ENTANGLEMENT AND COMPLEXITY OF INTERACTING QUBITS SUBJECT TO ASYMMETRIC NOISE

Eliot Kapit

Colorado School of Mines

OUTLINE

- **Hard sampling problems and noise**
- **Proposed experiment**
- **Benchmarks:** Volume entanglement, IPR, distance from P-T, output heaviness
- **Difficulty and fidelity estimates**
- **Conclusions**

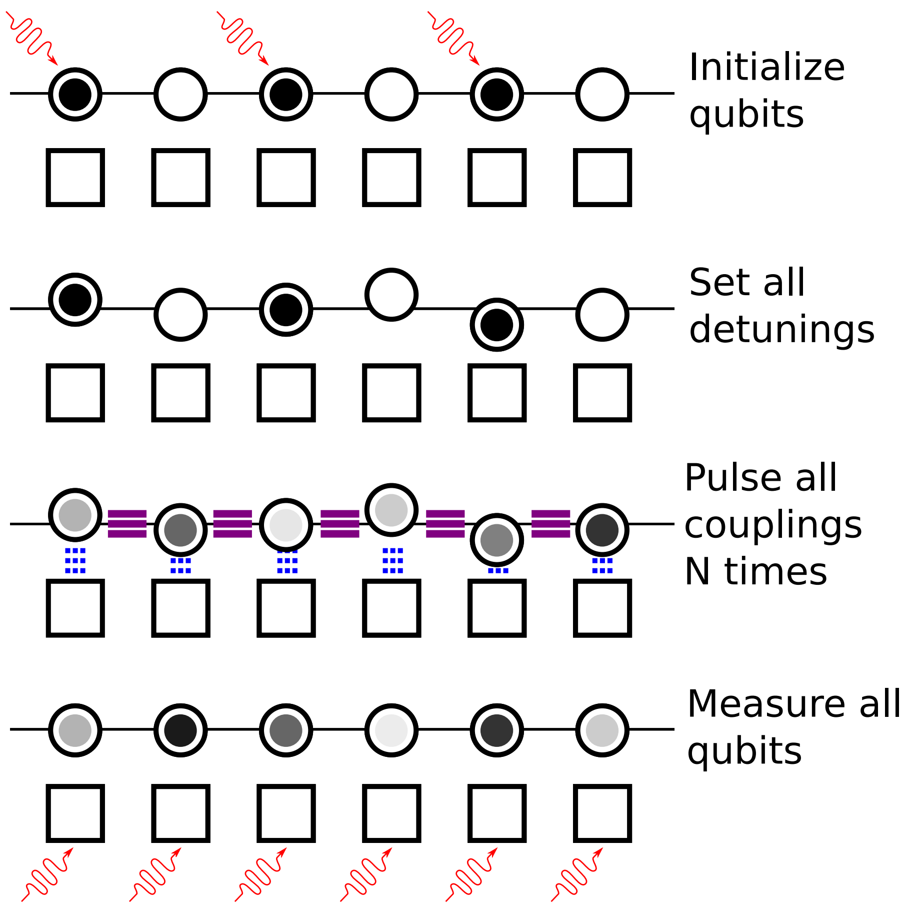

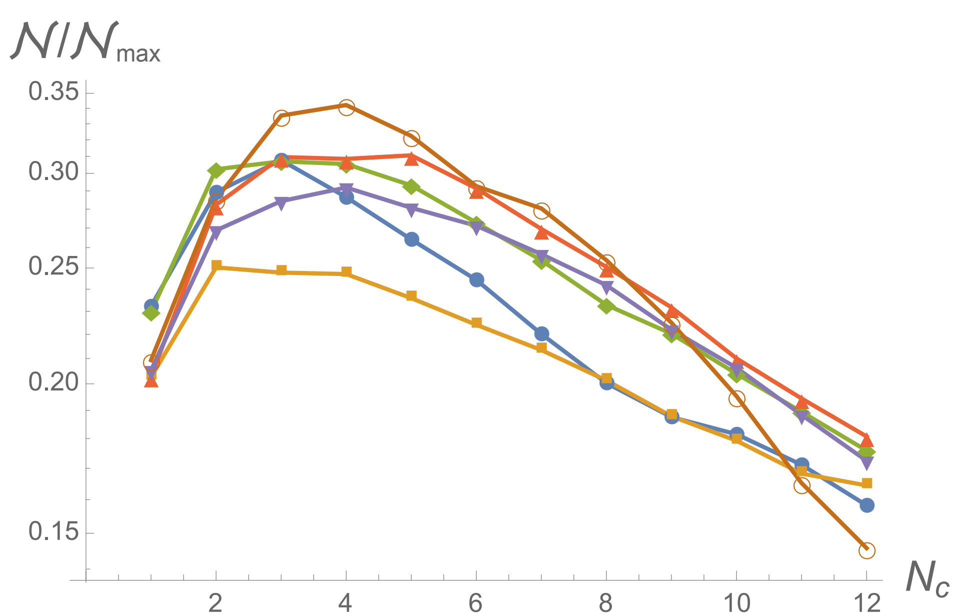

QUANTUM SAMPLING PROBLEMS

- Quantum sampling problems present the best near-term chance of demonstrating quantum "supremacy" in real, noisy hardware
    - Boson Sampling, IQP circuits, QAOA, random quantum circuits, Bose-Hubbard chains, etc.
- Goal: sample the output of a system evolved from a simple initial state through a quantum entangling process. Exponentially hard for classical machines!
- But: rapid progress in simulation algorithms & classical hardware make this a fast-moving target
- Guiding principle: ***maximize*** *simulation complexity and* ***minimize*** *quantum hardware complexity*

