Noise vs quantum algorithms

Yihui Quek
MIT → EPFL
quekpottheories

Noise is the defining characteristic of NISQ computation!

Example: Google's fidelity for their 2019 quantum advantage demonstration was just 0.002.

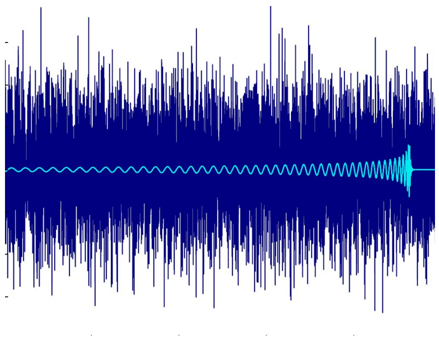

Can we compute with noisy devices?
= can we extract any signal from quantum noise?

Noise is the defining characteristic of NISQ computation!

Example: Google's fidelity for their 2019 quantum advantage demonstration was just 0.002.

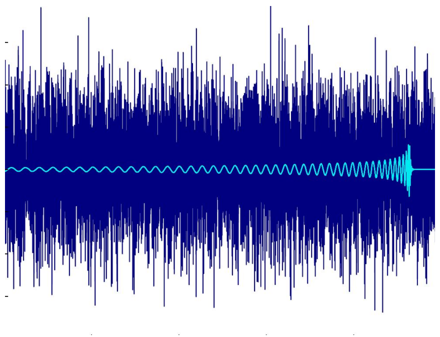

Can we compute with noisy devices?
= can we extract any signal from quantum noise?

This talk is about **going beyond conventional wisdom** regarding this question.

Conventional wisdom

Can quantum computation be made robust against errors?

~1999 [Shor, Aharonov, Ben-Or ++]: **YES, in theory** – can do quantum error correction if noise is low enough!

Conventional wisdom

Can quantum computation be made robust against errors?

~1999 [Shor, Aharonov, Ben-Or ++]: **YES, in theory** – can do quantum error correction if noise is low enough!

~2015 – **NISQ** era: **MAYBE** – We don't have enough qubits for quantum error correction and we also can't make measurements in the middle of a computation.

Conventional wisdom

Can quantum computation be made robust against errors?

~1999 [Shor, Aharonov, Ben-Or ++]: **YES, in theory** – can do quantum error correction if noise is low enough!

~2015 – **NISQ** era: **MAYBE** – We don't have enough qubits for quantum error correction and we also can't make measurements in the middle of a computation.

~2020 [Some people including me]: **In that case, NO** (see next slide)

Conventional wisdom

Can quantum computation be made robust against errors?

~1999 [Shor, Aharonov, Ben-Or ++]: **YES, in theory** – can do quantum error correction if noise is low enough!

~2015 – **NISQ** era: **MAYBE** – We don't have enough qubits for quantum error correction and we also can't make measurements in the middle of a computation.

~2020 [Some people including me]: **In that case, NO** (see next slide)

~ 2024/2025 [This talk]: **But what about NISQ+?**

Depolarizing noise is certifiably bad news
(if you cannot error correct)

nature physics

Explore content ⌄ About the journal ⌄ Publish with us ⌄

nature › nature physics › articles › article

Article | Published: 21 October 2021

Limitations of optimization algorithms on noisy quantum devices

Daniel Stilck França & Raul García-Pa

nature communications

Explore content ⌄ About the journal ⌄ Publish with us ⌄

nature › nature communications › articles › article

Article | Open access | Published: 29 November 2021

Noise-induced barren plateaus in variational quantum algorithms

Samson Wang, Enrico Fontana, M. Cerezo, Kunal Sharma, Akira Sone, Lukasz Cincio & Patrick J. Coles

Depolarizing noise is certifiably bad news

(if you cannot error correct)

nature physics

Explore content ⌄ About the journal ⌄ Publish with us ⌄

nature › nature physics › articles › article

Article | Published: 21 October 2021

Limitations of optimization algorithms on noisy quantum devices

Daniel Stilck França & Raul García-Pa

nature communications

Explore content ⌄ About the journal ⌄ Publish with us ⌄

nature › nature communications › articles › article

Article | Open access | Published: 29 November 2021

Noise-induced barren plateaus in variational quantum algorithms

Samson Wang, Enrico Fontana, M. Cerezo, Kunal Sharma, Akira Sone, Lukasz Cincio & Patrick J. Coles

nature physics

Explore content ⌄ About the journal ⌄ Publish with us ⌄

nature › nature physics › articles › article

Article | Open access | Published: 25 July 2024

Exponentially tighter bounds on limitations of quantum error mitigation

Yihui Quek, Daniel Stilck França, Sumeet Khatri, Johannes Jakob Meyer & Jens Eisert

Depolarizing noise is certifiably bad news
(if you cannot error correct)

nature physics

Explore content ⌄ About the journal ⌄ Publish with us ⌄

nature › nature physics › articles › article

Article | Published: 21 October 2021

Limitations of optimization algorithms
devices

Daniel Stilck França

nature physics

Explore content ⌄ About the journal ⌄ Publish with us ⌄

nature › nature physics › articles › article

Article | Open access | Published: 25 July 2024

hter bounds on
ntum error

, Sumeet Khatri, Johannes Jakob

And all because of this
one weird fact!

Article | Open access | Published: 29 November 2021

Noise-induced barren plateaus in variational quantum algorithms

Samson Wang, Enrico Fontana, M. Cerezo, Kunal Sharma, Akira Sone, Lukasz Cincio & Patrick J. Coles

Depolarizing noise increases entropy!

Depolarizing noise increases entropy!

Depolarizing noise
acts on *each* qubit
as: $\mathcal{D}(\rho) = \ldots \mathbb{I}$

Maximally-mixed state $\mathbb{I}$ ~ uniform distribution on a single bit.
Contains no information; maximally entropic!

Depolarizing noise increases entropy rapidly!

Depolarizing noise acts on *each* qubit as:

● = single-qubit noise

Maximally-mixed state ~ uniform distribution on a single bit.
Contains no information, maximally entropic!

'Gate-based' noise in a circuit: noise acts after every gate!

ρ_{in}
(n qubits)

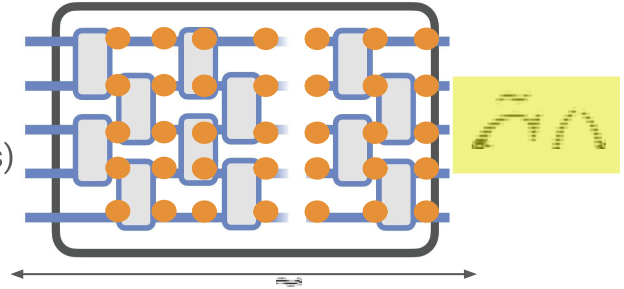

depth( )

What state does this circuit output?

As circuit depth increases:

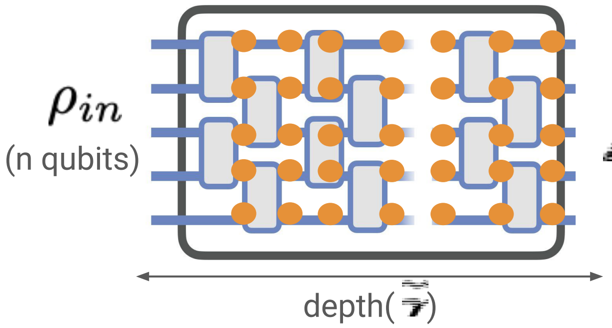

Noisy circuit's output state

A state that contains no information and can't be used for computation

As circuit depth increases, depolarizing noisy circuit eventually outputs maximally-mixed state

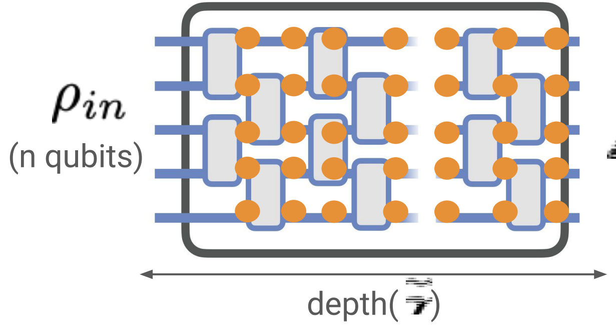

Noisy circuit's output state

A state that contains no information and can't be used for computation

How fast do they converge?

Trace norm – measures how 'far apart' states are.

