Title: Characterizing useful quantum computers

Date: Feb 28, 2018 04:00 PM

URL: http://pirsa.org/18020110

Abstract: <p>Quantum computers can only offer a computational advantage when they have sufficiently many qubits operating with sufficiently small error rates. In this talk, I will show how both these requirements can be practically characterized by variants of randomized benchmarking protocols. I will first show that a simple modification to protocols based on randomized benchmarking allows multiplicative-precision estimates of error rates. I will then outline a new protocol for estimating the fidelity of arbitrarily large quantum systems using only single-qubit randomizing gates.
<!--[if !supportLineBreakNewLine]-->
<!--[endif]--></p>

Characterizing useful quantum computers

Joel Wallman

U. Innsbruck: A. Erhart, R. Stricker, E. Martinez, D. Nigg, P. Schindler, T. Monz, R. Blatt
U. Waterloo: J. Emerson
U. Sydney: S. Flammia, R. Harper

Perimeter Seminar
February 28th 2018

Outline

- What does a quantum computer need to be useful?
- Estimators and sequence lengths for randomized benchmarking
- Limitations of randomized benchmarking
- Cycle benchmarking
- Experimental implementation in ion traps

Useful quantum computers

Useful quantum computers need:

1. Many qubits
2. Universal(ish) operations
3. Low error rates

If the quantum computer has too few qubits or too restricted a set of operations, it can be efficiently simulated

If the quantum computer has too much noise, the output is unreliable

How do we characterize these 3 features? Randomized benchmarking

Randomized benchmarking

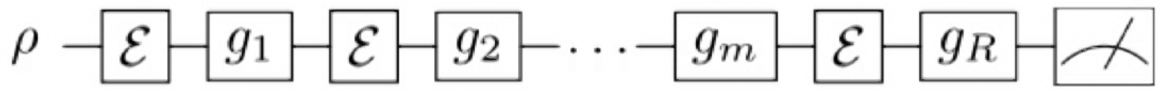

Apply a random sequence of m+1 gates from a group that multiplies to the identity.

Average probability of an outcome z (or a set of outcomes) over all sequences of length m is

$$\mathrm{Pr(z|m) = Ap^m + B}$$

Decay parameter is linearly related to the fidelity

$$\mathcal{F}(\tilde{\mathcal{E}}) = \int d\psi \, \mathrm{Tr} \left[\psi \tilde{\mathcal{E}}(\psi) \right]$$

Two key issues:

1. How does the error in the estimate of the error rate scale with the error rate?
2. How do we efficiently characterize a universal gate set?

Magesan, Gambetta, and Emerson, PRL 106, 180504 (2011)

Randomized benchmarking

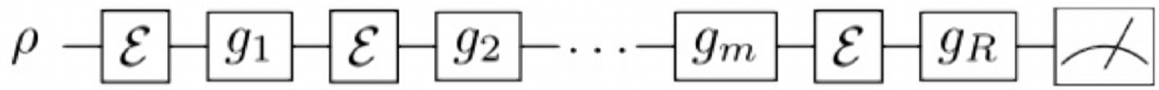

SPAM parameters are nuisances.

To set B to ½, randomly choose to compile to any gate X that flips the final outcome and set

$$\Pr(\mathrm{z}|\mathrm{m}) \rightarrow 1 - \Pr(\mathrm{z}|\mathrm{m})$$

Knill *et al.*, PRA 77, 012307 (2008)

Estimating decay rates

Two parameters, so can solve for two values of m:

$$p^{m_2-m_1} = \frac{2\mathrm{Pr}(\mathrm{z}|\mathrm{m}_2)-1}{2\mathrm{Pr}(\mathrm{z}|\mathrm{m}_1)-1}$$

Choose m_1=4 to avoid gate-dependent perturbations

Experimentally, have estimates $\hat{\mathrm{Pr}}(z|m_i) = \mathrm{Pr}(\mathrm{z}|\mathrm{m_i}) + \epsilon_\mathrm{i}$

Conservative uncertainty on estimate provided $\epsilon_2 \ll \hat{p}^{m_2-m_1}$ is

$$p \approx \hat{p} + \frac{2\epsilon_1 + 2\epsilon_2}{m_2 - m_1}$$

Choosing $m_2 \approx 1/r$ where r = 1-f gives multiplicative precision:

$$\hat{r} \approx r + 2r(\epsilon_1 + \epsilon_2)$$

JJW, Quantum 2, 47 (2018), Helsen *et al*, arXiv:1701.04299

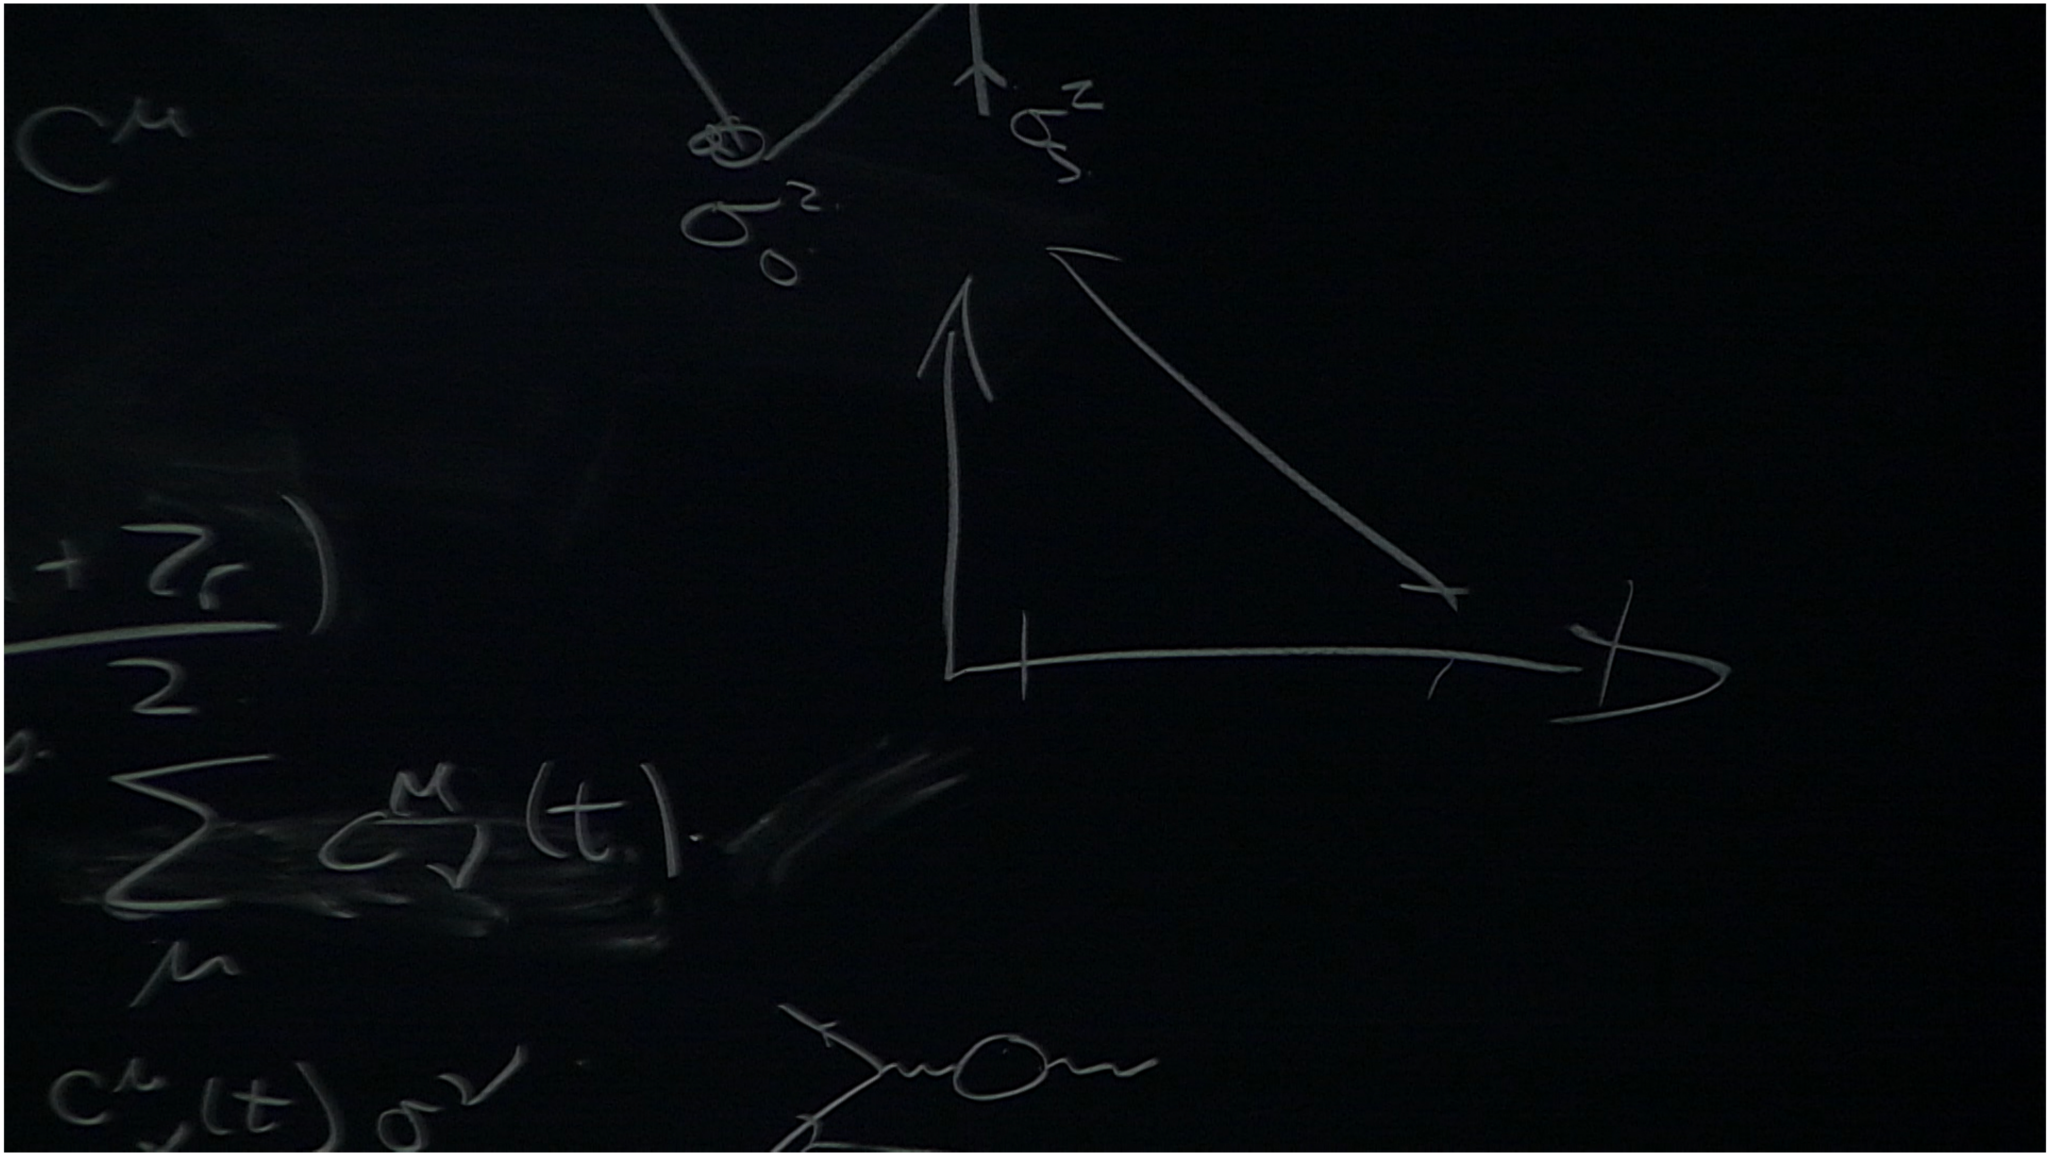

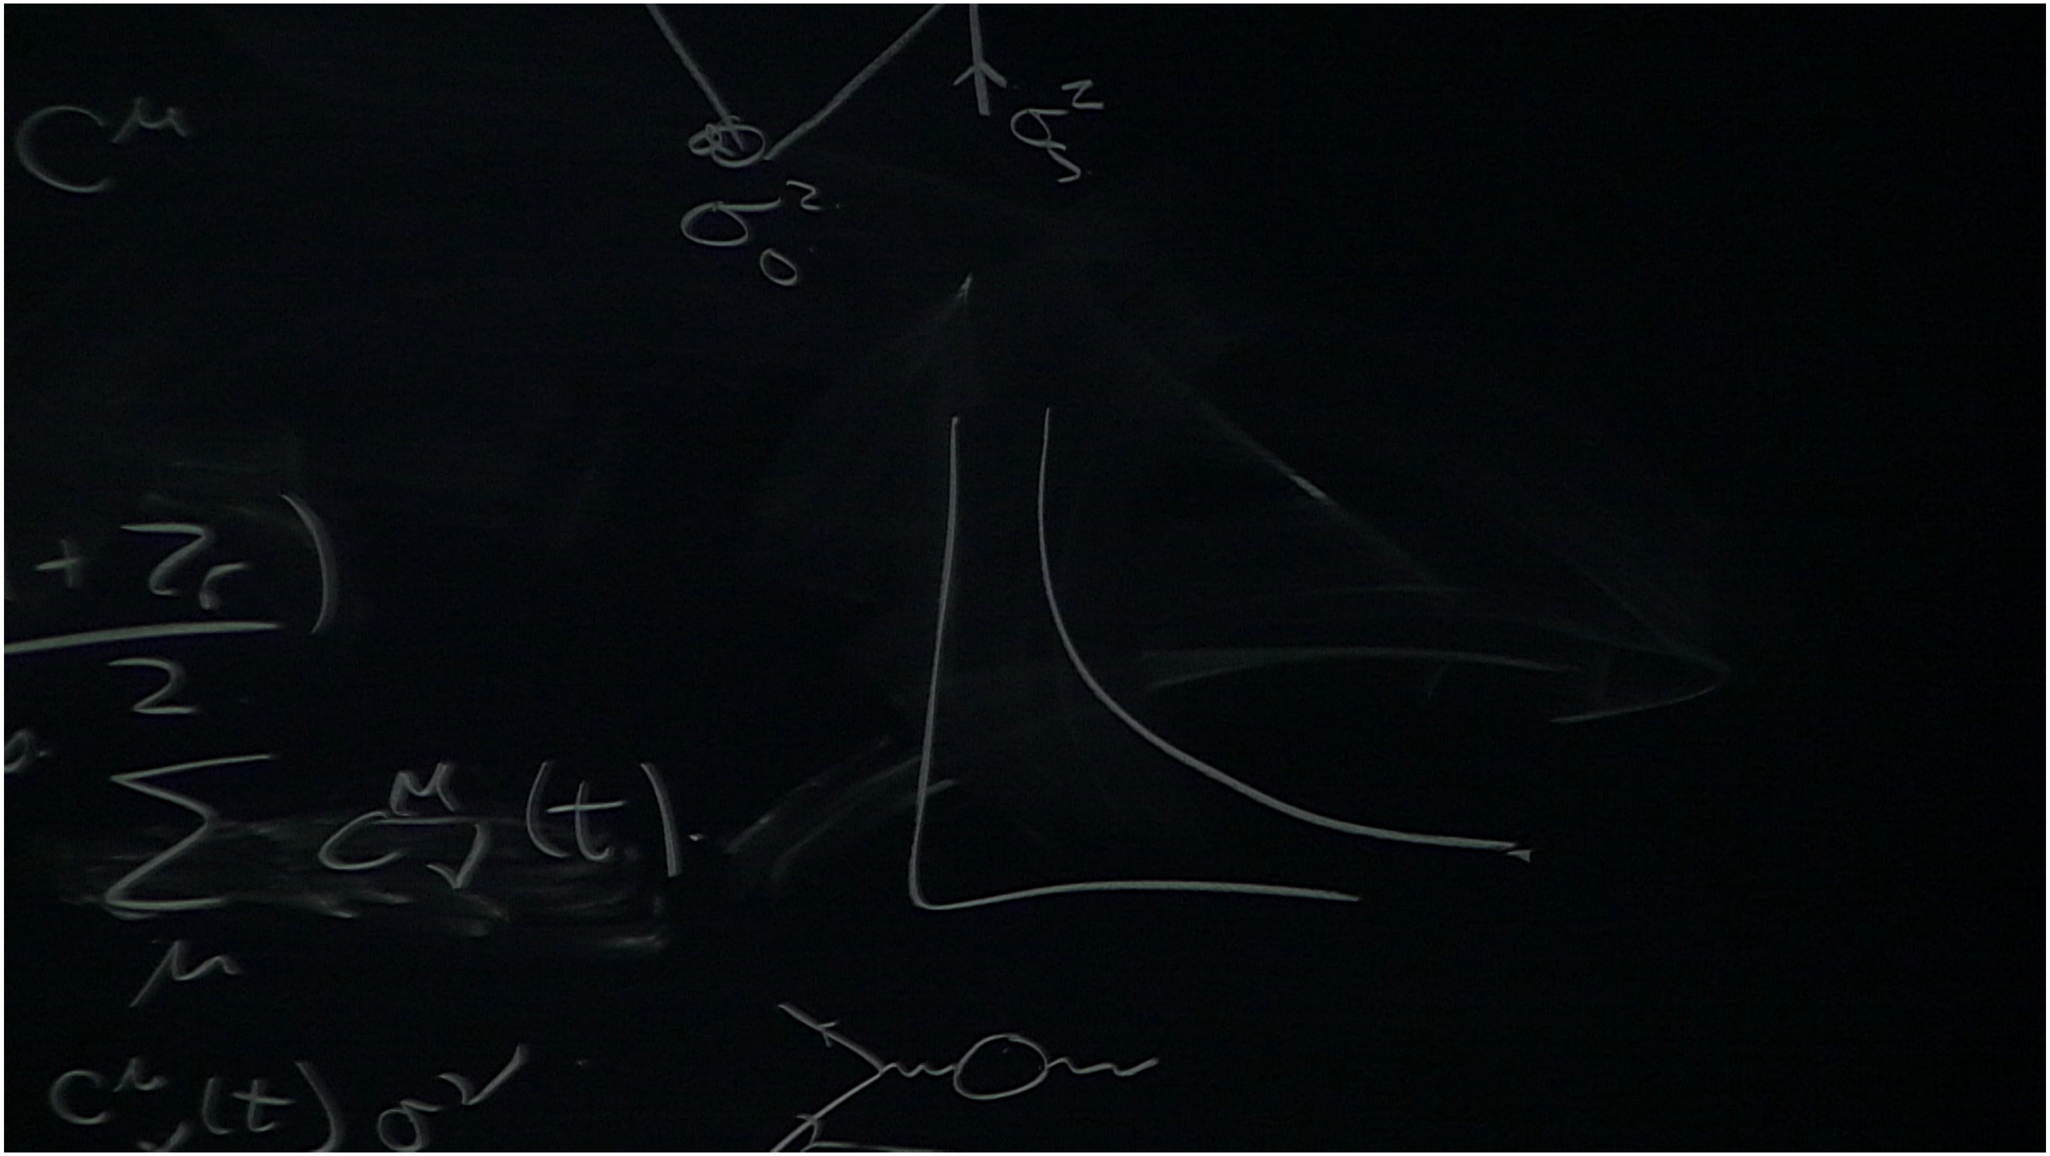

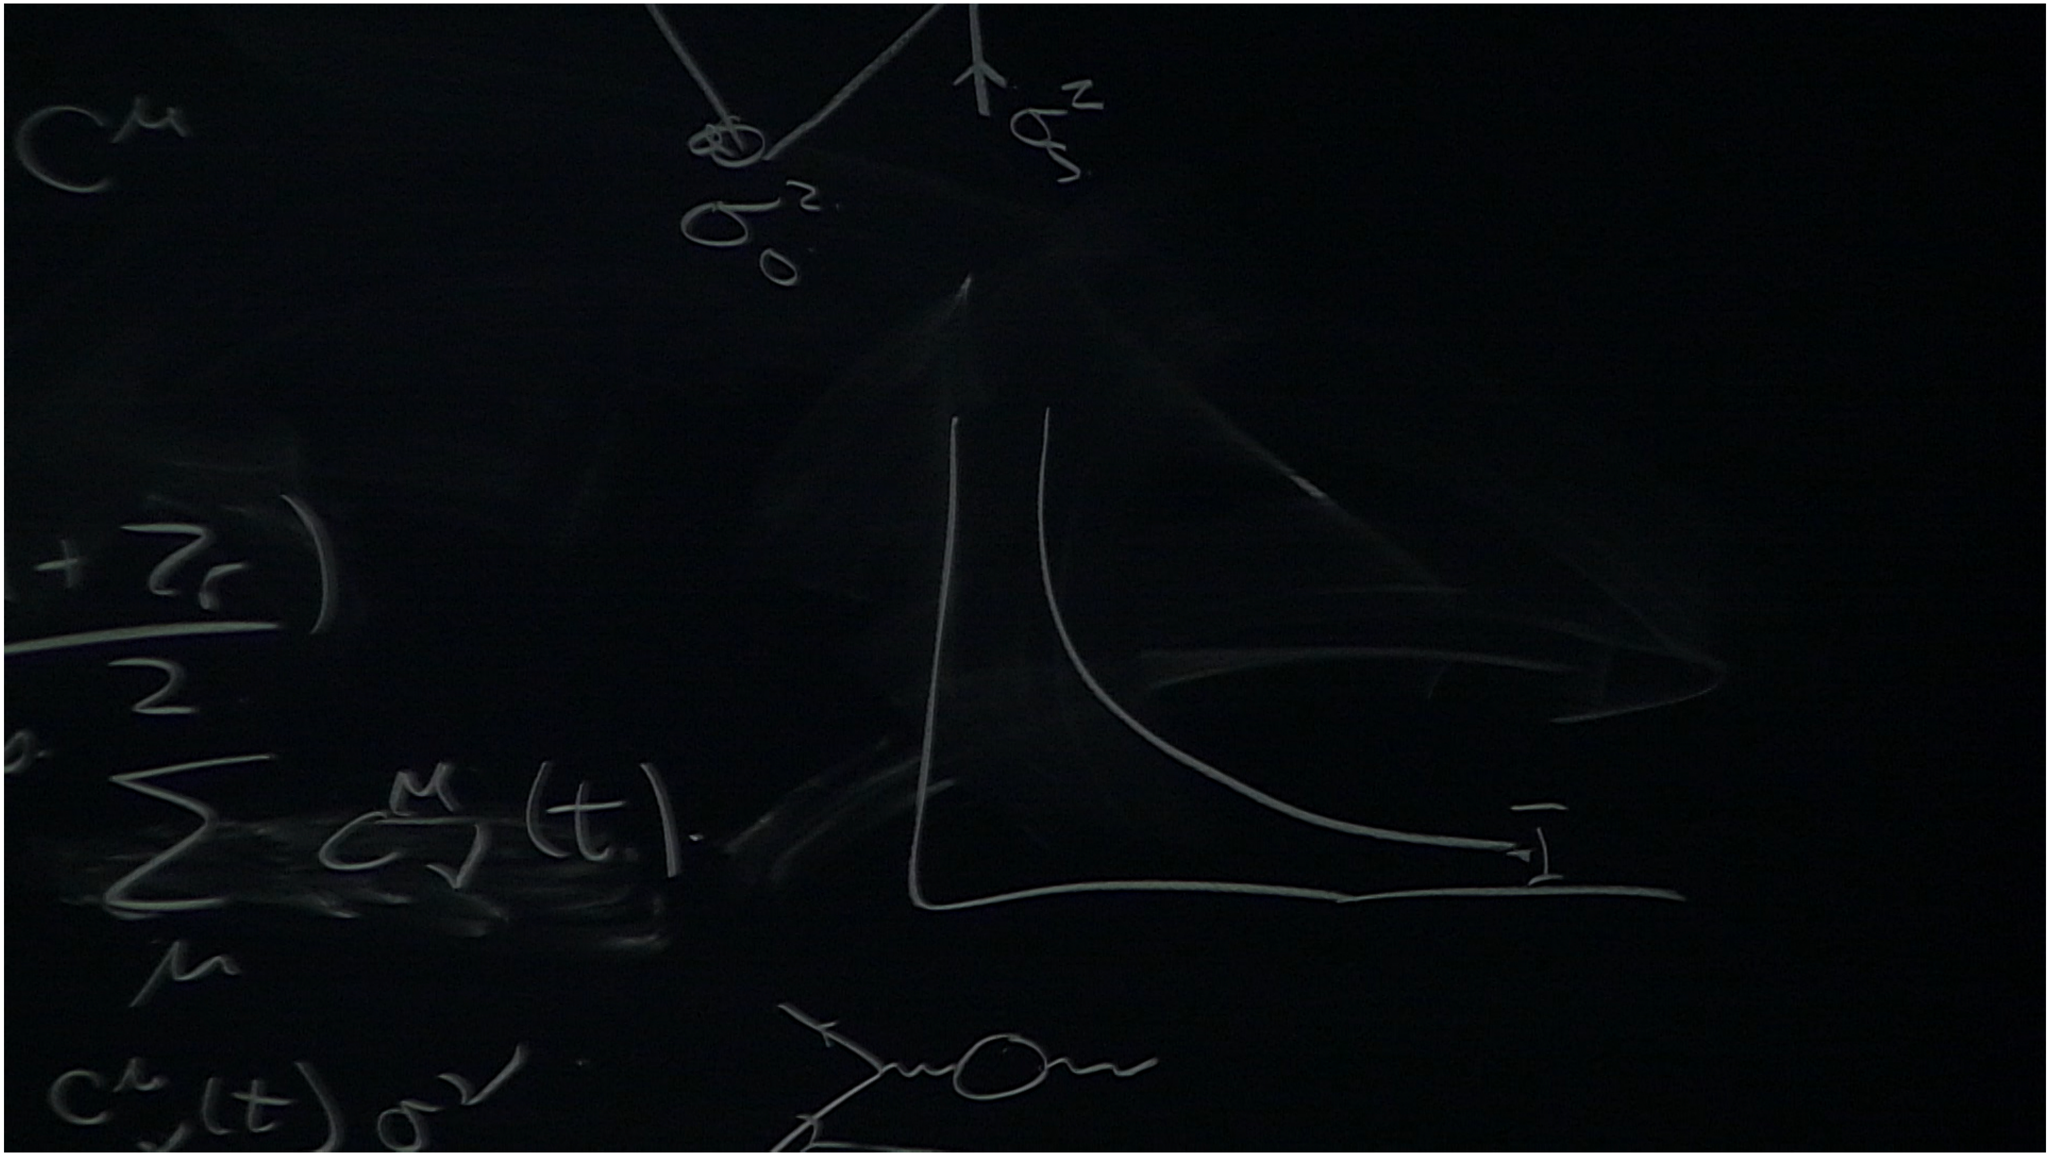

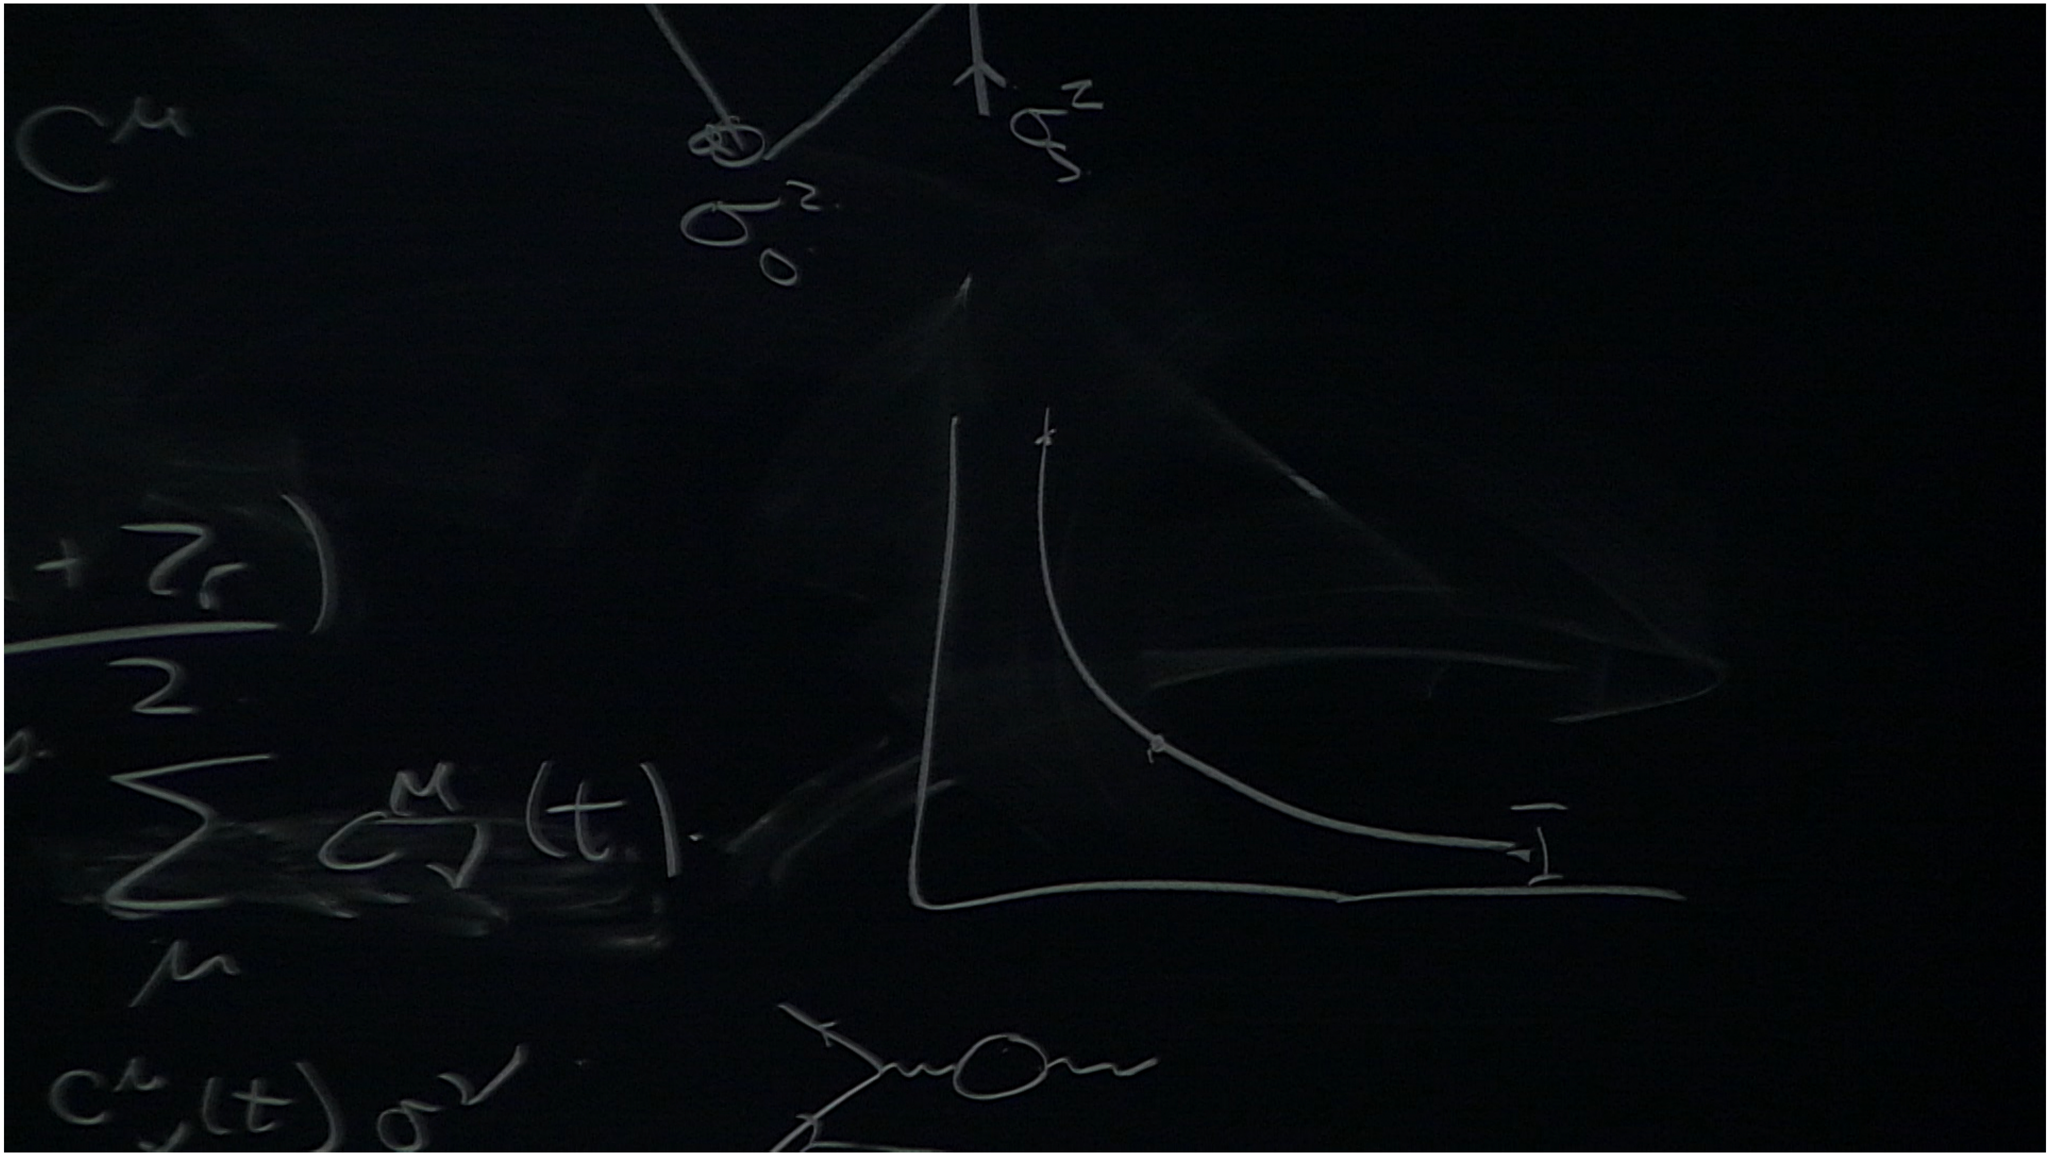

Joel_Wallman_APS_2018.pdf

Page: 6 / 24 Find:

Estimating decay rates

Two parameters, so can solve for two values of m:

$$p^{m_2-m_1} = \frac{2\mathrm{Pr}(\mathrm{z}|\mathrm{m}_2)-1}{2\mathrm{Pr}(\mathrm{z}|\mathrm{m}_1)-1}$$

Choose m_1=4 to avoid gate-dependent perturbations

Experimentally, have estimates $\hat{\mathrm{Pr}}(z|m_i) = \mathrm{Pr}(\mathrm{z}|\mathrm{m_i}) + \epsilon_\mathrm{i}$

Conservative uncertainty on estimate provided $\epsilon_2 \ll \hat{p}^{m_2-m_1}$ is

$$p \approx \hat{p} + \frac{2\epsilon_1+2\epsilon_2}{m_2-m_1}$$

Choosing $m_2 \approx 1/r$ where r = 1-f gives multiplicative precision:

$$\hat{r} \approx r + 2r(\epsilon_1 + \epsilon_2)$$

JJW, Quantum 2, 47 (2018), Helsen *et al*, arXiv:1701.04299

Fitting with two points?

Pros:

- Under the decay model, two sequence lengths provides the minimal uncertainty for fixed resources
- The estimator is straightforward and unambiguous (it is the least-squares, MLE, etc)
- Don't need to weight data points, which naturally assigns low weight to the most informative data points
- Clearly motivated choices of sequence lengths
- Simple error analysis
- Relatively insensitive to the distribution of probabilities over sequences

Cons:

- No model validation! However, with precise estimates of A and p can do hypothesis

Fitting with two points?