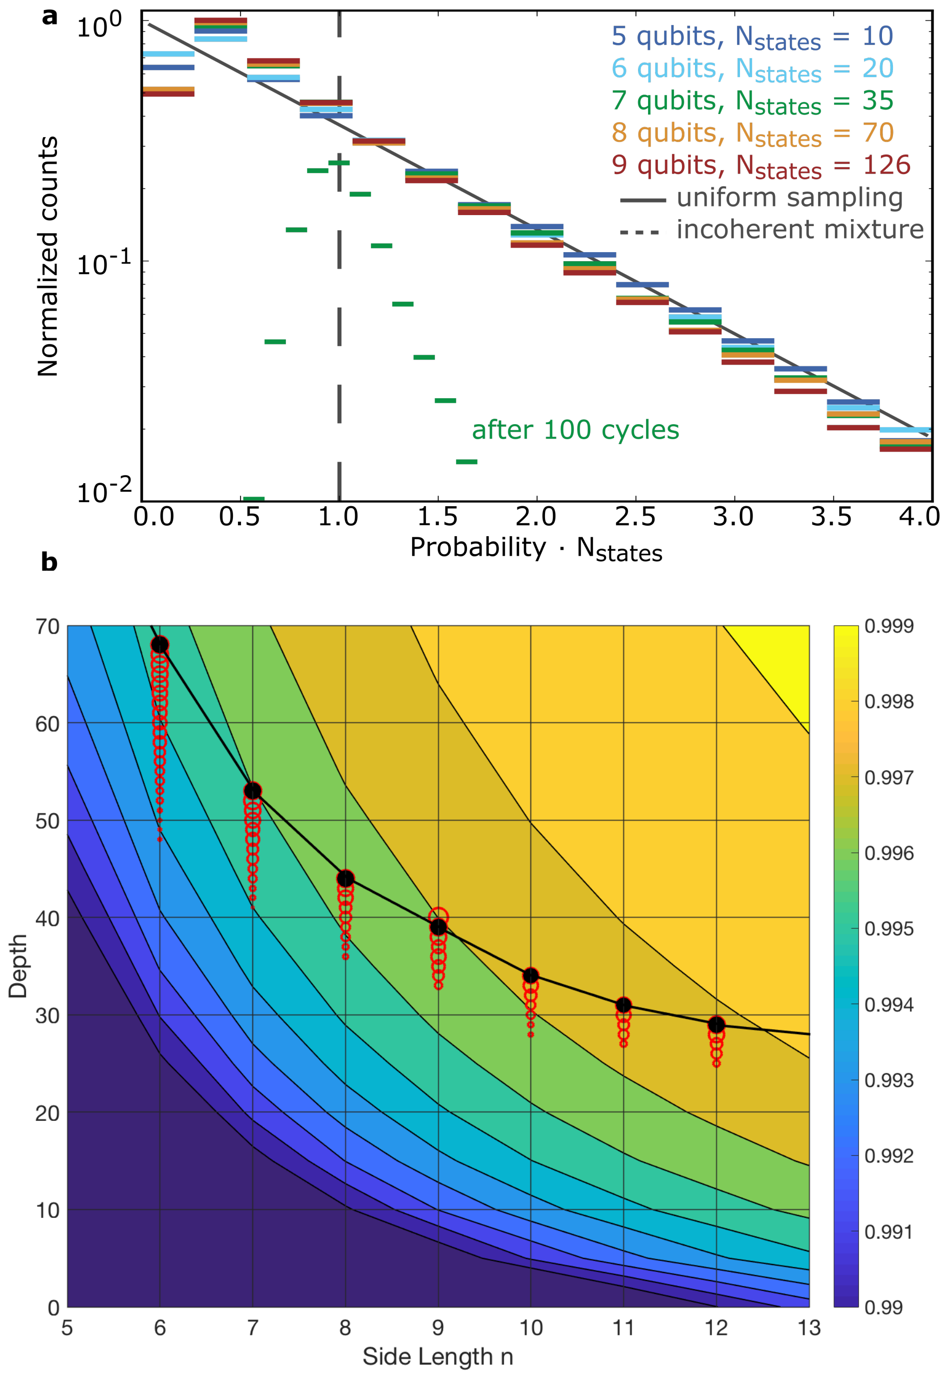

Image credits: Neill et al 2018, Chen et al 2018

QUANTUM SAMPLING PROBLEMS

- Continuous time evolution is **harder** to simulate than gates/optical circuits

- Experimental demonstration of c.t. sampling problem (9 qubit "gmon" chain): *Neill et al, Science 2018.* (see talk section A42, yesterday) System evolves under:

$$H_Q(t) = -g(t)\sum_{i=1}^{L-1}\left[a_i^\dagger a_{i+1} + \mathrm{H.c.}\right] + \sum_{i=1}^{L}\left[h_i a_i^\dagger a_i - \sum_{n=2}^{n_{max}} \delta_n \left|n_i\right\rangle\left\langle n_i\right|\right].$$

- Partitioning/tensor contraction methods not generally applicable to evolution under continuously varying H.