As circuit depth increases, depolarizing noisy circuit eventually outputs maximally-mixed state

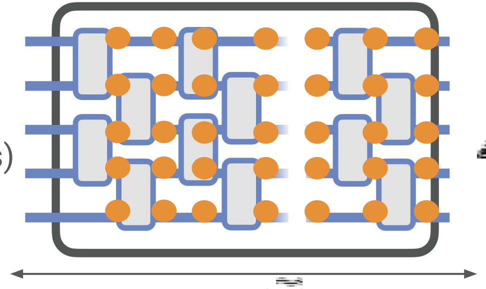

Noisy circuit's output state

A state that contains no information and can't be used for computation

How fast do they converge?

$$\left\| \tilde{\mathcal{C}}(\rho_{in}) - \frac{\mathbb{I}}{2^n} \right\| \leq c^{\text{depth}(\tilde{\mathcal{C}})}$$ for any ρ_{in}, any $\vec{\gamma}$, c<1

Convergence in log depth for *all circuits* :(

As circuit depth increases, depolarizing noisy circuit eventually outputs maximally-mixed state

ρ_{in}
(n qubits)

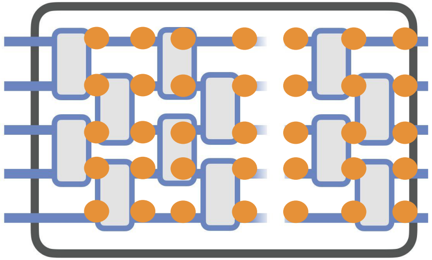

depth($\tilde{C}$)

$\tilde{\Phi}(\rho_{in})$

$\tilde{\Phi}(\rho_{in})$
Noisy circuit's output state

$I/2^n$

A state that contains no information and can't be used for computation

How fast do they converge?

$$\left\| \tilde{\Phi}(\rho_{in}) - \frac{I}{2^n} \right\| \leq c^{\mathrm{depth}(\tilde{C})}$$ for any ρ_{in}, any $\tilde{C}$, c<1

Convergence in log depth for *all circuits* :(

Huge problem for quantum advantage!

What is error mitigation?

In a world with noiseless quantum computers:

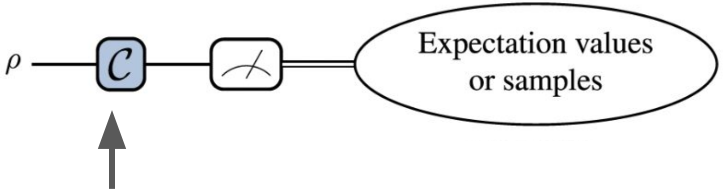

Quantum circuit
running
algorithm of
interest

What is error mitigation?

In the real world, C is noisy:

- qubit decoherence
- gate errors

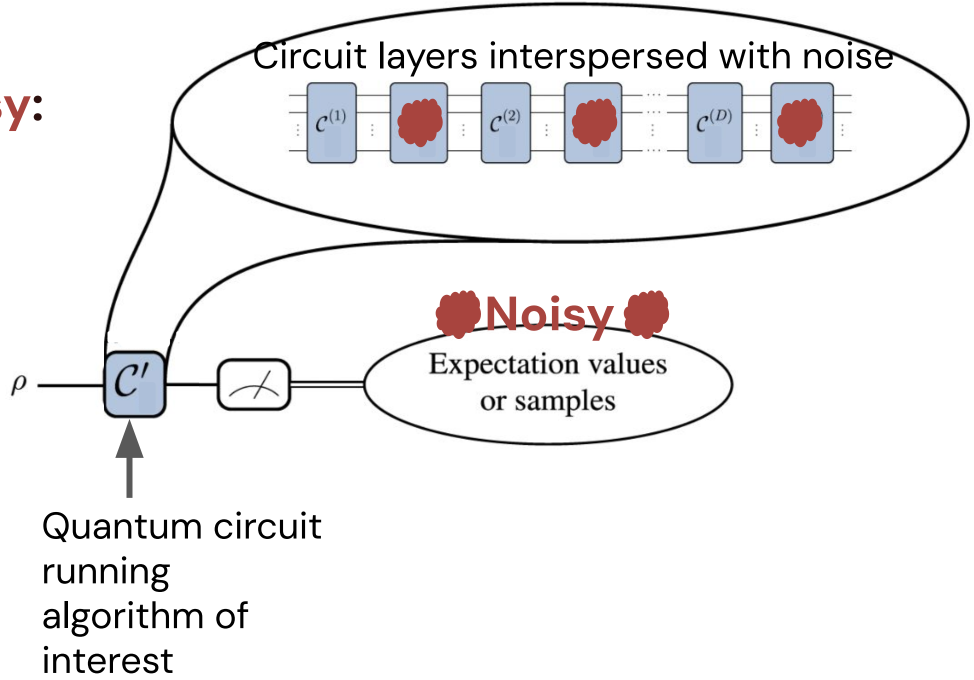

What is error mitigation?

In the real world, C is noisy:

Proposal: revert the effect of noise on the computation result, with classical post-processing.

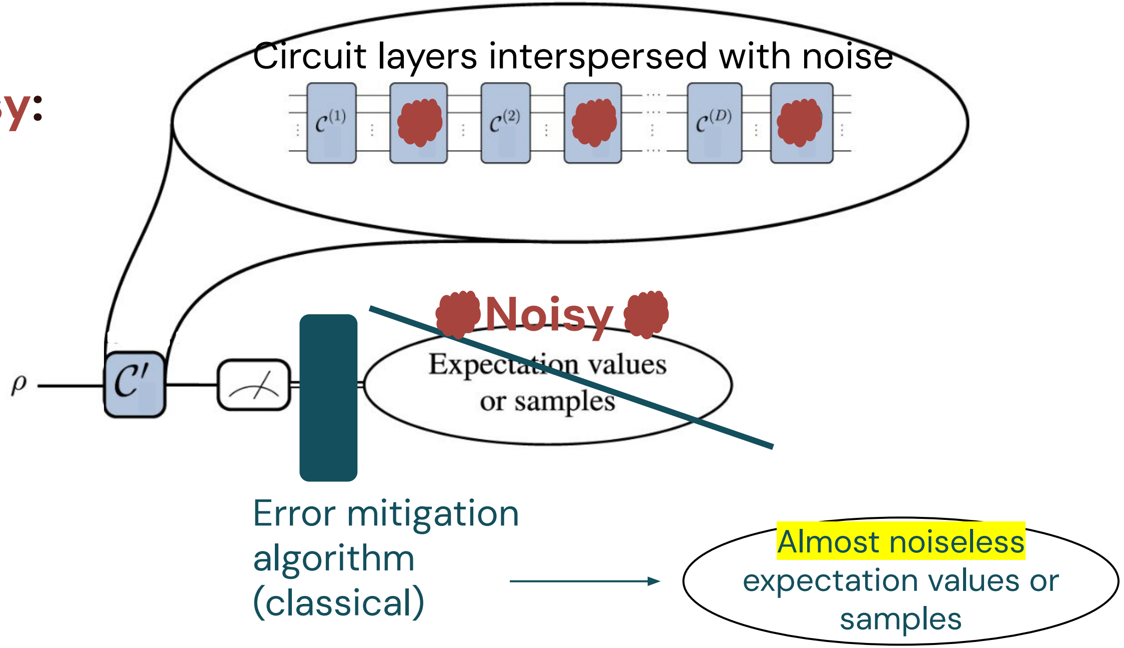

Example of error mitigation protocol

Zero-noise extrapolation:

1) Run the circuit of interest at amplified noise level λ (call this).
2) Measure

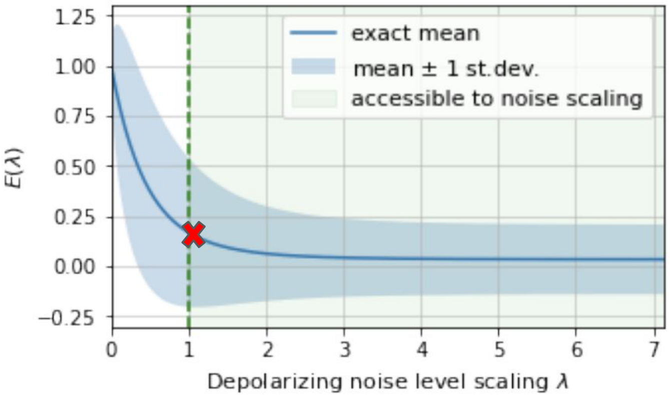

Plot taken from Giurgica-Tiron et al, 2020 IEEE International Conference on Quantum Computing and Engineering (QCE)

Example of error mitigation protocol

Zero-noise extrapolation:

1) Run the circuit of interest at amplified noise level λ (call this).
2) Measure
3) Repeat steps 1, 2 for different λ.

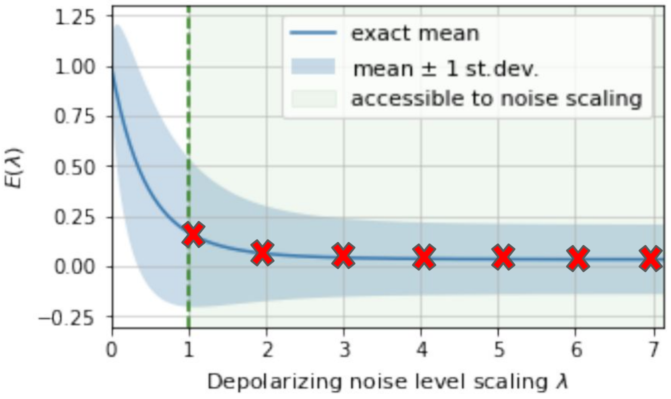

Plot taken from Giurgica-Tiron et al, 2020 IEEE International Conference on Quantum Computing and Engineering (QCE)

Example of error mitigation protocol

Zero-noise extrapolation:

1) Run the circuit of interest at amplified noise level λ (call this).
2) Measure
3) Repeat steps 1, 2 for different λ.
4) Output the extrapolated value