Pros:

- Under the decay model, two sequence lengths provides the minimal uncertainty for fixed resources
- The estimator is straightforward and unambiguous (it is the least-squares, MLE, etc)
- Don't need to weight data points, which naturally assigns low weight to the most informative data points
- Clearly motivated choices of sequence lengths
- Simple error analysis
- Relatively insensitive to the distribution of probabilities over sequences

Cons:

- No model validation! However, with precise estimates of A and p can do hypothesis testing on independent sequence lengths (and use Bayesian methods to update posterior if consistent)

What is wrong with standard randomized benchmarking?

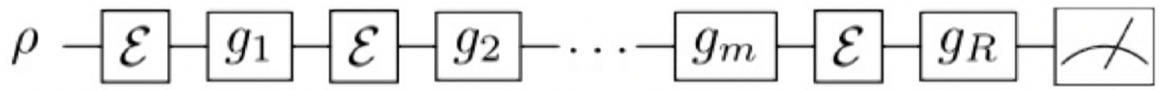

While it is efficient, polynomial factors matter!

A typical n-qubit Clifford gate requires $O(n^2/\log n)$ generating gates with infidelity r.

The fidelity per Clifford is roughly $1 - r\,n^2$, could be anywhere in $[1 - r\,n^4, 1]$

Cannot fit a decay that goes straight to zero!

Also does not apply to a universal gate set!

Cycle benchmarking

Randomized benchmarking requires O(n^2/log n) to twirl noise to depolarizing noise completely characterized by the fidelity.

Solution: perform a weaker twirl

Bonus: allows characterization of universal gate sets

Trade-off: resulting noise is more complicated, described by 4^n Pauli fidelities

$$\mathcal{F}_P(\tilde{\mathcal{H}}, \mathcal{H}) = \mathrm{Tr}\left[\mathcal{H}(P)\tilde{\mathcal{H}}(P)\right]/d$$

The fidelity is essentially the average Pauli fidelity

$$(d+1)\mathcal{F}(\tilde{\mathcal{H}}, \mathcal{H}) = 1 + d\ \mathbb{E}_P \mathcal{F}_P(\tilde{\mathcal{H}}, \mathcal{H})$$

So we can learn a small random set of Pauli fidelities and average to estimate the fidelity. Number of sampled Paulis needed for multiplicative precision is independent of r because

$$|\mathcal{F}(\tilde{\mathcal{H}}, \mathcal{H}) - \mathcal{F}_P(\tilde{\mathcal{H}}, \mathcal{H})| \leq r$$

Estimating Pauli fidelities

For a fixed P, we estimate $\mathcal{F}_P(\tilde{\mathcal{H}}, \mathcal{H}) = \mathrm{Tr}\left[\mathcal{H}(P)\tilde{\mathcal{H}}(P)\right]/d$ as follows.

1. Prepare $2^{-n}(I \pm P)$
2. Apply Q, H, $HQ^{-1}H^{-1}$ m times with m independent Q's from a unitary 1-design
3. Measure the expectation value of $\pm\mathcal{H}^m(P)$ (via coarse-graining)
4. Average over twirls and signs and fit to

$$Af(\mathcal{H}(P), \tilde{\mathcal{H}}(P)^m + 1/2$$

To make the above robust to state-preparation and measurement errors, we:
1. add a random Q' that commutes with P to the first step; and
2. choose values of m that give the same H^m

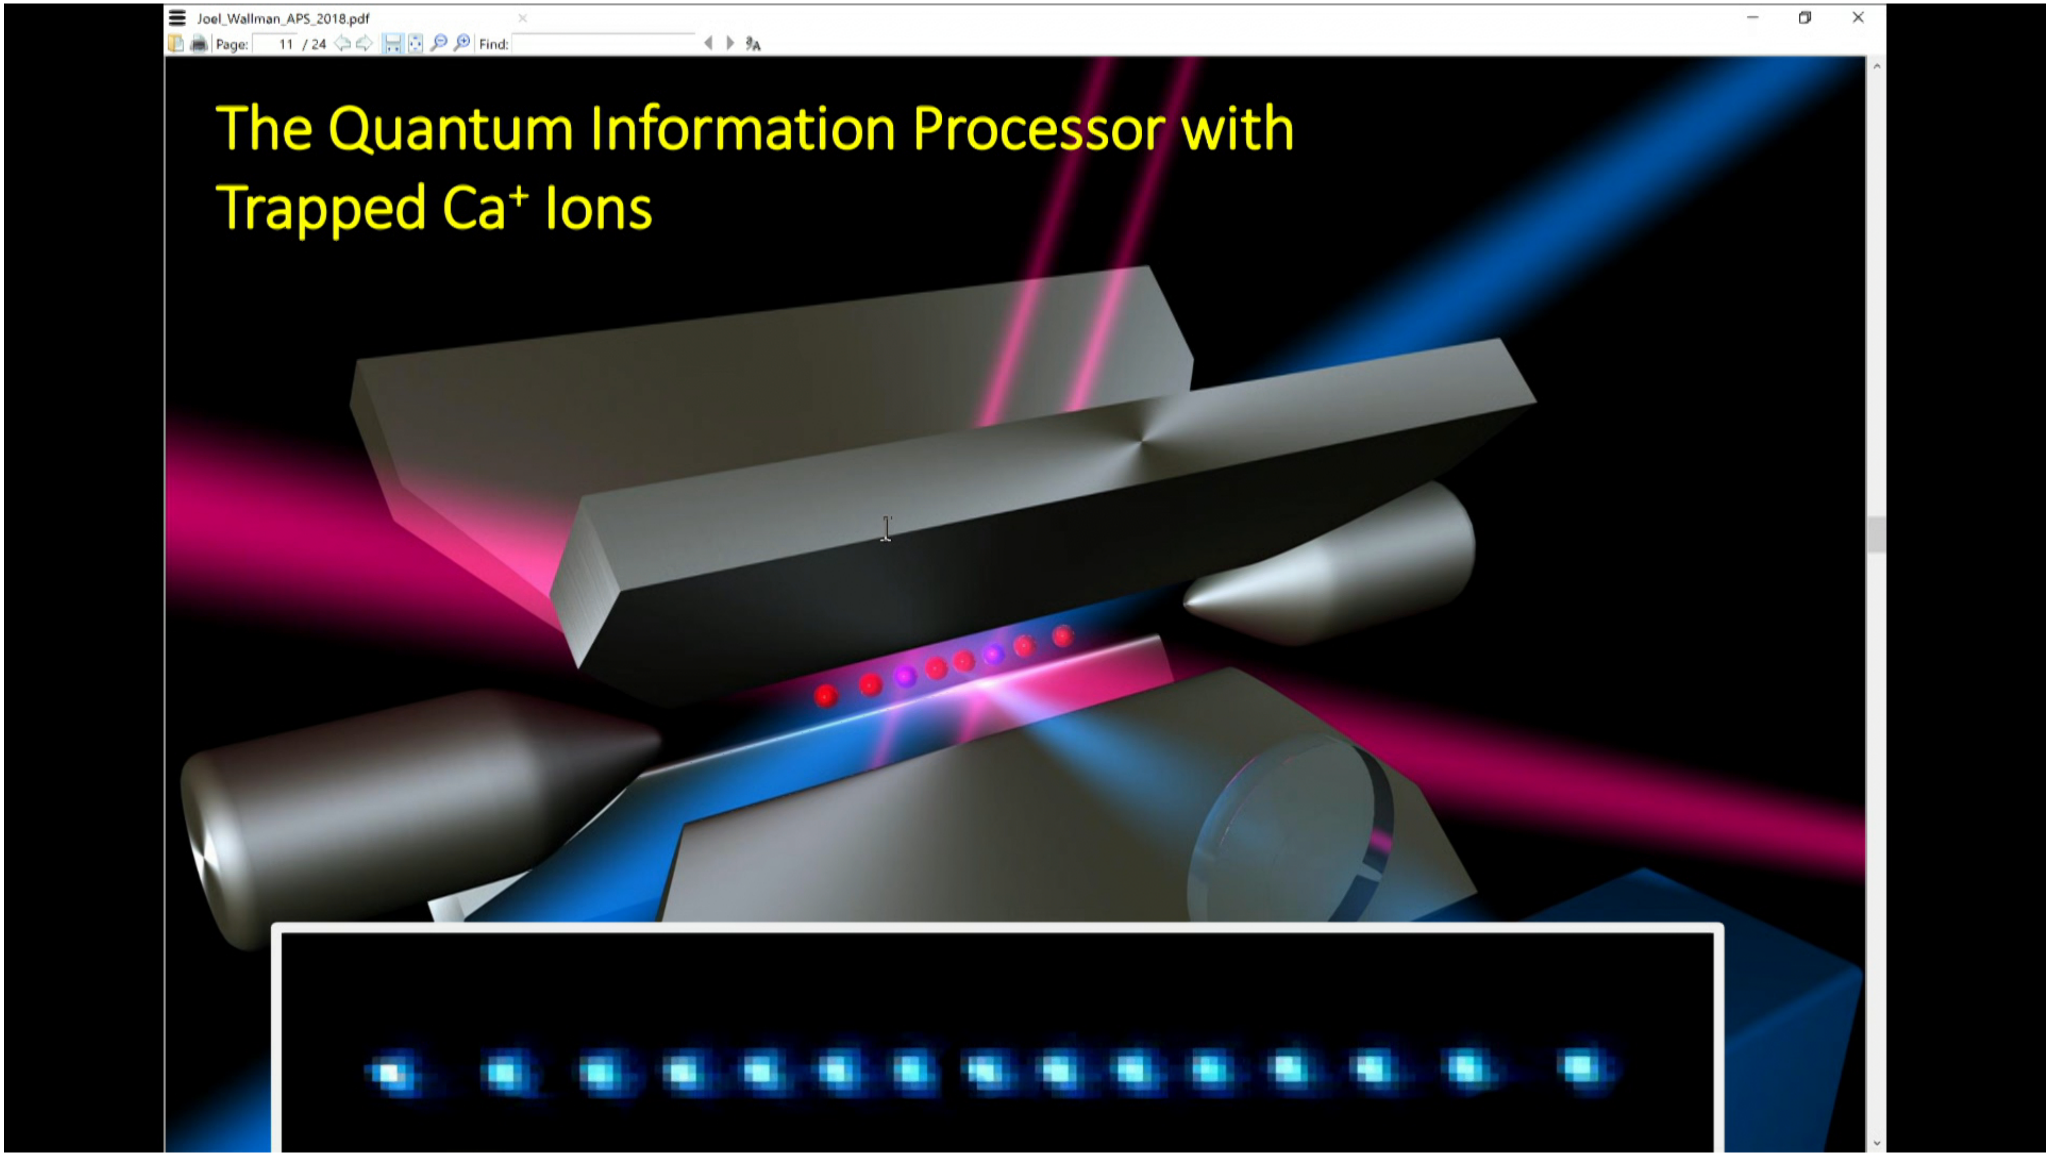
Joel_Wallman_APS_2018.pdf
Page: 11 / 24
Find:
The Quantum Information Processor with
Trapped Ca⁺ Ions