- Entanglement-scaling schemes (MPS, PEPS, etc) fail in volume entangled limit

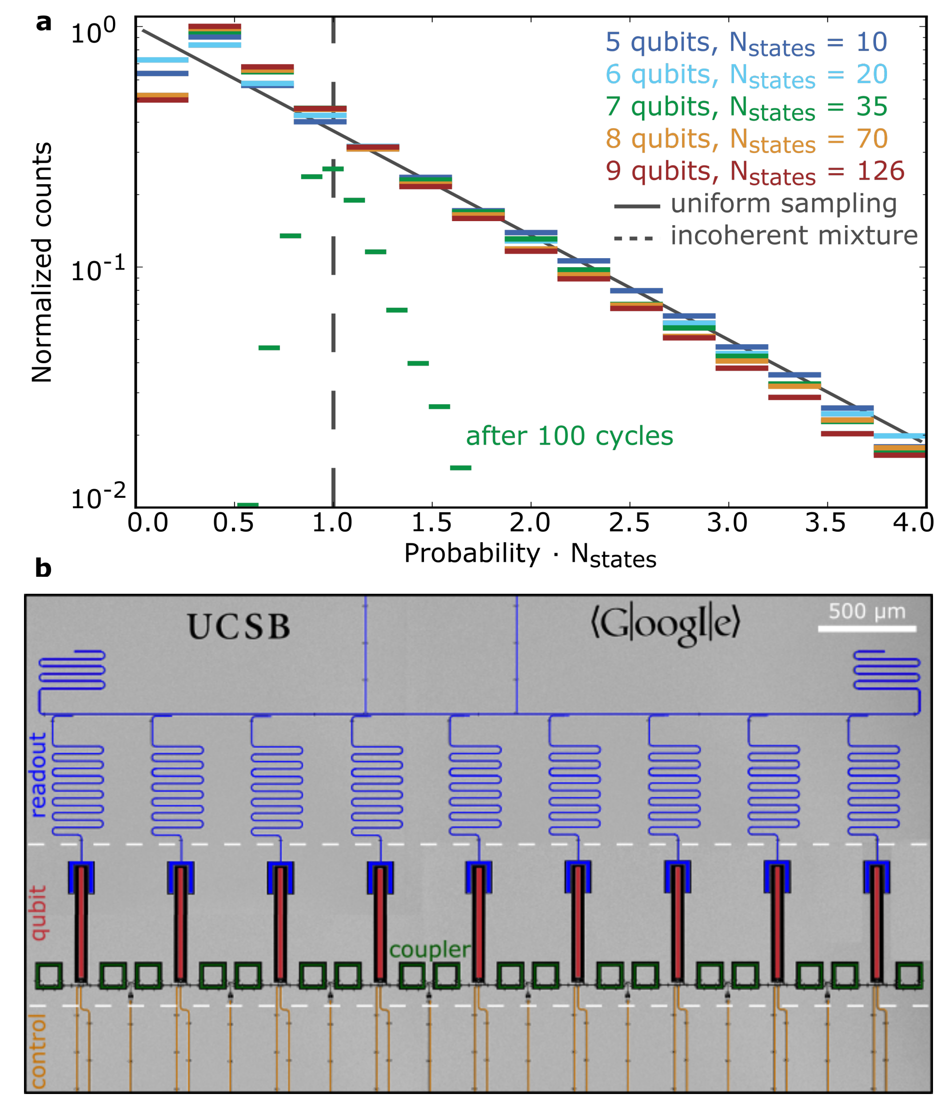

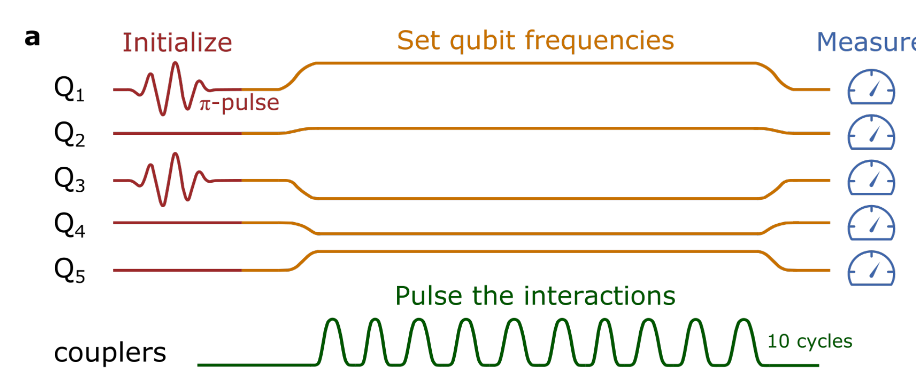

Image credits: Neill et al 2018

BOSE-HUBBARD (GMON) CHAIN

- gmon chain experiment: **initialize** product state, **pulse** tunable exchange couplers (or bring tunable qubits w/ fixed couplings in and out of resonance) **N times,** then **measure** state
- Notice: each qubit has a CPW resonator for final state readout. Otherwise these do not participate during evolution.
- **Half of the physical quantum degrees of freedom are left idle!**
- My proposal: drive parametric qubit-cavity couplings during evolution to increase complexity

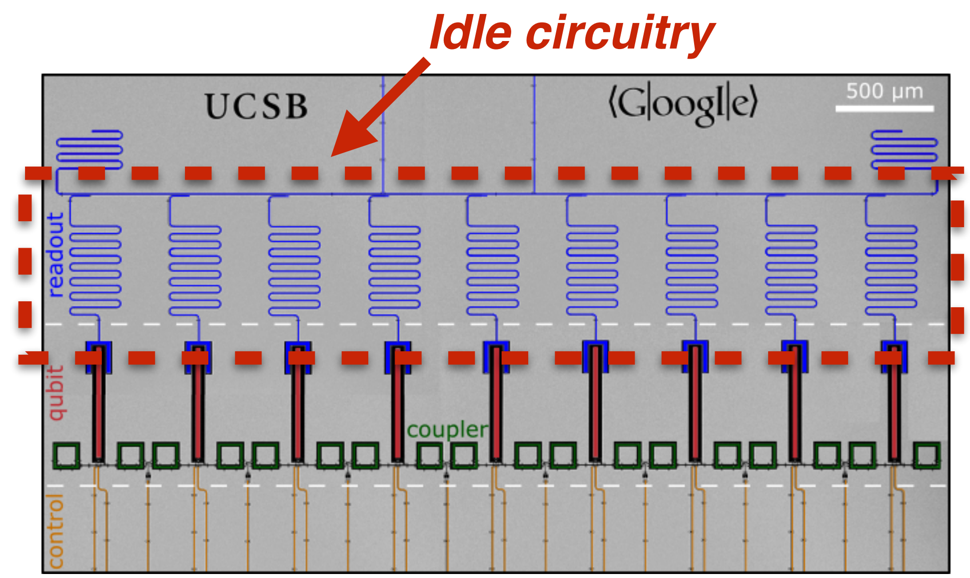

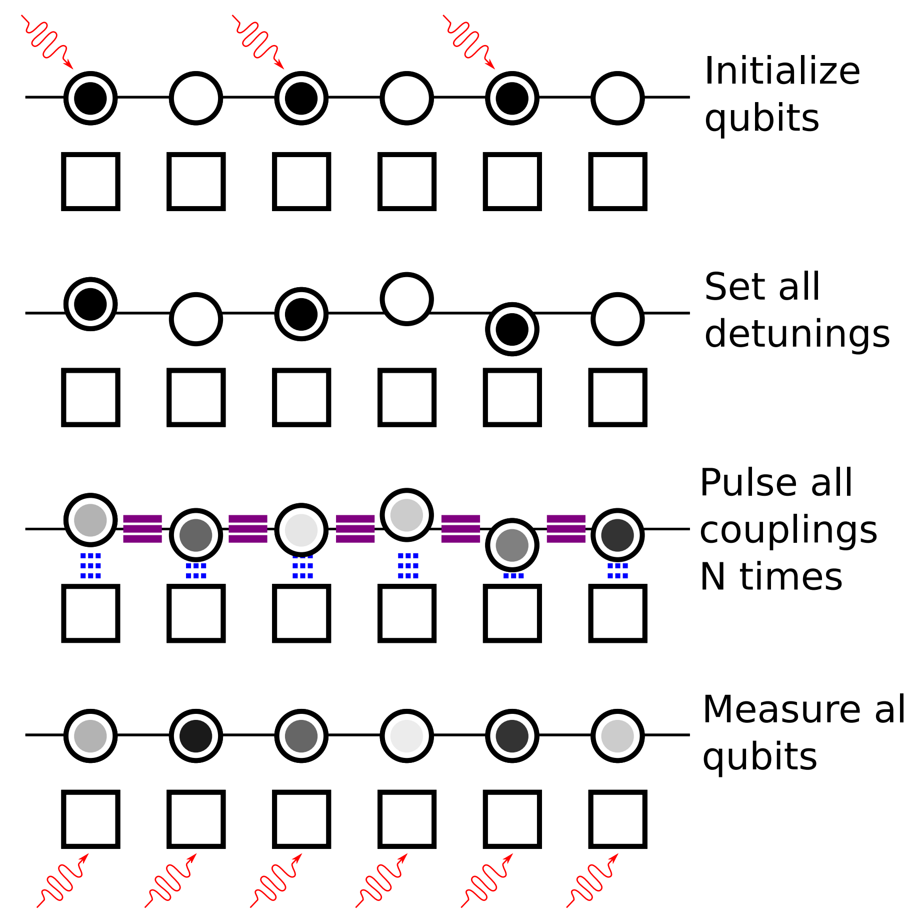

Image credits: Neill et al 2018

PROPOSED EXPERIMENT

- Proposal: simultaneously with the coupler pulses, drive red and/or blue sideband couplings to cavities (*Murch et al PRL 2012, Strand et al PRB 2013, Kapit PRA 2015, Li et al PR App. 2018,* others)