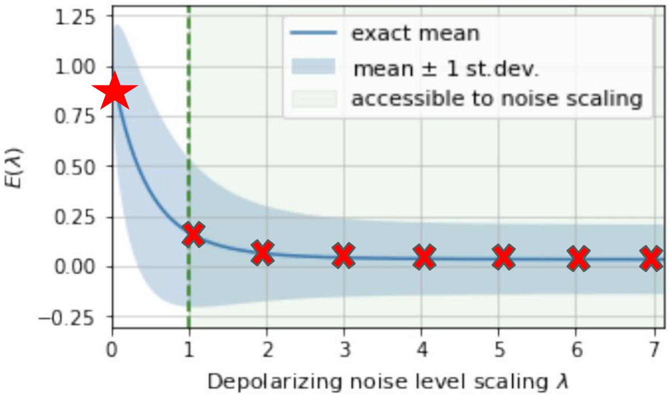

Plot taken from Giurgica-Tiron et al, 2020 IEEE International Conference on Quantum Computing and Engineering (QCE)

Our model of error mitigation

$\mathcal{A}$ = error mitigation algorithm

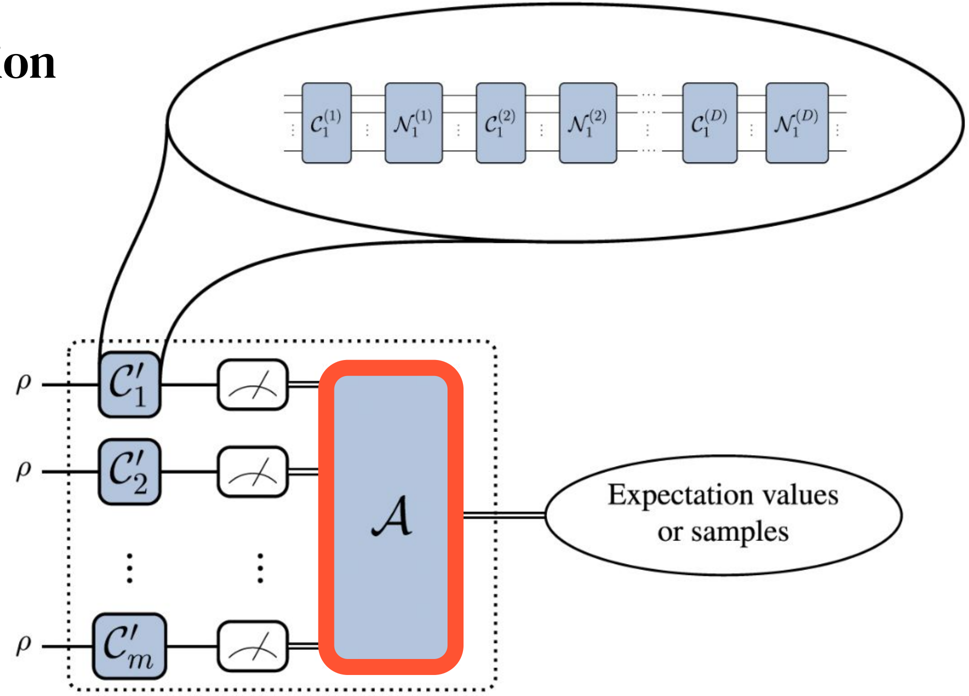

Our model of error mitigation

Quantum circuits:
n qubits, depth **D**

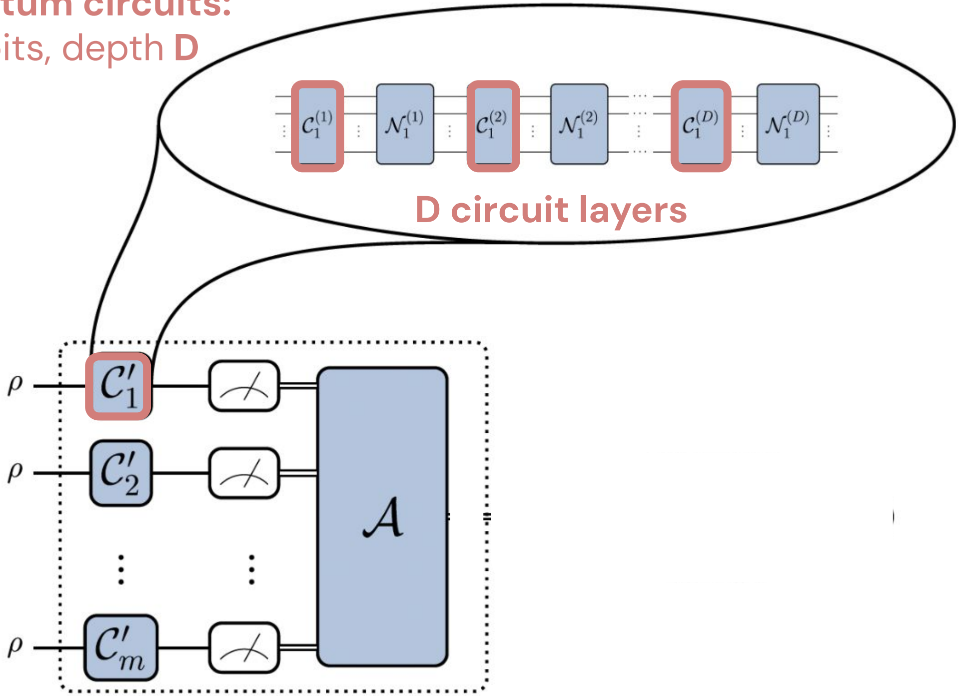

Our model of error mitigation

Quantum circuits:
n qubits, depth **D**

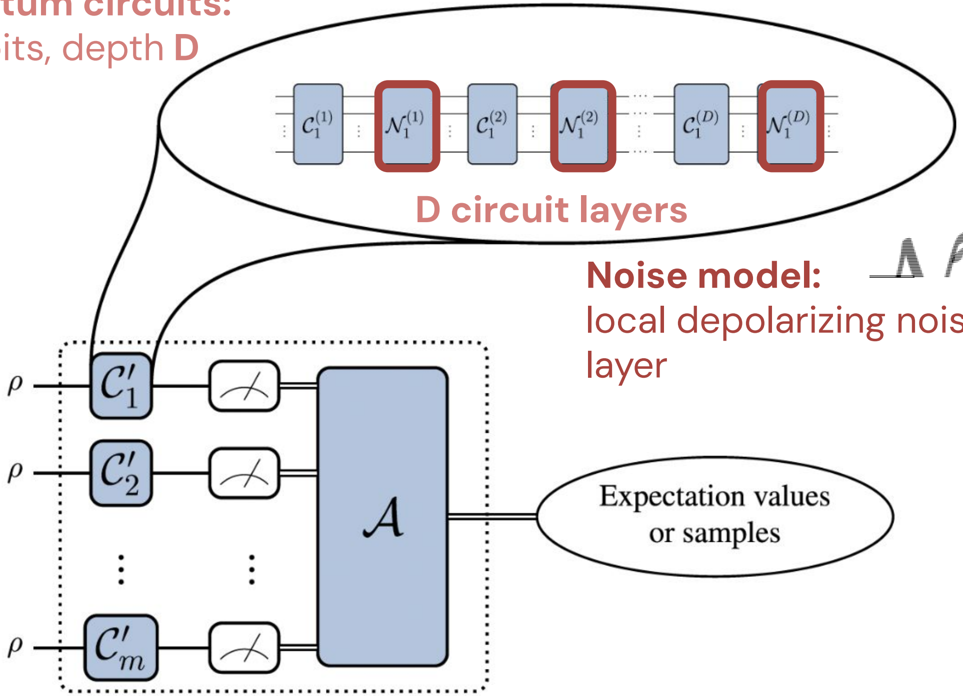

Noise model:
local depolarizing noise of strength **p** after each layer

Input to $\mathcal{A}$

Input:

Copies of [illegible] output by circuits with depolarizing noise of strength p

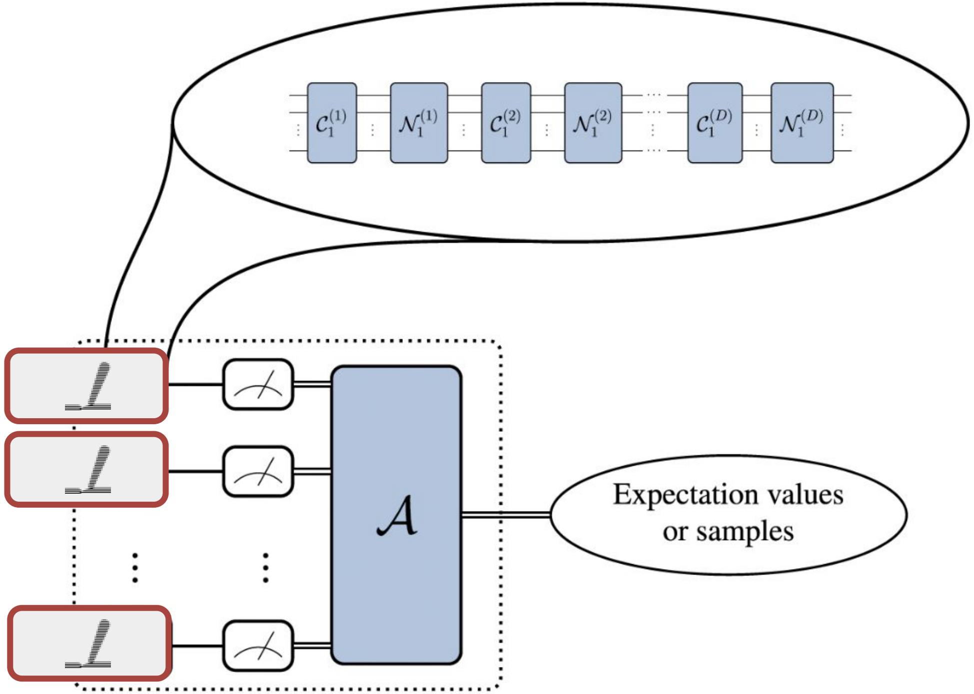

Output of $\mathcal{A}$

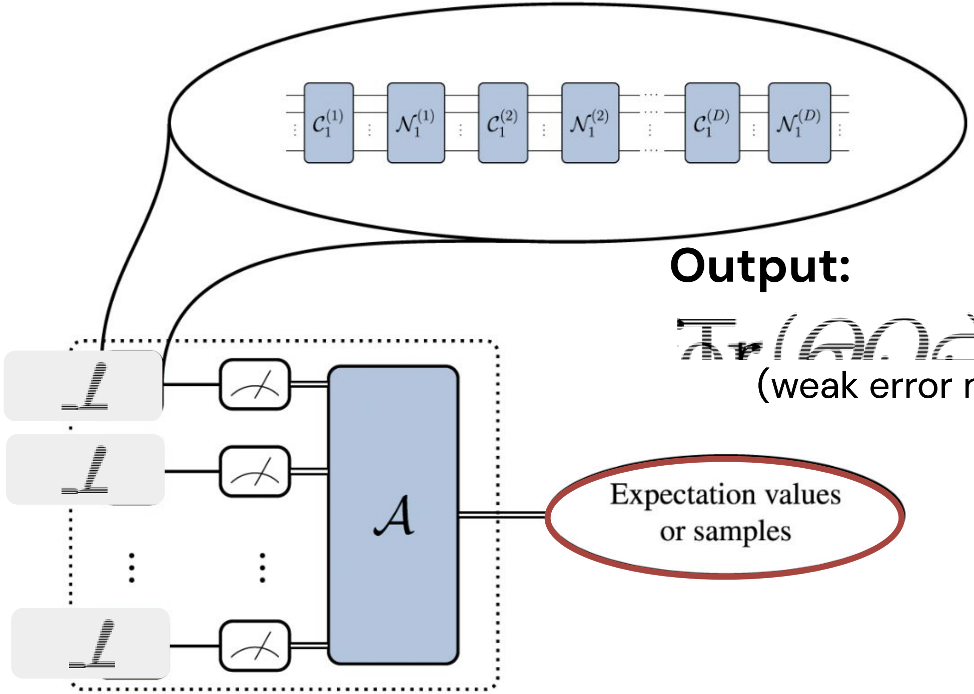

Output:

(weak error mitigation → e.g. VQE)

Output of $\mathcal{A}$

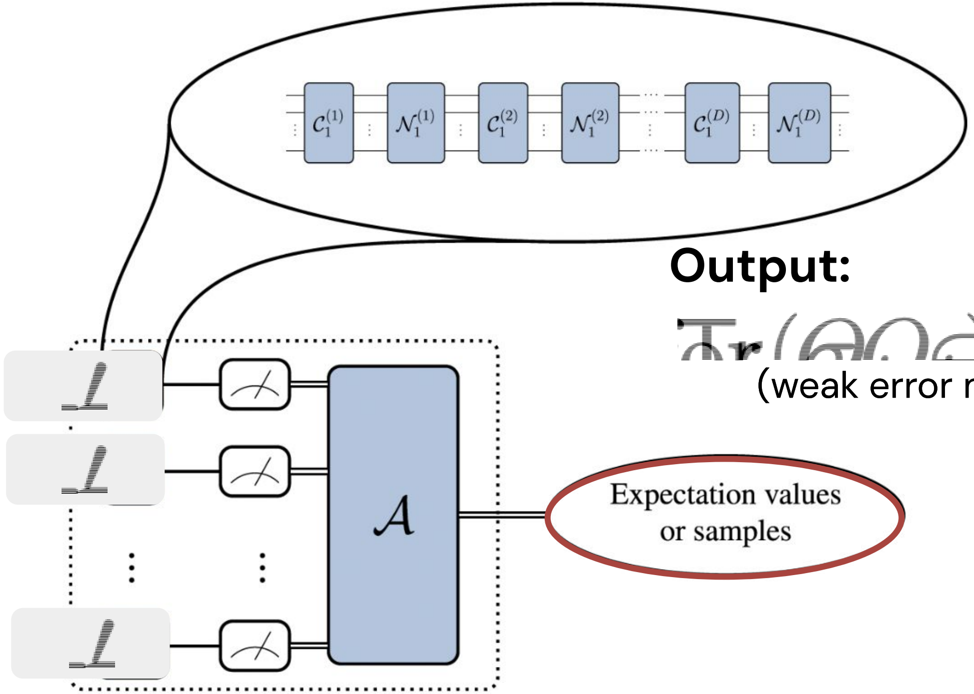

= noisy circuit output state

= noise*less* circuit output state

Output:

Tr([illegible]

(weak error mitigation → e.g. VQE)

Our question: sample complexity of error mitigation?

How many copies of [image] does [image] need...

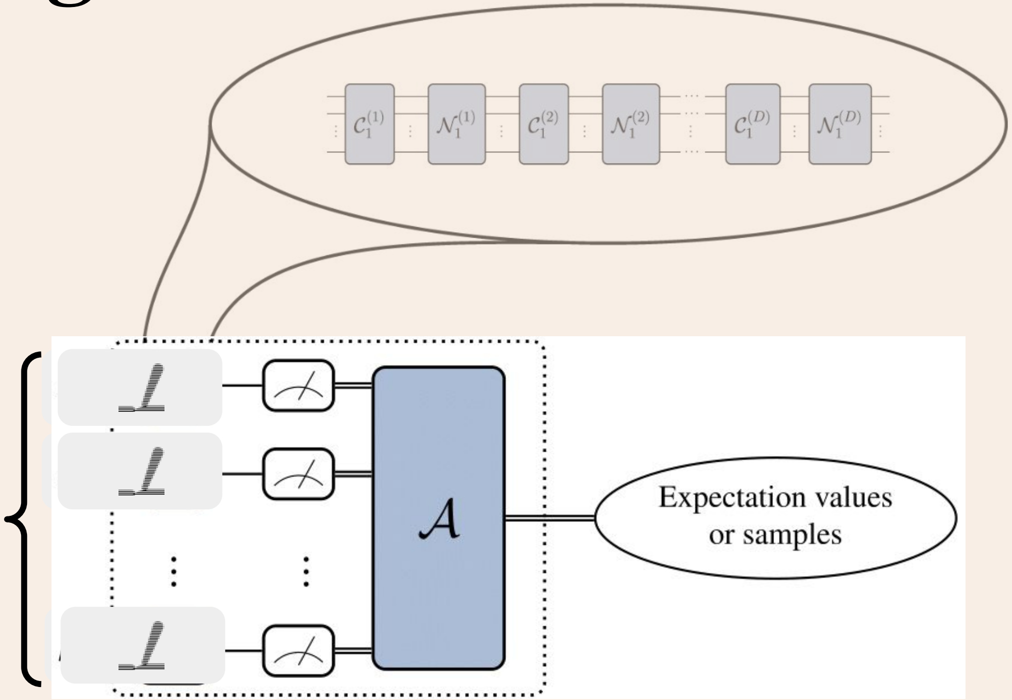

Our question: sample complexity of error mitigation?

How many copies of does need...

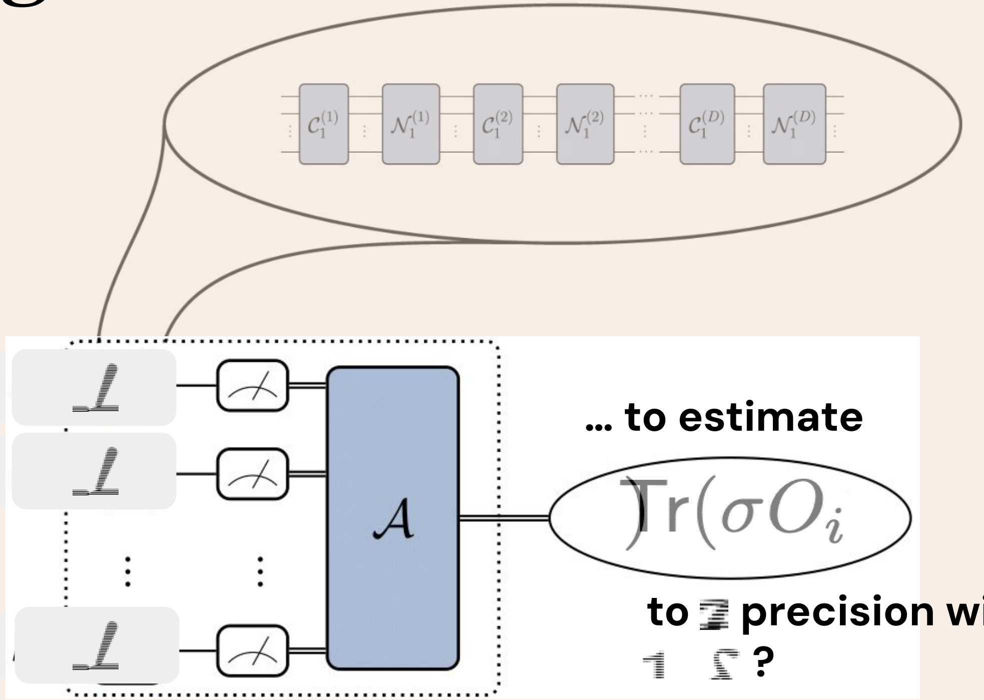

... to estimate

to precision with probability ?