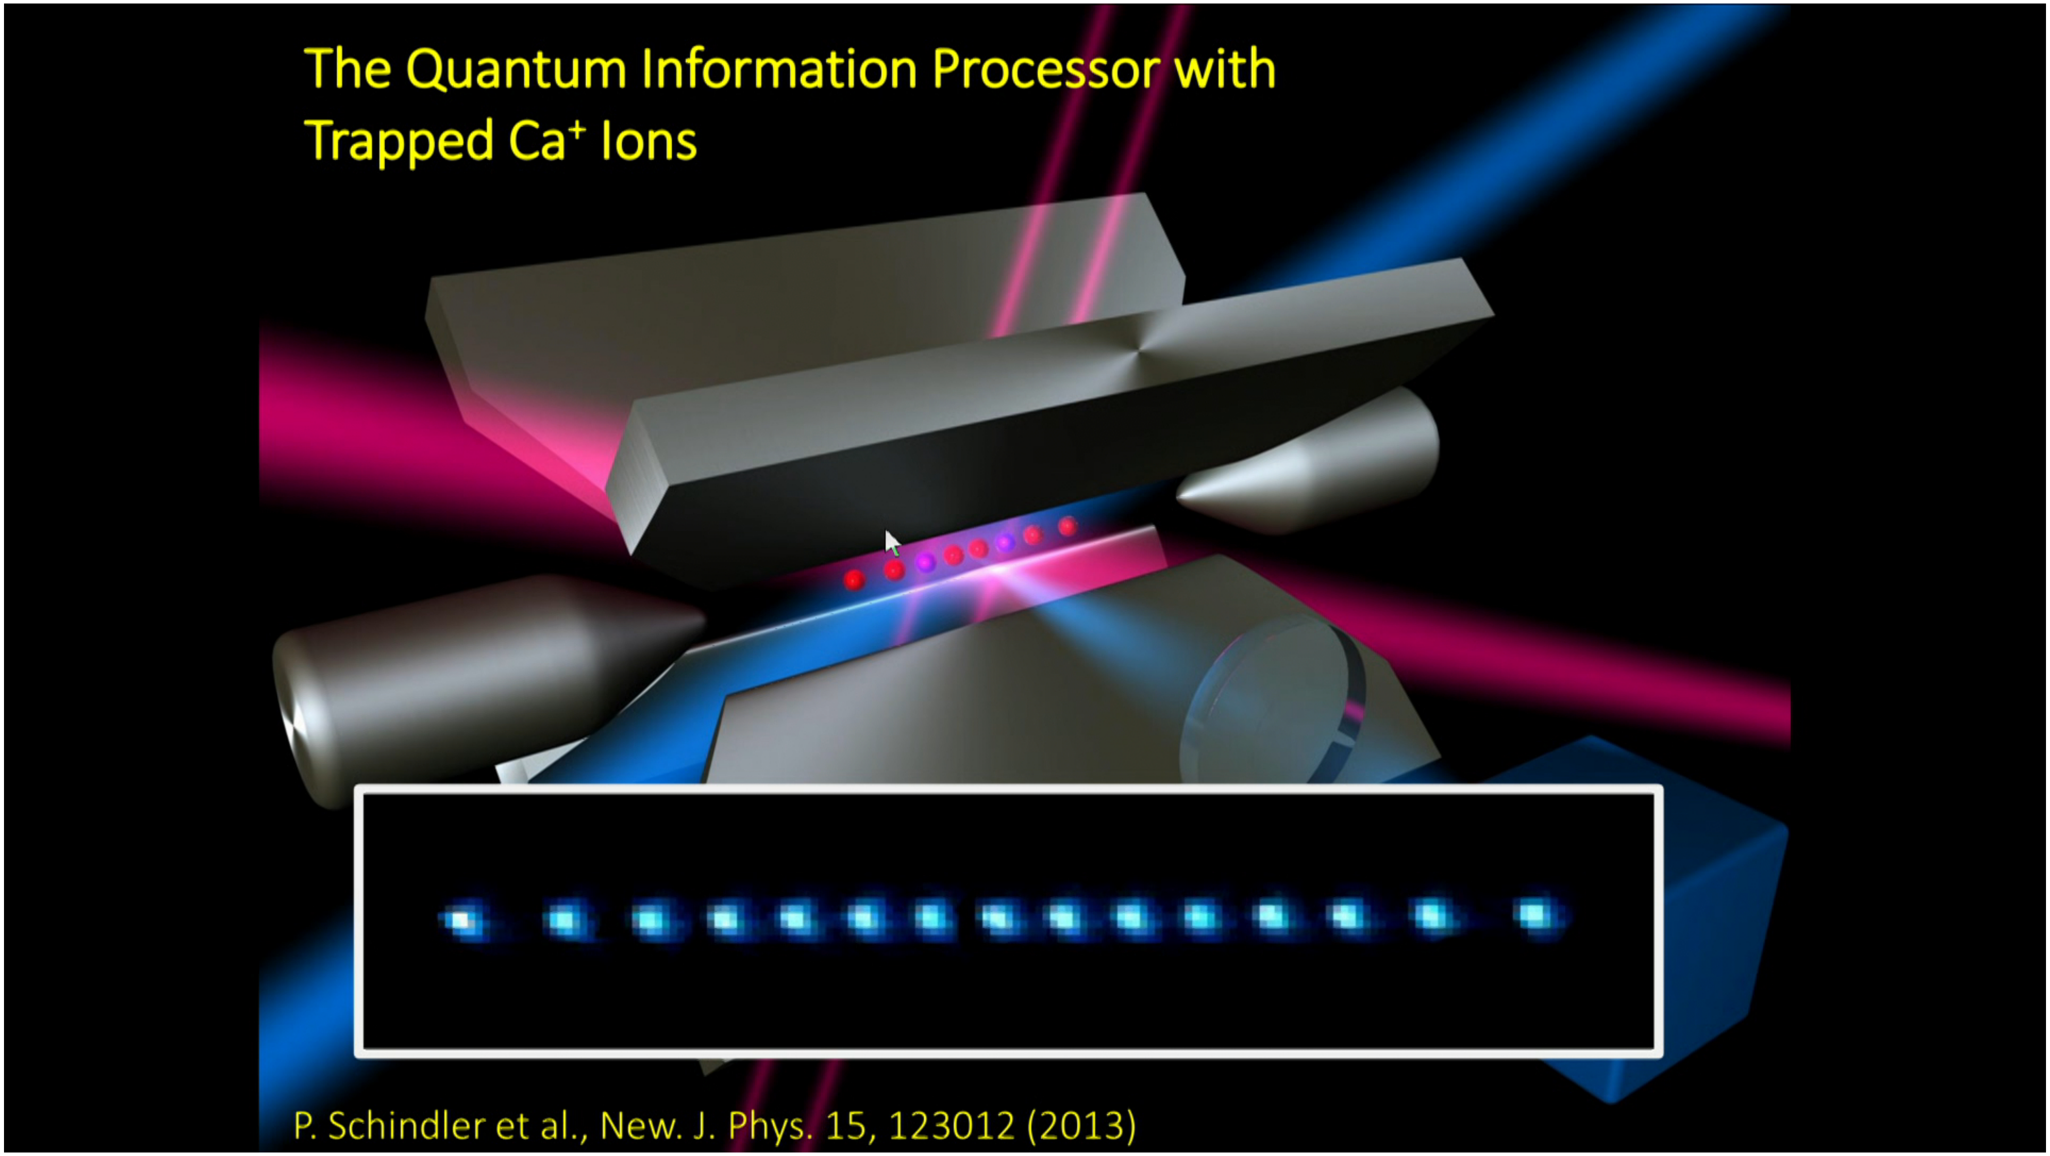
The Quantum Information Processor with
Trapped Ca⁺ Ions
P. Schindler et al., New. J. Phys. 15, 123012 (2013)