- Total Hamiltonian is:

$$H_Q(t) = -g(t)\sum_{i=1}^{L-1}\left[a_i^\dagger a_{i+1} + \text{H.c.}\right] + \sum_{i=1}^{L}\left[h_i a_i^\dagger a_i - \sum_{n=2}^{n_{max}} \delta_n \left|n_i\right\rangle\left\langle n_i\right|\right].$$

$$H_{QC}(t) = \sum_{i=1}^{L}\left[h_{Ci} a_{Ci}^\dagger a_{Ci} + \Delta a_{Ci}^\dagger a_{Ci} a_i^\dagger a_i\right] + \sum_{i=1}^{L}\left[\Omega_{QCi}^R(t)\, a_{Ci}^\dagger a_i + \Omega_{QCi}^B(t)\, a_{Ci}^\dagger a_i^\dagger + \text{H.c.}\right].$$

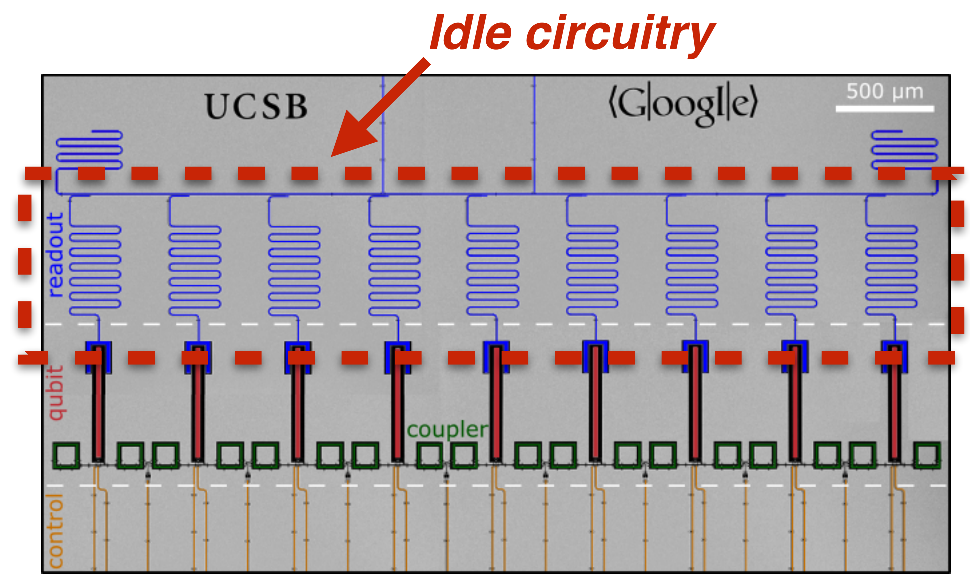

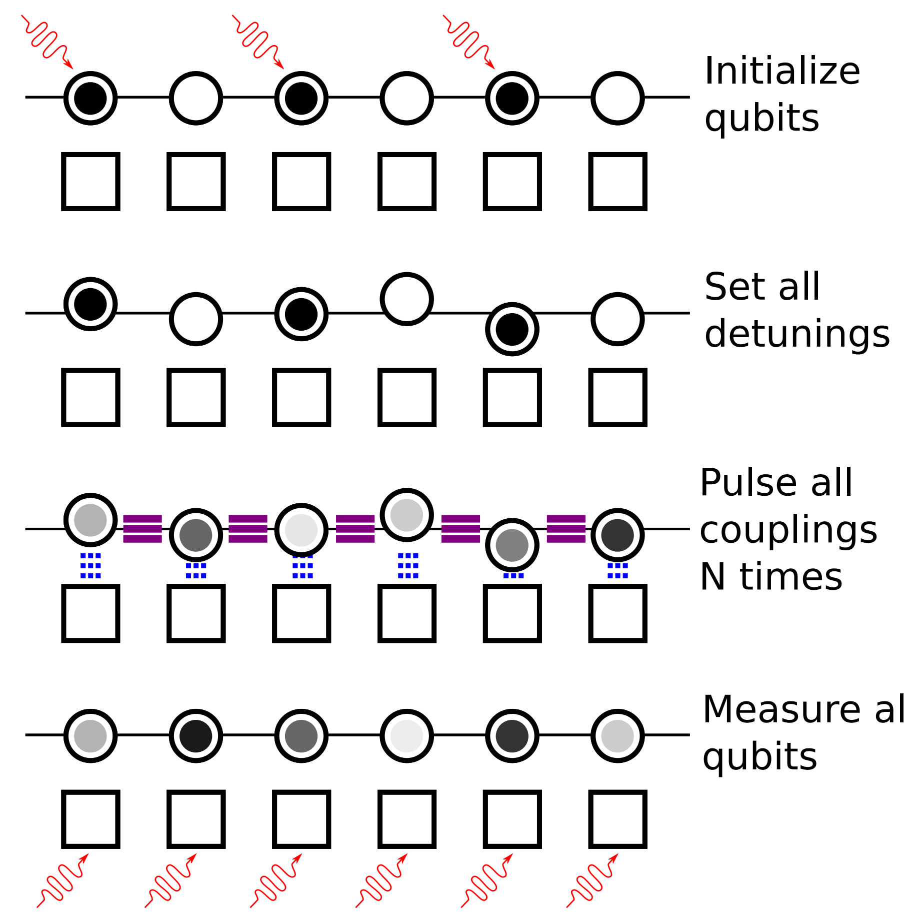

Image credits: Neill et al 2018

PROPOSED EXPERIMENT

- Origin of these terms sketched below; green errors are action of qubit-resonator exchange coupling, red and blue oscillations in qubit energy or resonator driving. **Important: we want these terms to be weak compared to coupler g!**