Our question: sample complexity of error mitigation?

Relevant params:
n, D (circuit width and depth)
p (noise strength)

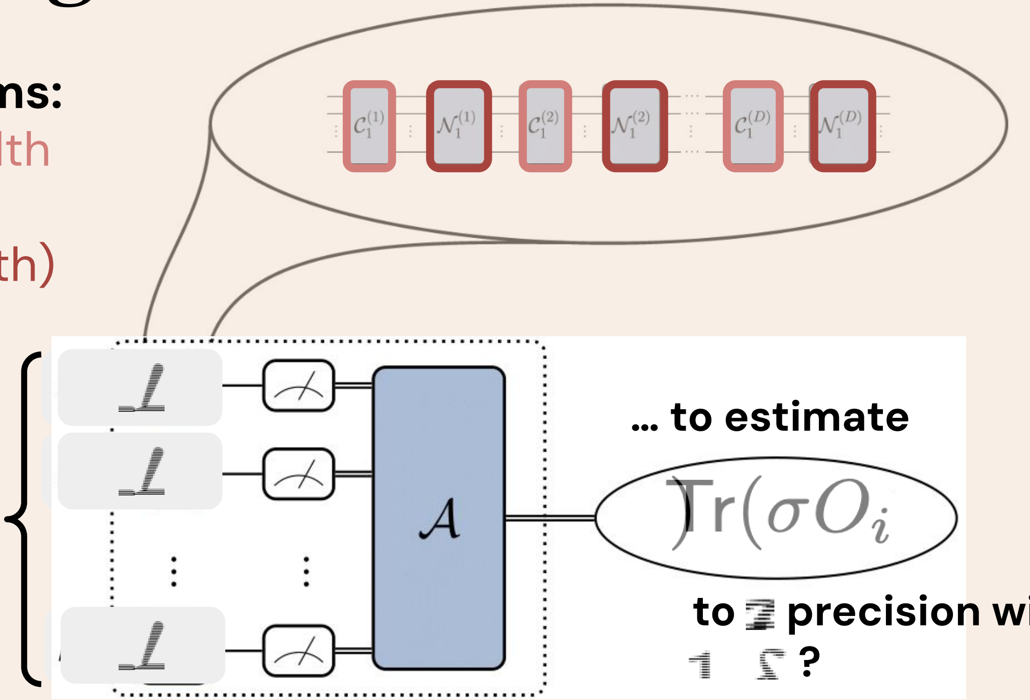

How many copies of σ **does** $\mathcal{A}$ **need...**

... to estimate $\mathrm{Tr}(\sigma O_i)$ **to** ϵ **precision with probability** $1-\delta$ **?**

Our lower bounds

How many copies of does need...

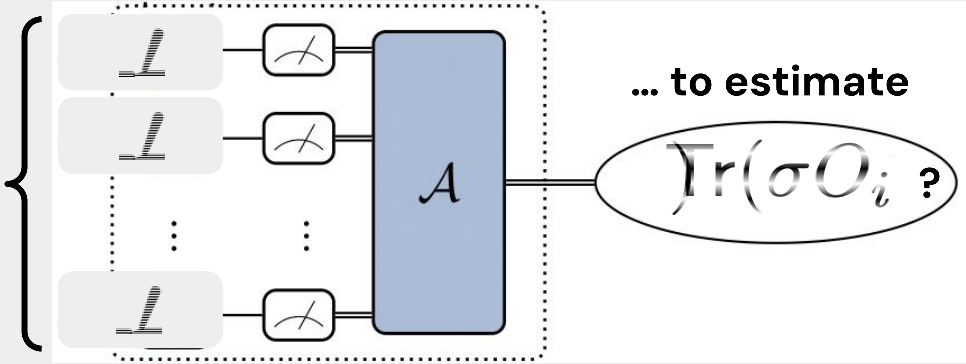

Relevant parameters:
n, D (circuit width/depth);
p (depolarizing noise strength)

Thm 1: if circuit outputting has depolarizing noise and is of depth .

How to interpret our results

We show: $\exp(\Omega(nD))$ runs of a depolarizing-noisy circuit are required for good error mitigation.

How to interpret our results

We show: $\exp(\Omega(nD))$ runs of a depolarizing-noisy circuit are required for good error mitigation.

- Previous belief: $\exp(\Omega(D))$ copies required.

How to interpret our results

We show: $\exp(\Omega(nD))$ runs of a depolarizing-noisy circuit are required for good error mitigation.

- Previous belief: $\exp(\Omega(D))$ copies required.
- But NISQ circuits are depth [illegible]: our result is exponentially stronger.

How to interpret our results

We show: $\exp(\Omega(nD))$ runs of a depolarizing-noisy circuit are required for good error mitigation.

- Previous belief: $\exp(\Omega(D))$ copies required.
- But NISQ circuits are depth [illegible]: our result is exponentially stronger.
- Loss of quantum advantage for error-mitigated algorithms may occur earlier than expected in the presence of noise

Proof intuition

Question: How many copies of $\perp$ are needed for EM?

Proof intuition

Question: How many copies of are needed for EM?

We construct such circuits:

Depolarizing noise (also causes rapid convergence to m.m. state)

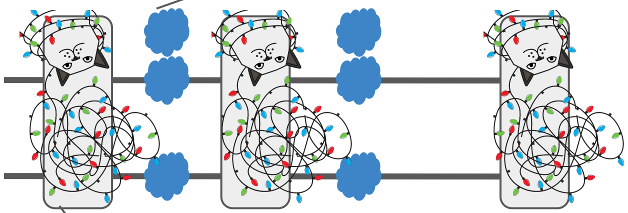

Extremely rapidly scrambling Clifford circuits
→ quick convergence to m.m. state

Proof intuition

Question: How many copies of  are needed for EM?

We construct such circuits:

Depolarizing noise (also causes rapid convergence to m.m. state)

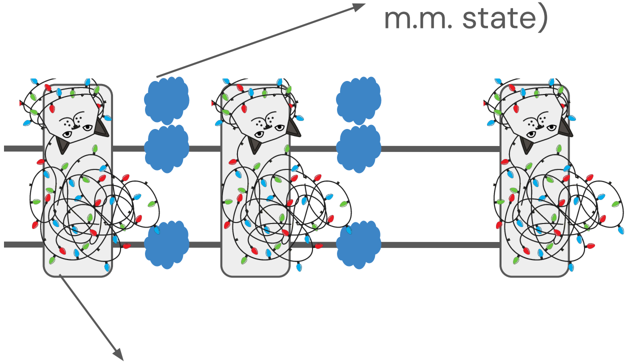

Extremely rapidly scrambling Clifford circuits
→ quick convergence to m.m. state

What this captures: **depolarizing noise and the circuits 'conspire' to push the state toward maximally mixed.** With no intermediate measurements, you are at the mercy of entropy generation.

What about…non-unital noise?

Not all noise is depolarizing!

See: T1 decay,
atom loss, photon loss.

Superconducting qubits (Google, IBM, Rigetti)

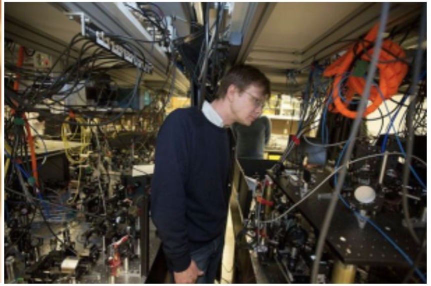

Neutral atom arrays (e.g. Lukin group/QuEra)

What about...non-unital noise?

Not all noise is depolarizing!

See: T1 decay,
atom loss, photon loss.

Superconducting qubits

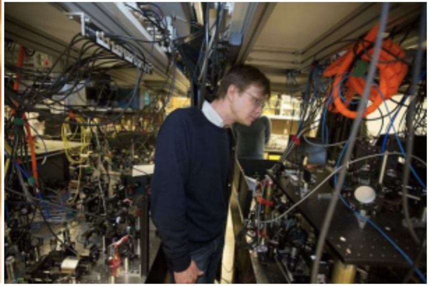

Neutral atom arrays

Error source	Time optimal	Smooth amplitude	Error type X, Y, Z, LG, AL**
Scattering* $\|1\rangle$	0.103% / 0.043%	0.036%	6%, 6%, 25%, 47%, 15%
Scattering $\|0\rangle$	0.019%	0.025%	7%, 7%, 14%, 62%, 10%
Rydberg $T_1 = 88\,\mu s$	0.113%	0.085%	2%, 2%, 6%, 23%, 67%
Rydberg $T_2^* = 3\,\mu s$	0.134%	0.089%	0%, 0%, 75%, 0%, 25%
Position fluct.	0.012%	0.054%	0%, 0%, 96%, 0%, 4%
Rydberg $m_J = -\frac{1}{2}$	0.06 - 0.15%	0.01%	
Total fidelity	**99.53 - 99.62%**	**99.70%**	**2%, 2%, 41%, 17%, 38%**

From [High-fidelity parallel entangling gates on a neutral-atom quantum computer]

What about…non-unital noise?

Not all noise is depolarizing!

See: T1 decay,
atom loss, photon loss.