Level Scheme of $^{40}Ca^{+}$

qubit on narrow S - D quadrupole transition

$\tau \approx 1s$

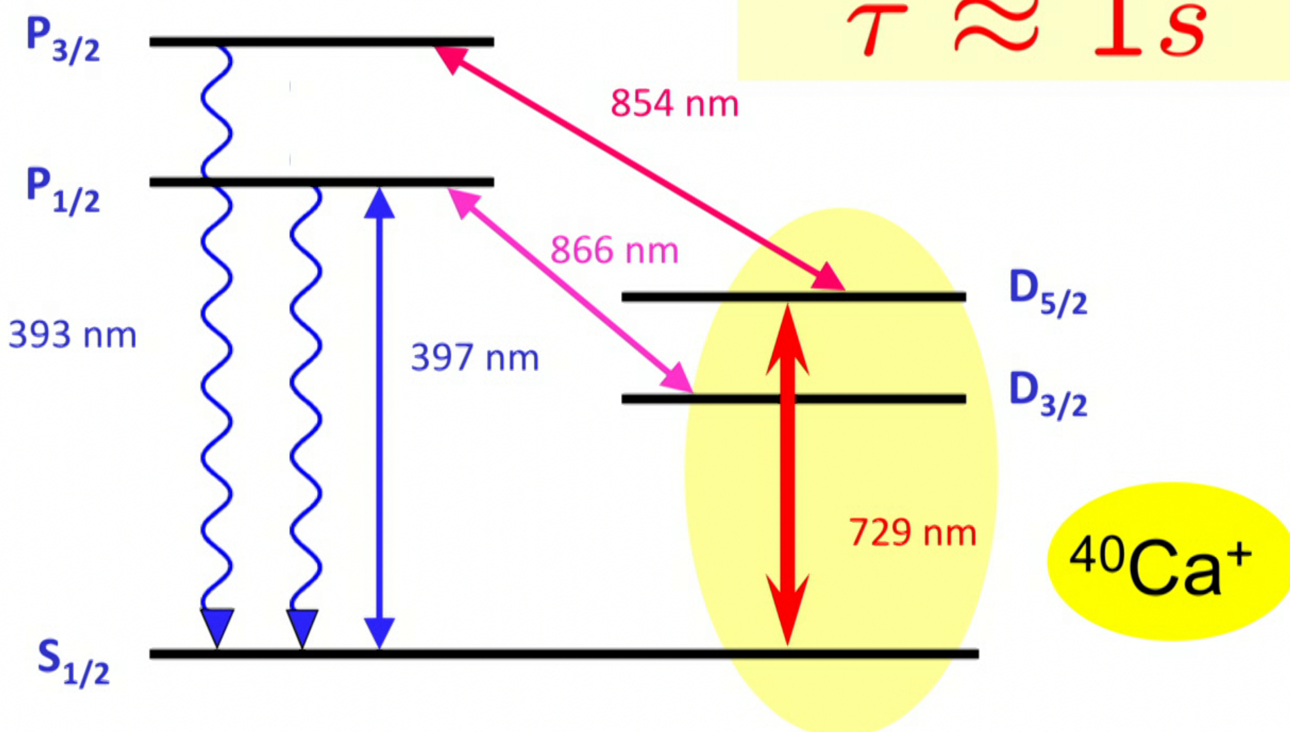

Quantum computing with global and local operations

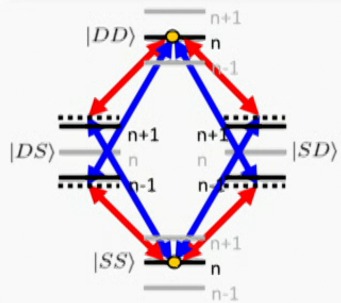

$S^2_{x,y}(\theta)$

Bichromatic excitation: entangling operations

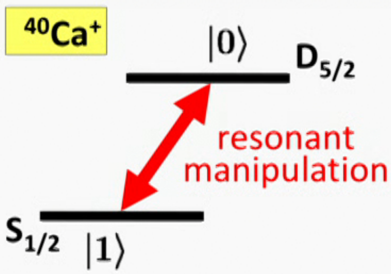

$S_{x,y}(\theta)$

Resonant excitation: collective local operations

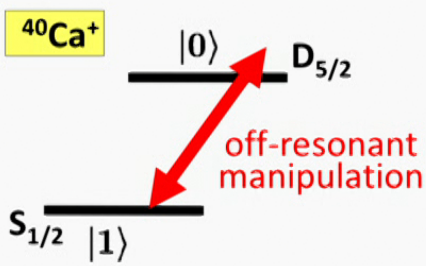

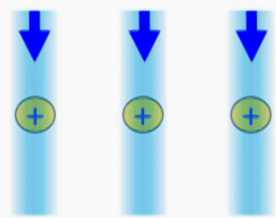

off-resonant excitation:
individual local operations
(AC Stark shifts)

$\sigma_z^{(i)}(\theta)$

Quantum gate operations – unitaries

Quantum circuits:

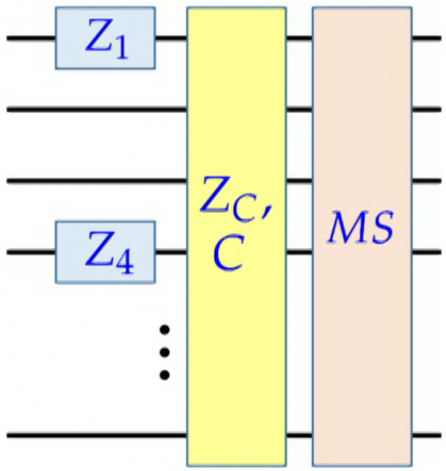

$$Z_n(\theta, j) = e^{-i\theta\sigma_z^n}$$

$\sigma_z^{(n)}(\theta)$
local Stark shifts

$$Z_C(\theta) = e^{-i\theta\sum_i \sigma_z^i}$$

$S_z(\theta)$
collective Stark shifts

$$C(\theta, \phi) = e^{-i\theta\sum_i \sigma_\phi^i}$$

$S_\phi(\theta)$
collective local ops.

$$MS(\theta, \phi) = e^{-i\theta\sum_{i<j} \sigma_\phi^i \sigma_\phi^j}$$

$S_\phi^2(\theta)$
entangling MS ops.

$$\sigma_\phi = \cos\phi\sigma_x + \sin\phi\sigma_y$$

σ_k^j *k*-th Pauli matrix acting on *j*-th qubit

additionally available:

- hiding operations (reduce comp. subspace)
- dephasing operations (open systems)
- initialization/reset operation
- quantum (cache) memory

System Capabilities at a glance: Qubits & Gates

Parameter	Description	Value
Number of qubits (ions)	$^{40}Ca^{+}$ |S>, |D> states	up to 14
Qubit preparation fidelity & time	Sideband cool ~ 2ms Raman cool ~ 400µs	99.5% fidelity $<n> \sim 1$
Coherence times	T_1 T_2	1.1 sec 30 ms to 0.55 sec
Measurement time	PMT CCD	400 µs 2 ms
Gate time	$R(\pi,\phi)$ $Z(\pi)$ $MS(\pi/2,\phi)$	10 µs 10 µs 40-100 µs

The Mølmer-Sørensen gate

Primarily interested in the n-qubit Mølmer-Sørensen gate

$$\mathrm{MS} = \exp\left(-i\pi S_x^2/8\right)$$

$$S_x = \sum_j X_j$$

For even n,

$$\mathrm{MS} = \tfrac{1-i}{2} I + (-1)^{n/2} \tfrac{1+i}{2} X^{\otimes n}$$

For any Pauli P that does not commute with the MS gate,

$$\mathrm{MS} P \mathrm{MS}^\dagger = X^{\otimes n} P$$

Twirling set can be independent Paulis + rotations about the X axis
(isomorphic to the dihedral group)

Compiling independent single-qubit gates

Twirling set can be independent Paulis + rotations about the X axis.

Abstractly prefer including $\pi/4$ X rotations as non-Pauli gates systematically convolves Pauli errors between time steps, larger twirl averages such errors at each time step.

However, such Pauli scrambling makes a small contribution to the average error rate.

Compiling independent single-qubit gates

A bigger constraint: primitive gates contain only collective rotations and single-qubit Z rotations.

Such rotations are sufficient to generate a cycle of arbitrary independent single-qubit gates, e.g., $Z(\alpha)X\left(\frac{\pi}{2}\right)Z(\beta)X\left(\frac{\pi}{2}\right)Z(\gamma)$.