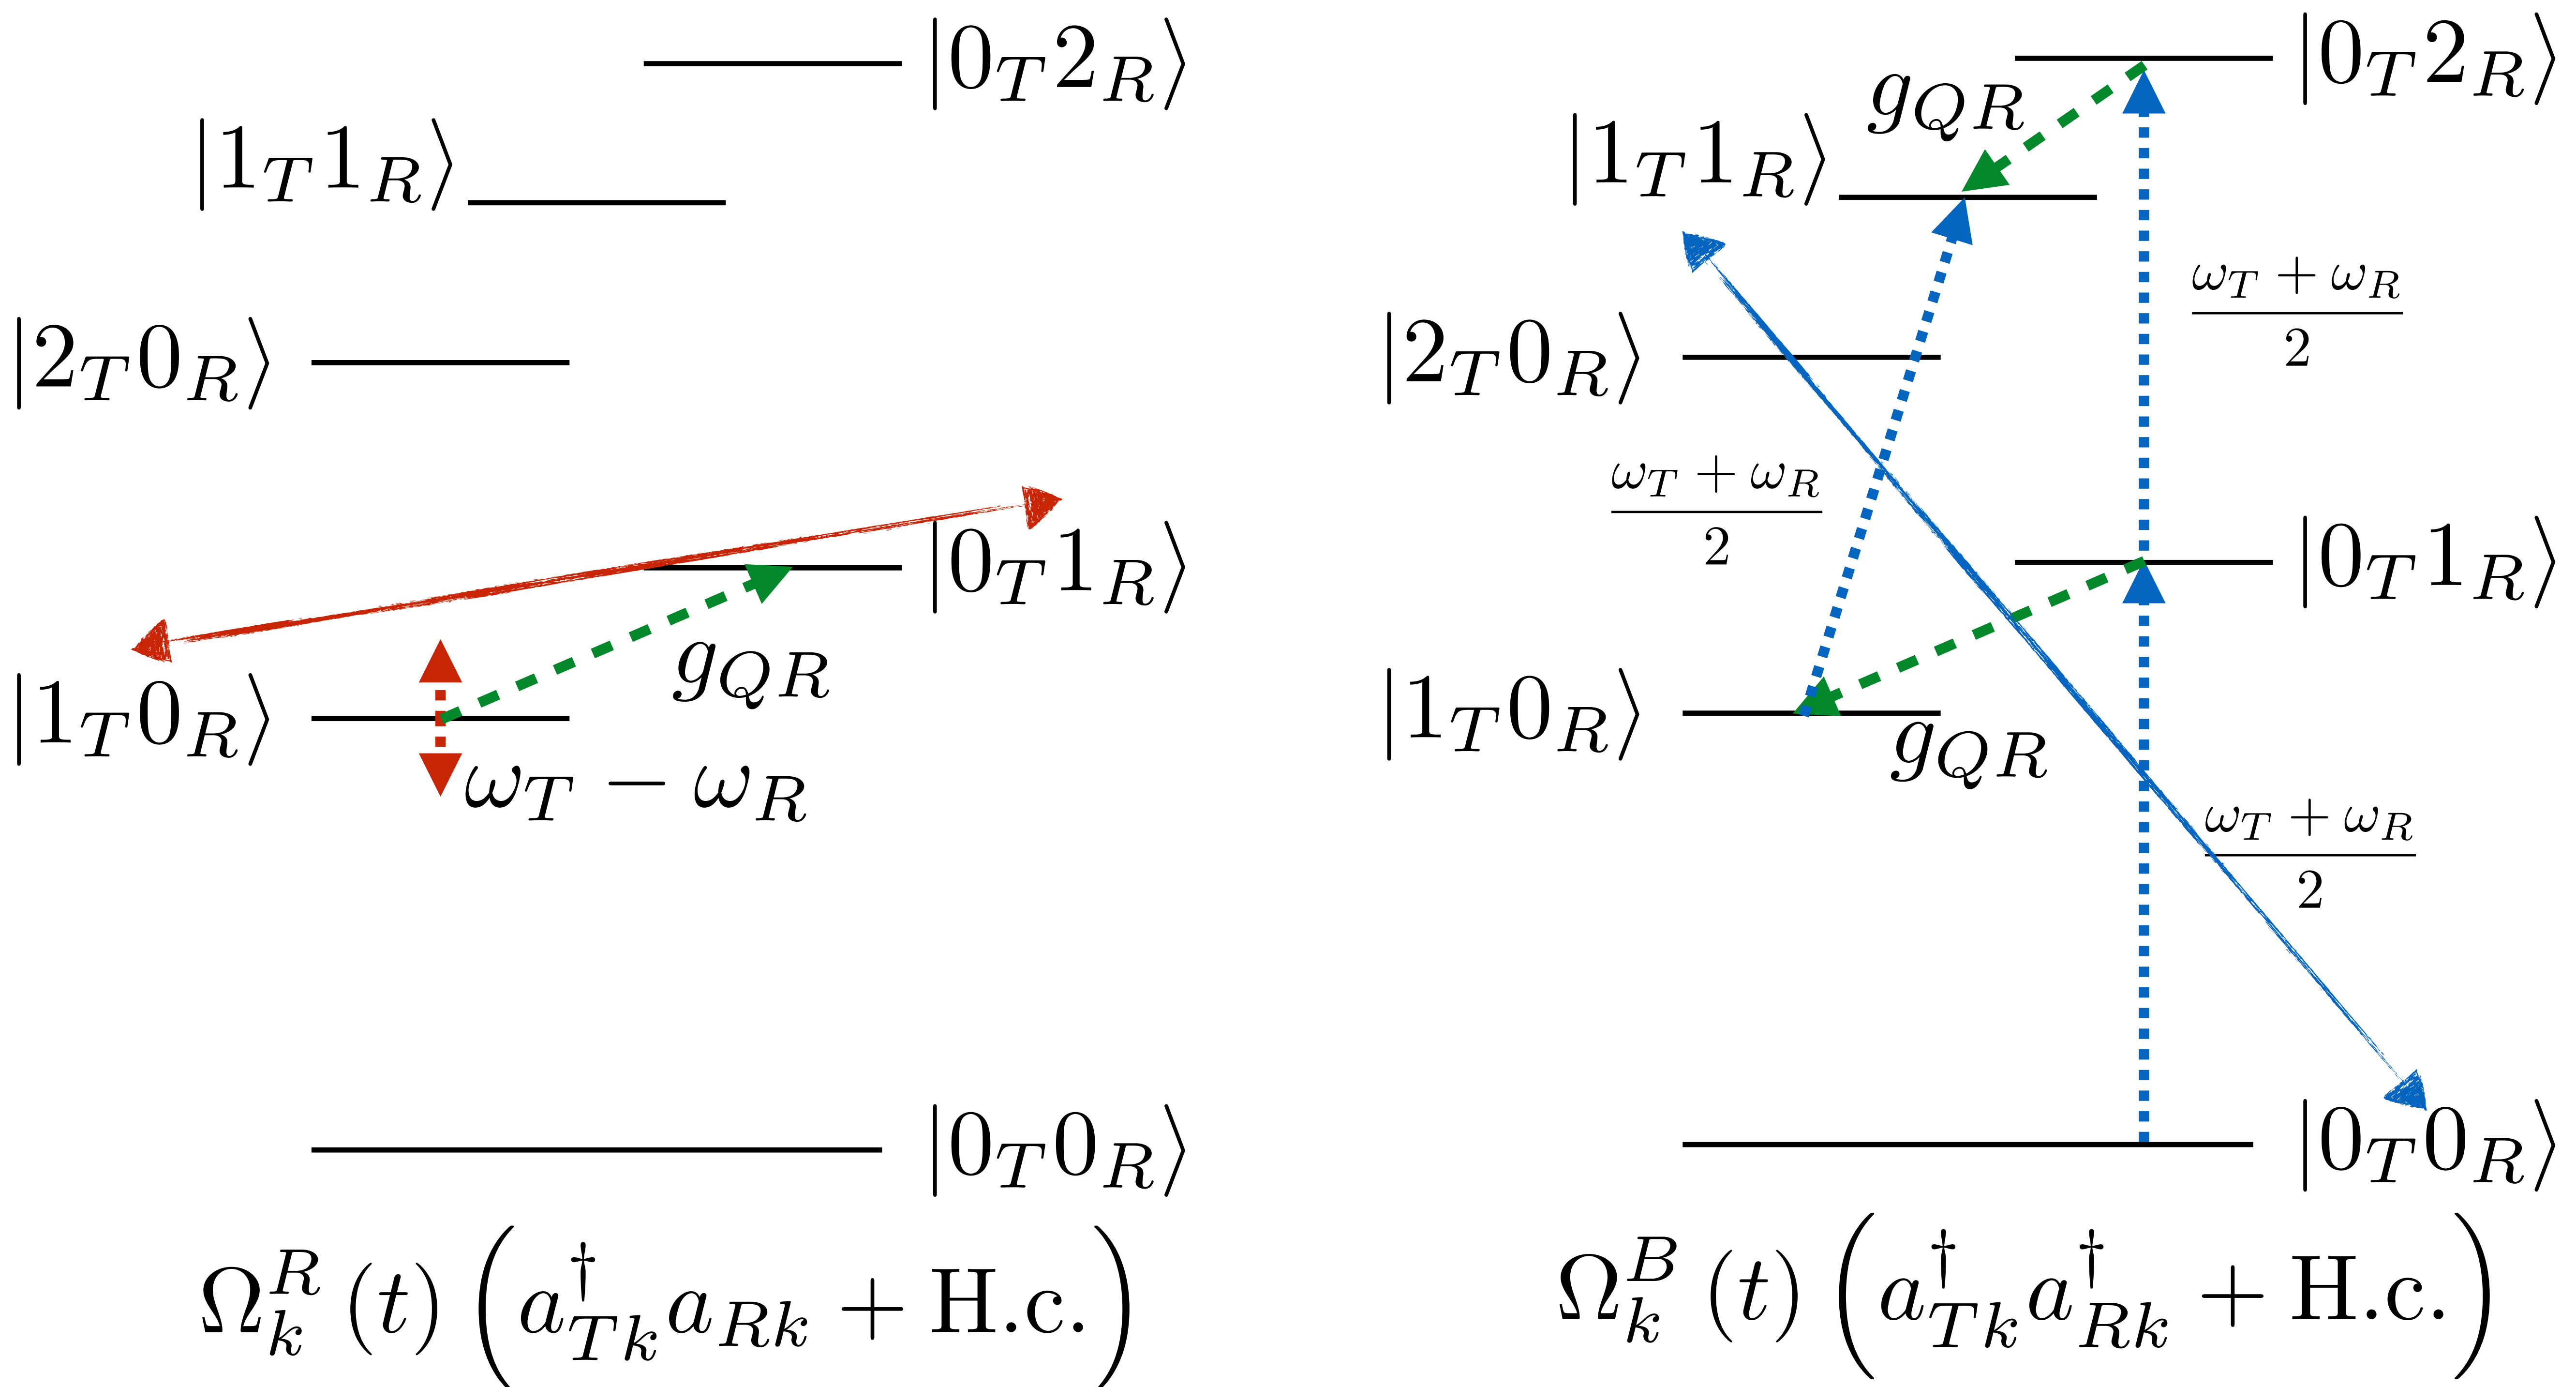

NOISY SAMPLING PROBLEMS

- Obvious concern: cavities are lossy, with typical loss rates $\Gamma_C \sim 10\text{MHz}$
- Can include this noise in the definition of the sampling problem (assume qubits are noise-free). Sample diagonals of a density matrix evolving as:

$$\partial_t \rho = i\left[H\left(t\right), \rho\right] + \sum_{i=1}^{K} \left(O_i \rho O_i^\dagger + \frac{1}{2}\left\{ O_i^\dagger O_i, \rho \right\} \right)$$