Superconducting qubits (Google, IBM, Rigetti)

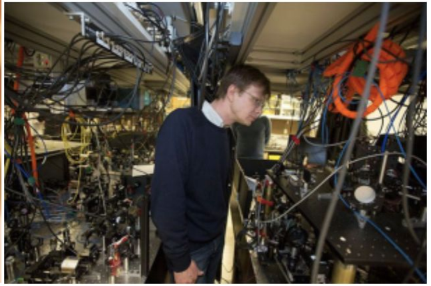

Neutral atom arrays (e.g. Lukin group/QuEra)

Definition of non-unital quantum channel: doesn't preserve the m.m. state.

Canonical example: **amplitude damping noise!** ~ partial **reset**-to-|0>

What about...non-unital noise?

Not all noise is depolarizing!

See: T1 decay,
atom loss, photon loss.

Superconducting qubits (Google, IBM, Rigetti)

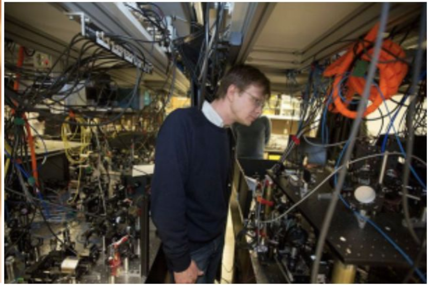

Neutral atom arrays (e.g. Lukin group/QuEra)

Definition of non-unital quantum channel: doesn't preserve the m.m. state.

Canonical example: **amplitude damping noise!** ~ partial **reset**-to-|0>

Non-unital noise: physically important yet shockingly understudied!

Fantastic features of non-unital noise and where to find them:

I will use amplitude-damping noise as a concrete example; behaviors hold for general non-unital noise.

1. Amplitude-damping and depolarizing noise *have different fixed points*

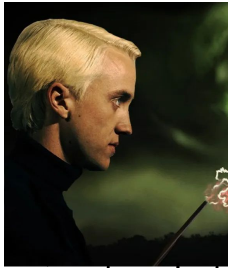

Depolarizing noise tends to “scramble” the distribution by **increasing entropy**.

Fixed point: maximally-mixed state

1. Amplitude-damping and depolarizing noise *have different fixed points*

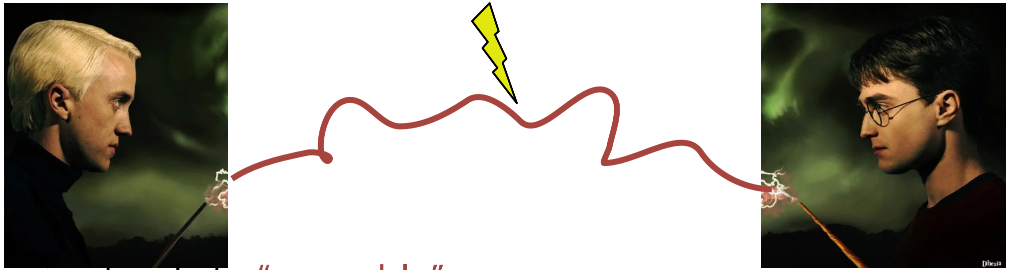

Depolarizing noise tends to "scramble" the distribution by **increasing entropy**.

Fixed point: maximally-mixed state

Amplitude damping noise tries to "unscramble" the distribution by **decreasing entropy**!

Fixed point: $|0\rangle$

1. Amplitude-damping and depolarizing noise *have different fixed points*

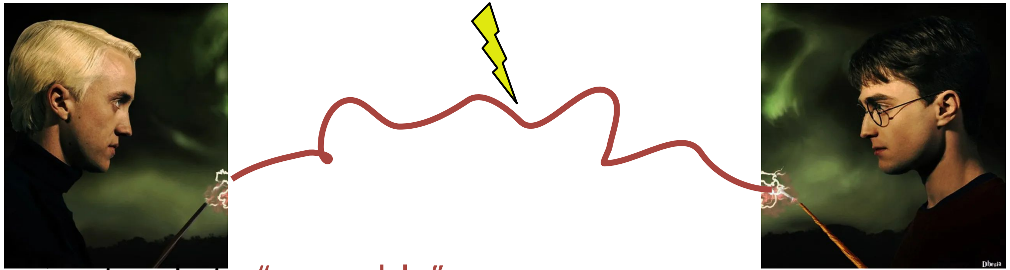

Depolarizing noise tends to "scramble" the distribution by **increasing entropy**.

Fixed point: maximally-mixed state

Amplitude damping noise tries to "unscramble" the distribution by **decreasing entropy**!

Fixed point: |0>

Consequence: The same circuit, with depolarizing noise and with amplitude damping noise, acts very differently on the Bloch sphere!

1. Amplitude-damping and depolarizing noise *have different fixed points*

Circuit with depolarizing noise

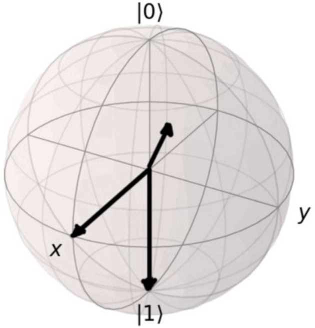

denotes noise fixed point

Noise drives input **towards the center** of Bloch sphere.

1. Amplitude-damping and depolarizing noise *have different fixed points*

Circuit with depolarizing noise

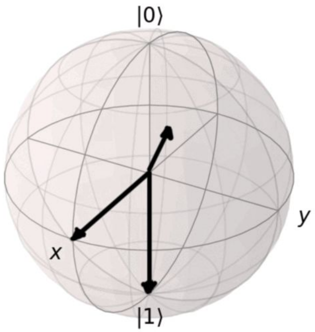

★ denotes noise fixed point

Noise drives input t**owards the center** of Bloch sphere.

Circuit with amplitude damping noise

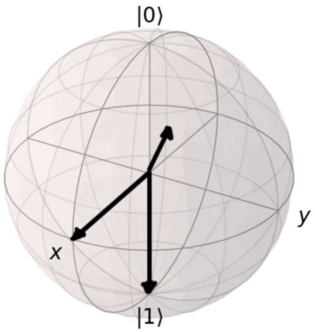

Noise drives input t**owards the north pole** of Bloch sphere.

Consequence: No barren plateaus

Circuits used for quantum machine learning often display barren plateaus, making them hard to optimize.

Consequence: No barren plateaus

Circuits used for quantum machine learning often display barren plateaus, making them hard to optimize.

We show: optimization landscape under non-unital noise is **not barren**.

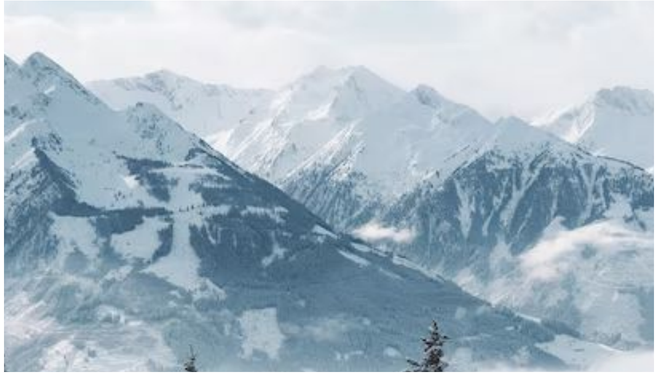

Consequence: No barren plateaus

Circuits used for quantum machine learning often display barren plateaus, making them hard to optimize.

We show: optimization landscape under non-unital noise is **not barren**.

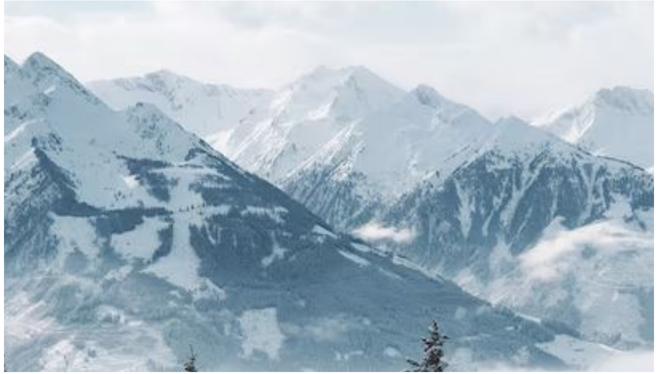

Important limitation: we don't show that the resulting peak is at the location of minimum energy, but others have tackled this.

2. The miracle of non-unital noise: fault-tolerance "for free"?

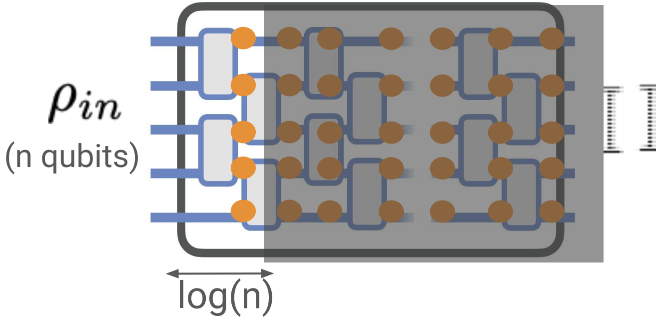

With no error correction, **depolarizing noise** erases all information within a circuit after log(n) depth.

