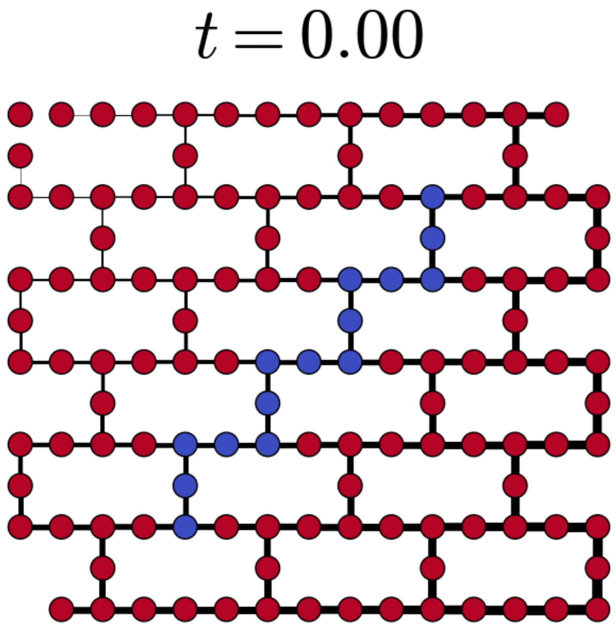

Rudolph, Manuel S., et al. "Classical surrogate simulation of quantum systems with LOWESA." arXiv:2308.09109