- But: photon loss in boson sampling or random quantum circuits (see right figure) drives problem toward triviality
- Must this always be the case?

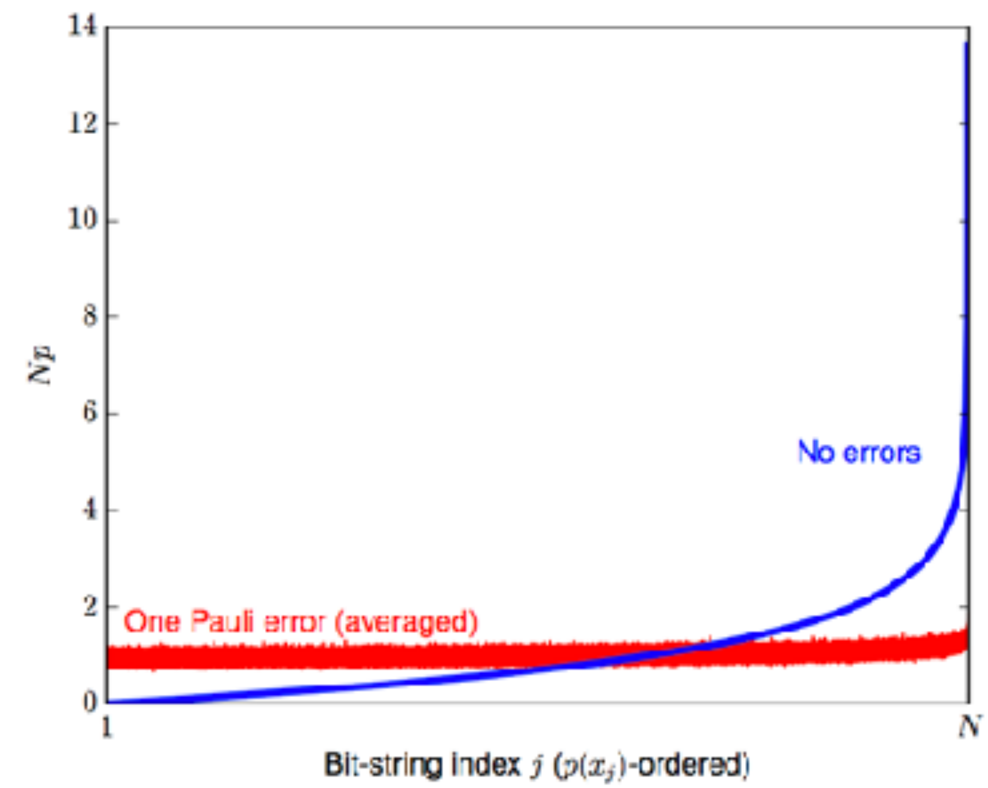

Image credits: Boixo et al, Nat. Phys. 2017

NOISY SAMPLING PROBLEMS

$$\partial_t \rho = i\left[H(t), \rho\right] + \sum_{i=1}^{K}\left(O_i \rho O_i^\dagger + \frac{1}{2}\left\{O_i^\dagger O_i, \rho\right\}\right)$$