2. The miracle of non-unital noise: fault-tolerance "for free"?

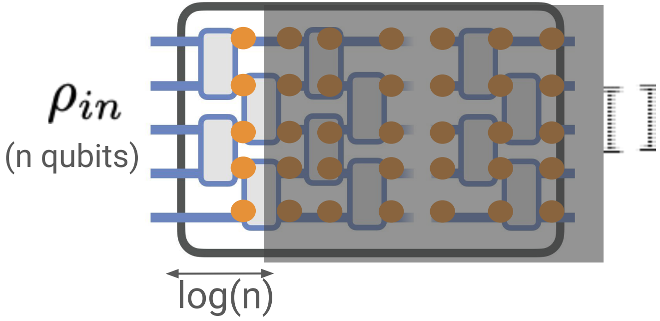

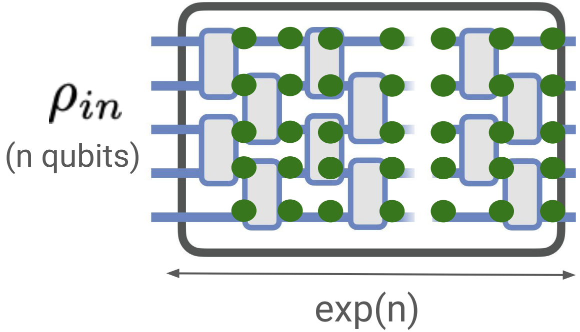

With no error correction, **depolarizing noise** erases all information within a circuit after log(n) depth.

Quantum refrigerator proposal:
Even with no error correction, **non-unital noise** can be leveraged to compute for exp(n) depth!

2. The miracle of non-unital noise: fault-tolerance “for free”?

Key idea [Aharonov/Ben-Or 1999]: it’s possible to do fault-tolerant quantum computation **without intermediate measurements** if *one supplies fresh auxiliary qubits in the state |0>.*

2. The miracle of non-unital noise: fault-tolerance "for free"?

Key idea [Aharonov/Ben-Or 1999]: it's possible to do fault-tolerant quantum computation **without intermediate measurements** if *one supplies fresh auxiliary qubits in the state |0>*.

Quantum refrigerator [Gottesman, Ben-Or 2013]: Let's use **non-unital noise** to supply us with qubits in the state |0>!

Remember: Amplitude-damping noise ~ **reset**-to-all-0s

Amplitude-damping noise ~ **reset**-to-all-0s

2. The miracle of non-unital noise: fault-tolerance "for free"?

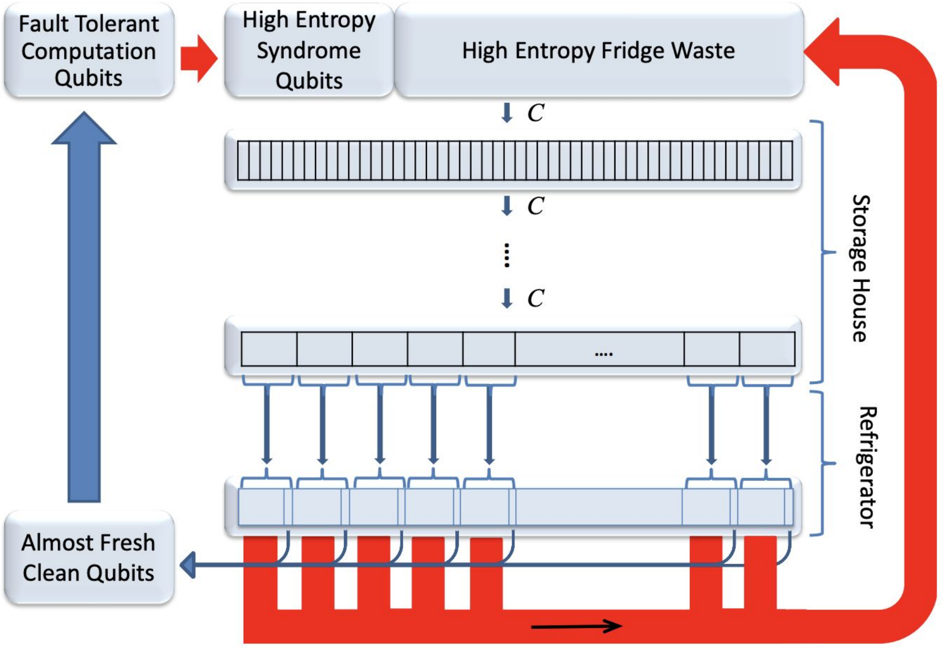

Quantum refrigerator, Ben-Or, Gottesman, Hassidim (arXiv 1301.1995)

Amplitude-damping noise ~ **reset**-to-all-0s

2. The miracle of non-unital noise: fault-tolerance "for free"?

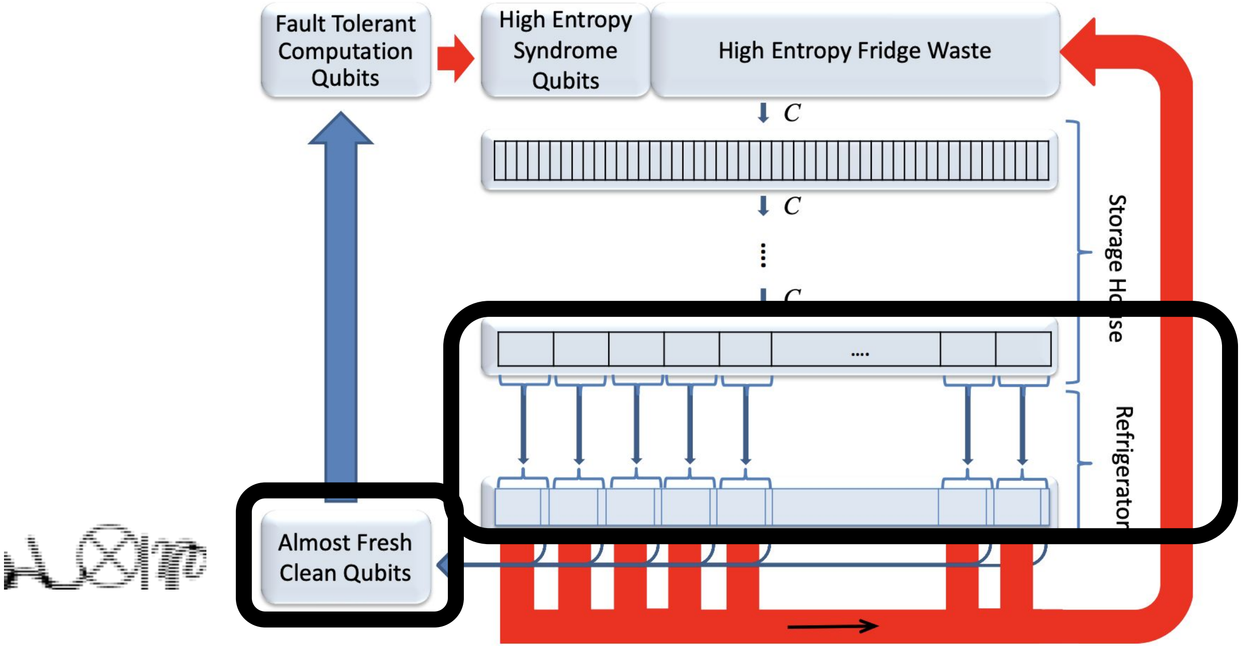

Implemented by non-unital noise (drives towards [illegible])

Quantum refrigerator, Ben-Or, Gottesman, Hassidim (arXiv 1301.1995)

Amplitude-damping noise ~ **reset**-to-all-0s

2. The miracle of non-unital noise: fault-tolerance "for free"?

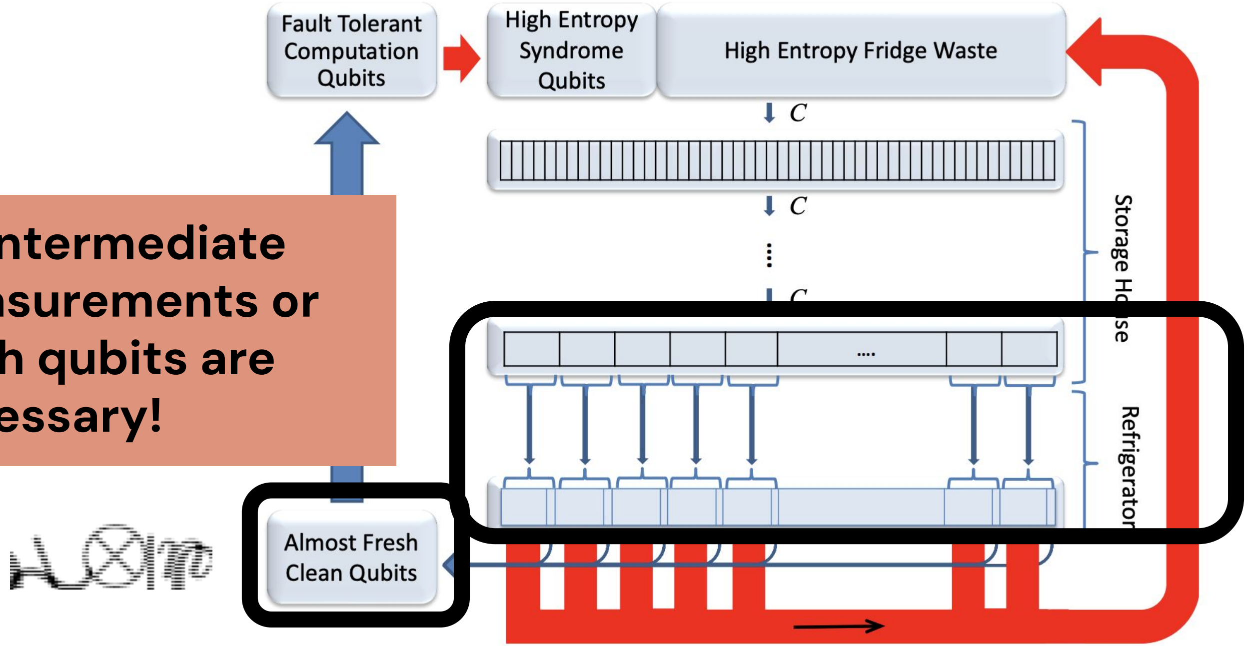

No intermediate measurements or fresh qubits are necessary!