However, Z gates are the noisiest gates, want to minimize how many are used. Much fewer gates required to achieve independent Pauli gates.

First pass:

1. look for the most common Pauli P
2. Choose 4 collective $X\left(\frac{\pi}{2}\right)$ /$Y\left(\frac{\pi}{2}\right)$ rotations that multiply to P (6 for the identity)
3. Insert Z(π) on qubits in-between collective pulses to obtain arbitrary Paulis.

The experimental implementation

- Mølmer-Sørensen gate is of order 4, so choose sequence lengths to be 4 and the largest multiple of 4 with a decent signal
- Sample Pauli fidelities exhaustively for 2 and 4 qubits, verify the number of Pauli fidelities required for an accurate estimate at 4 qubits (around 50) for larger numbers of qubits

Results

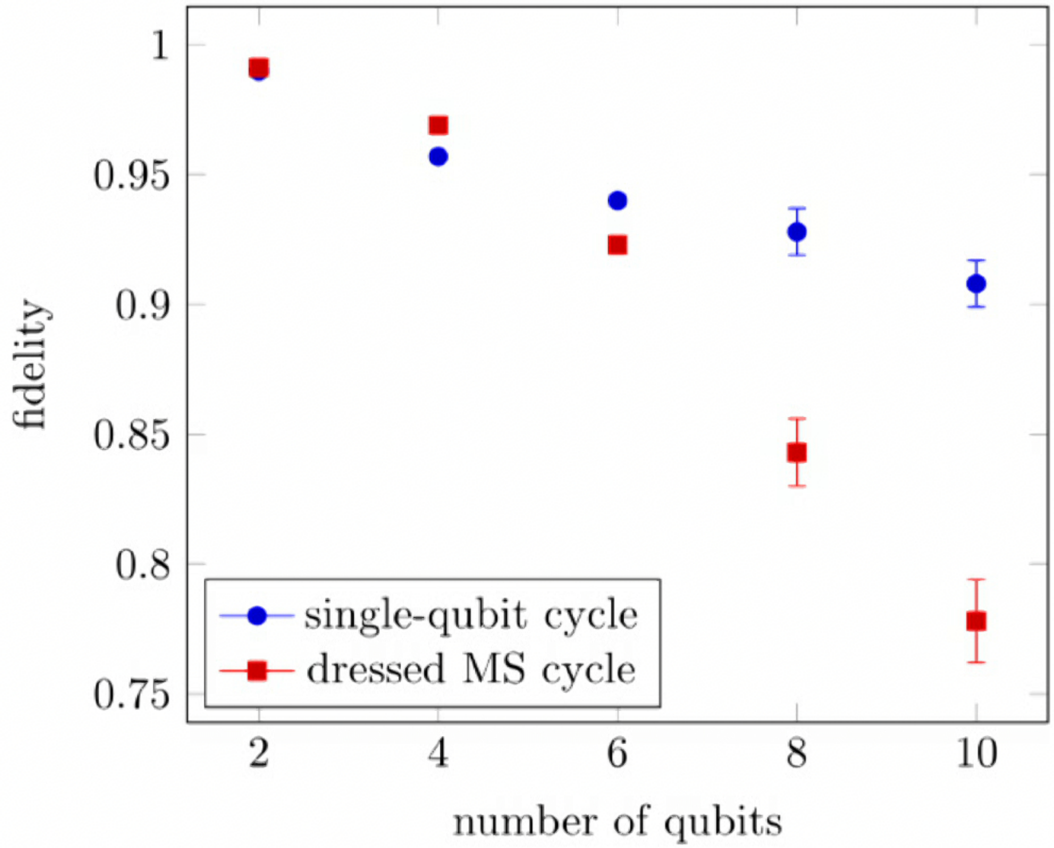

Approximately 1.75 Z pulses per qubit, reference decay is consistent with a fidelity of 0.995 including any crosstalk

Uncertainties obtained from bootstrap percentiles

Results

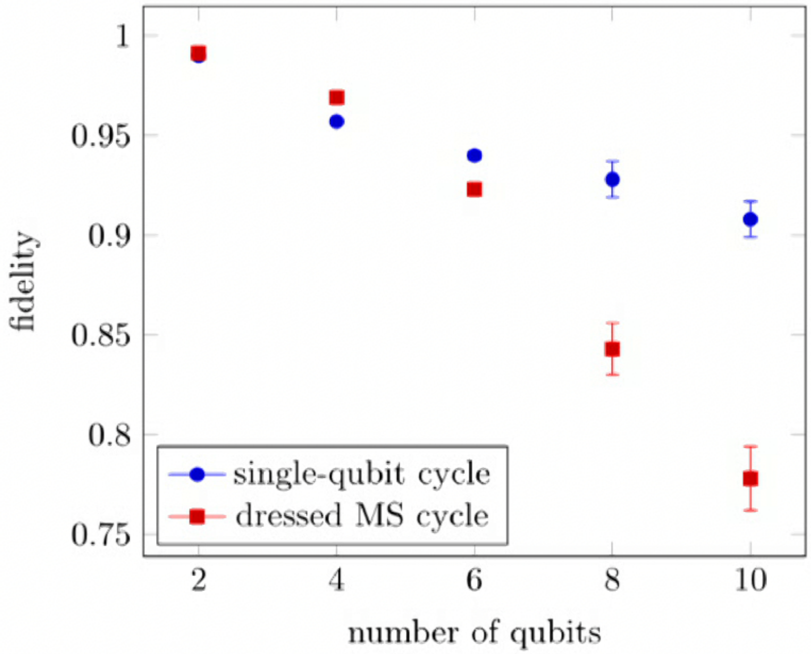

A standard (bad!) estimate of the fidelity of the interleaved gate is the difference between the interleaved and non-interleaved fidelity.

The fidelity is greater with an interleaved gate for n = 2, 4.

Possible explanations:

1) Coherent errors cancelling between MS gate and twirling gates
2) Context-dependent noise on the twirling gates

Results

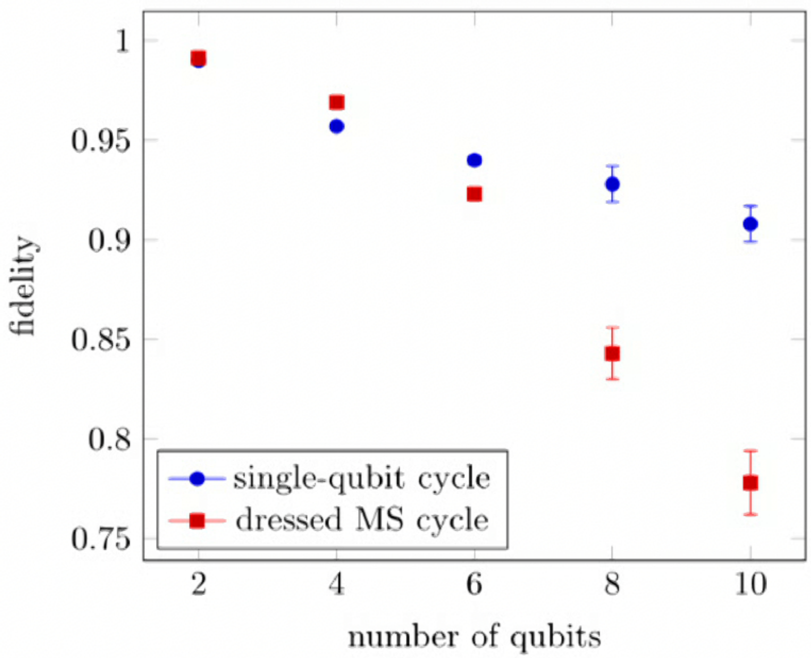

A standard (bad!) estimate of the fidelity of the interleaved gate is the difference between the interleaved and non-interleaved fidelity.

The fidelity is greater with an interleaved gate for n = 2, 4.

Possible explanations:

1) Coherent errors cancelling between MS gate and twirling gates
2) Context-dependent noise on the twirling gates

Detecting coherence (preliminary)

Want to determine whether the noise is coherent.

Apply the MS gate *m* times between twirling rounds and see how the fidelity decays as a function of *m* [Sheldon *et al,* PRA 93, 012301 (2016)]

"Quadratic" decay => coherent errors, linear decay => stochastic errors

A new laser arrived that should reduce intensity fluctuations… so we have worse data

MS^m	Fidelity
1	0.974(2)
5	0.937(4)
9	0.91(1)

Summary

- Presented a (relatively) complete analysis of estimates and errors from randomized benchmarking (and variants thereof)
- Standard randomized benchmarking, while efficient in principle, is impractical for larger numbers of qubits
- Developed cycle benchmarking, which is in-principle practical and directly outputs the performance under randomized compiling without further assumptions [JJW and Emerson, PRA 94, 052325 (2016)]
- Shown cycle benchmarking is practical for many qubits (2 hours for 10 qubit data)
- Extracting an accurate fidelity on an individual gate from cycle benchmarking takes more work, waiting for the laser to be set up
- Implementing the protocol for different fractions of the MS gate will show scalable characterization of non-Clifford gates

Thank you !