- Lindblad evolution is capable of universal quantum computation (Verstraete *et al* Nat. Phys. 2009), but this construction is artificial

- Realistic local noise is trickier: Hermitian operations (e.g. Pauli errors) produce an incoherent walk in Hilbert space, uniform photon loss is similarly trivializing

- However, **local noise does not mean easy simulation!** Finite temperature simulation of systems with an MC sign problem is exponentially hard; note that noise operators for a finite T bath are extremely complex, even if they arise from local couplings…

- Guiding principle: *added noisy elements must be capable of generating quantum correlations*

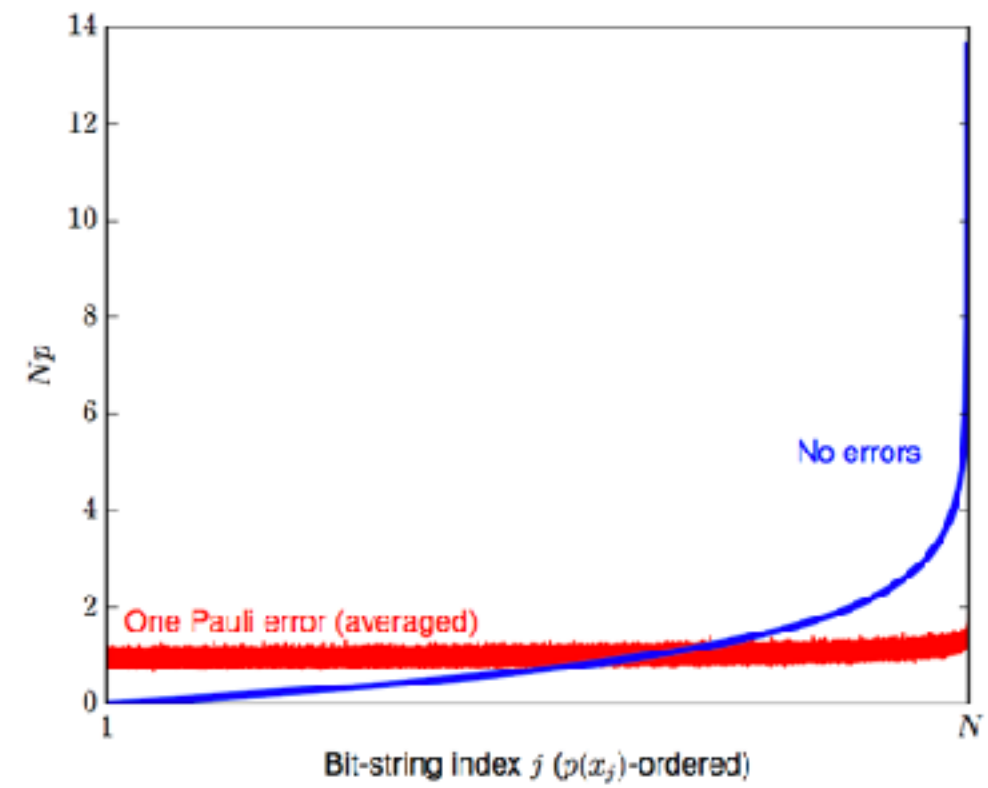

Image credits: Boixo et al, Nat. Phys. 2017

RESONANT COUPLING TO PROPAGATING MODES

- A way out: restrict loss to a subsystem (cavities). Turn on everything *simultaneously*. Make qubit-cavity couplings weak rel. to Q-Q couplings: $g \gg \Omega_{QC} \geq \Gamma_C$

- Resonance condition results: Photon addition/loss only significant when changing occupation of specific propagating modes

- Subsequent loss measures a highly nonlocal operator, with weight over entire lattice. Does not (necessarily) decorrelate state!