Implemented by non-unital noise (drives towards |0⟩)

Quantum refrigerator, Ben-Or, Gottesman, Hassidim (arXiv 1301.1995)

Application: Error recovery beyond error mitigation?

Error mitigation

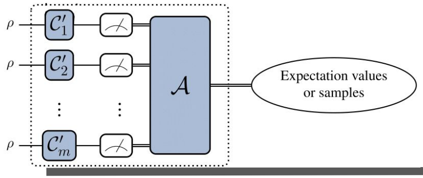

Error correction

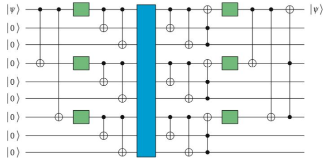

Few quantum resources but high sample complexity

Sample-efficient but large quantum resource burden

Application: Error recovery beyond error mitigation?

Error mitigation

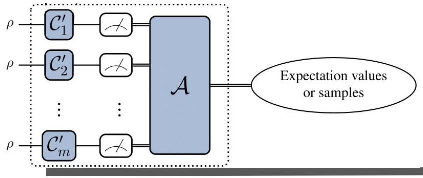

Goldilocks zone?

Error correction

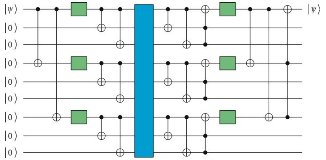

Few quantum resources but high sample complexity

Few measurements; no/few ancillas; some encoded gates
"Best of both worlds"!

Sample-efficient but large quantum resource burden

Is the "quantum refrigerator" exactly what we need?

A limitation to bear in mind

For a circuit with *random* gates, non-unital noise (like depolarizing noise) also creates **effectively shallow circuits**.

A limitation to bear in mind

For a circuit with *random* gates, non-unital noise (like depolarizing noise) also creates **effectively shallow circuits**.

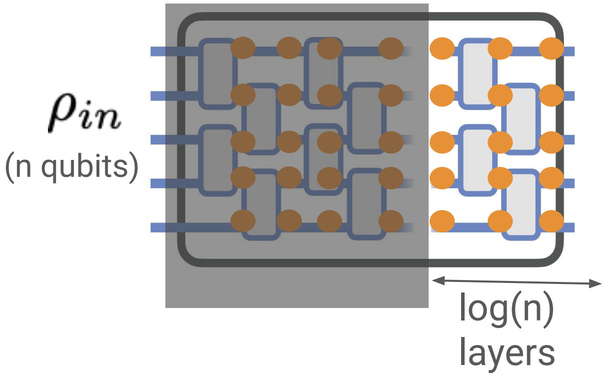

: **Single-qubit nonunital noise**

All gates more than log(n) layers from the last layer cannot influence expectation values by much!

Million-dollar question: Can you still get quantum advantage via nonunital noise for a sampling task?

Million-dollar question: Can you still get quantum advantage via nonunital noise for a sampling task?

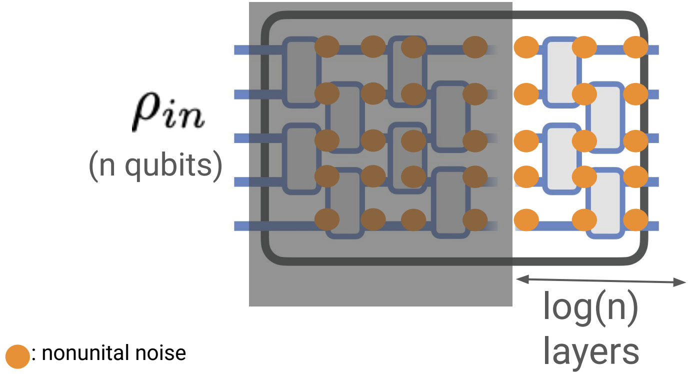

Effective shallow depth picture only holds for local expectation values.

Million-dollar question: Can you still get quantum advantage via nonunital noise for a *sampling* task?

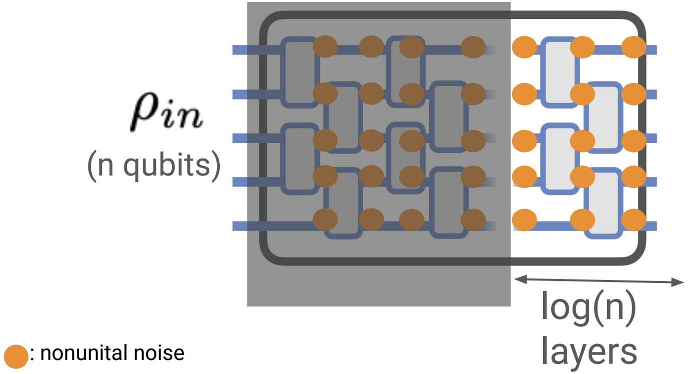

Effective shallow depth picture only holds for local expectation values.

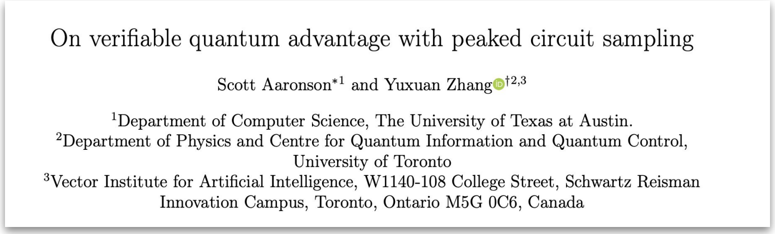
On verifiable quantum advantage with peaked circuit sampling

Scott Aaronson*[1] and Yuxuan Zhang†[2,3]

[1]Department of Computer Science, The University of Texas at Austin.
[2]Department of Physics and Centre for Quantum Information and Quantum Control, University of Toronto
[3]Vector Institute for Artificial Intelligence, W1140-108 College Street, Schwartz Reisman Innovation Campus, Toronto, Ontario M5G 0C6, Canada

Can non-unital noise help you to hide a peak in a **shallow depth circuit**? Hide = classical computer cannot tell from circuit description if peaked or random.

Exponentially tighter bounds on limitations of quantum error mitigation

Yihui Quek,[1] Daniel Stilck França,[2,3,1] Sumeet Khatri,[1] Johannes Jakob Meyer,[1] and Jens Eisert[1,4]

[1]*Dahlem Center for Complex Quantum Systems, Freie Universität Berlin, 14195 Berlin, Germany*
[2]*Department of Mathematical Sciences, University of Copenhagen, 2100 København, Denmark*
[3]*Univ Lyon, Inria, ENS Lyon, UCBL, LIP, F-69342, Lyon Cedex 07, France.*
[4]*Helmholtz-Zentrum Berlin für Materialien und Energie, 14109 Berlin, Germany*

(Dated: November 14, 2022)

Quantum error mitigation has been proposed as a means to combat unwanted and unavoidable errors in near-term quantum computing by classically post-processing outcomes of multiple quantum circuits. It does so in a fashion that requires no or few additional quantum resources, in contrast to fault-tolerant schemes that come along with heavy overheads. Error mitigation leads to noise reduction in small schemes of quantum computation. In this work, however, we identify strong limi-
tations to the degree to which quantum noise can be effectively 'undone' for larger system sizes. We

Noise-induced shallow circuits and absence of barren plateaus

Antonio Anna Mele,[1,*] Armando Angrisani,[2,3,†] Soumik Ghosh,[4]
Sumeet Khatri,[1] Jens Eisert,[1,5] Daniel Stilck França,[6,‡] and Yihui Quek[7,§]

[1]*Dahlem Center for Complex Quantum Systems, Freie Universität Berlin, 14195 Berlin, Germany*
[2]*LIP6, CNRS, Sorbonne Université, 75005 Paris, France*
[3]*Institute of Physics, Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland*
[4]*Department of Computer Science, University of Chicago, Chicago, Illinois 60637, USA*
[5]*Fraunhofer Heinrich Hertz Institute, 10587 Berlin, Germany*
[6]*Univ Lyon, ENS Lyon, UCBL, CNRS, Inria, LIP, F-69342, Lyon Cedex 07, France*
[7]*Departments of Mathematics and Physics, Massachusetts Institute of Technology, 182 Memorial Drive, Cambridge, MA 02138, USA*

(Dated: October 11, 2024)

Motivated by realistic hardware considerations of the pre-fault-tolerant era, we comprehensively study the impact of uncorrected noise on quantum circuits. We first show that any noise 'truncates' most quantum circuits to effectively logarithmic depth, in the task of estimating observable expectation values. We then prove that quantum circuits under any non-unital noise exhibit lack of barren plateaus for cost functions composed of local

EPFL

Postdoc, visitor and maybe PhD positions available!