- See *Kapit, Quant. Sci. Tech. 2017* for a review

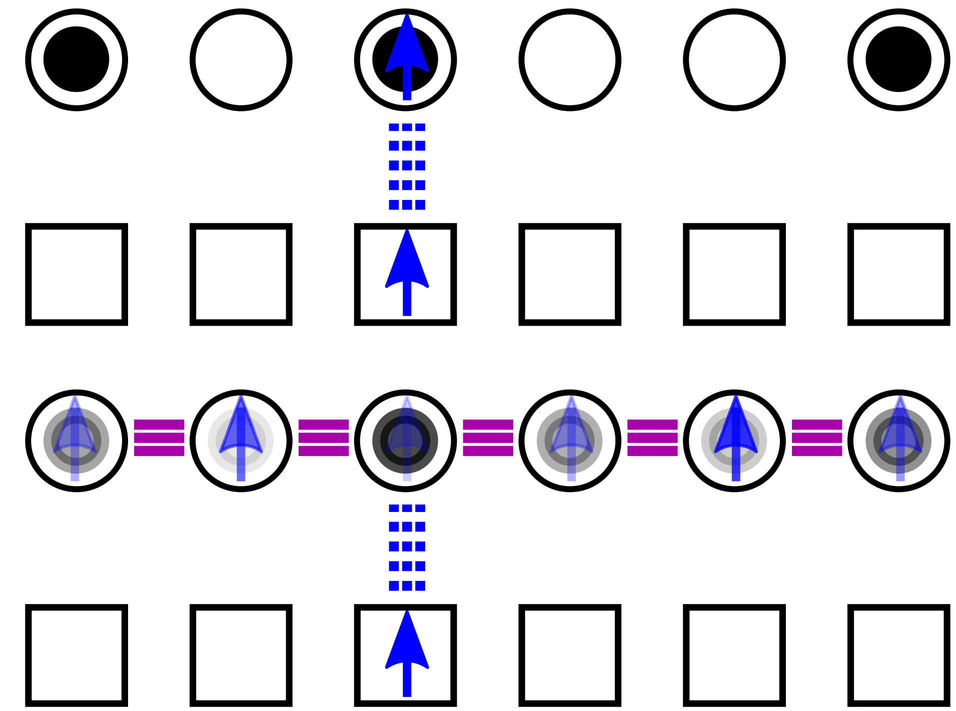

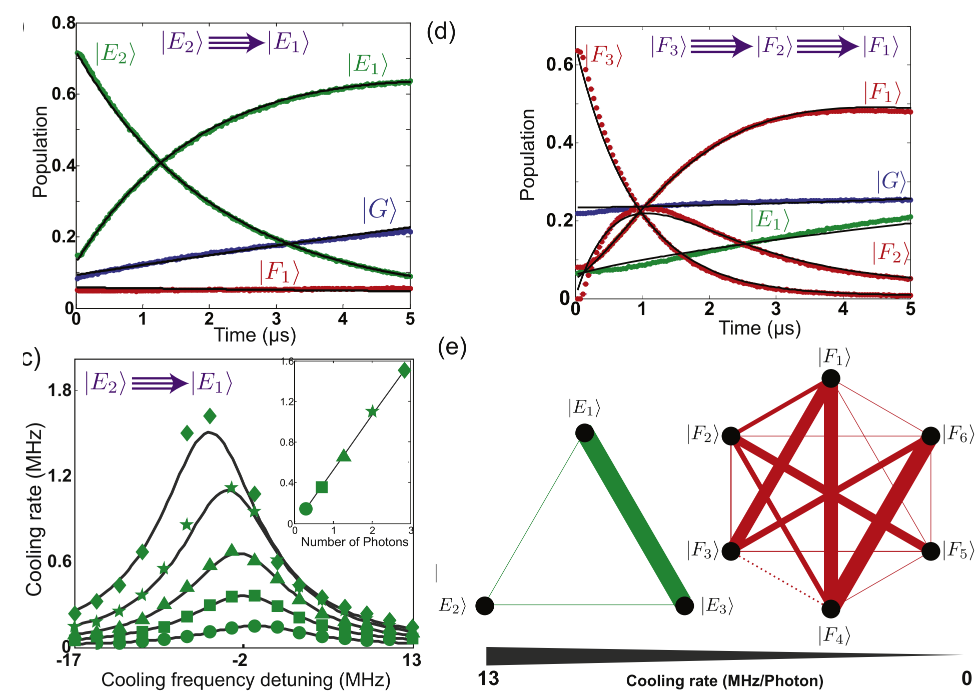

Image credits: Hacohen-Gourgy et al, PRL 2015

PROPOSED EXPERIMENT

- In summary, our protocol is:
- **Prepare** initial product state in z basis. Set qubit detunings.
- **Pulse** all qubit-qubit and qubit-cavity couplers on and off *simultaneously*, N times. QC much weaker than QQ.
- **Measure** state of all qubits (cavities not measured)
- **Repeat** many times to sample output distribution, compare to theoretical model.

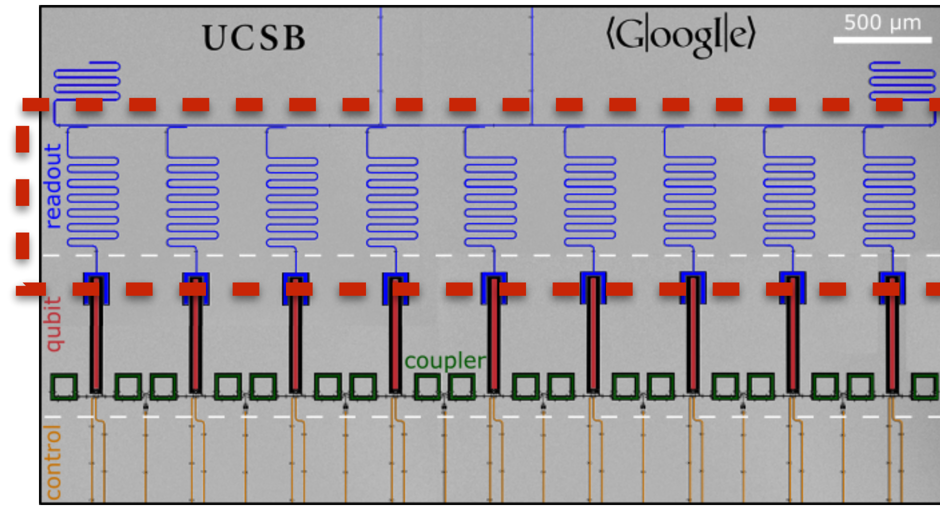

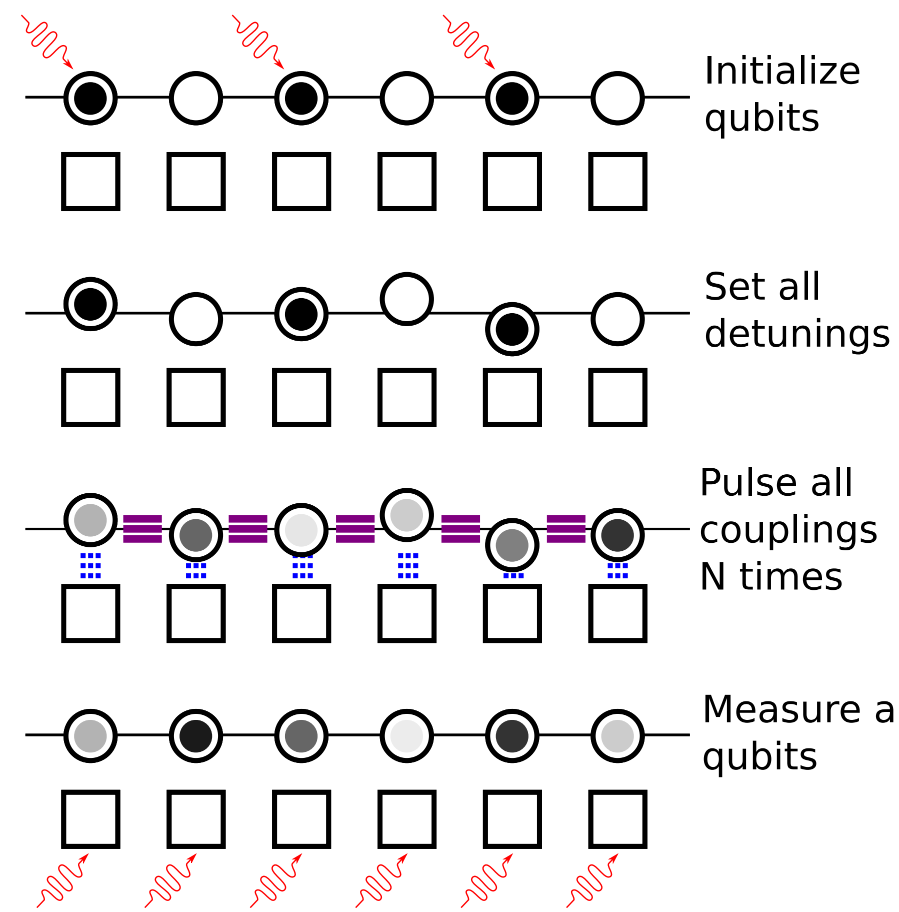

Image credits: Neill et al 2017

NUMERICAL BENCHMARKS

- **Output distribution details:** # fluctuations, loss rates, IPR,
- **Volume entanglement**
- **Fidelity loss** from qubit error & basis truncation

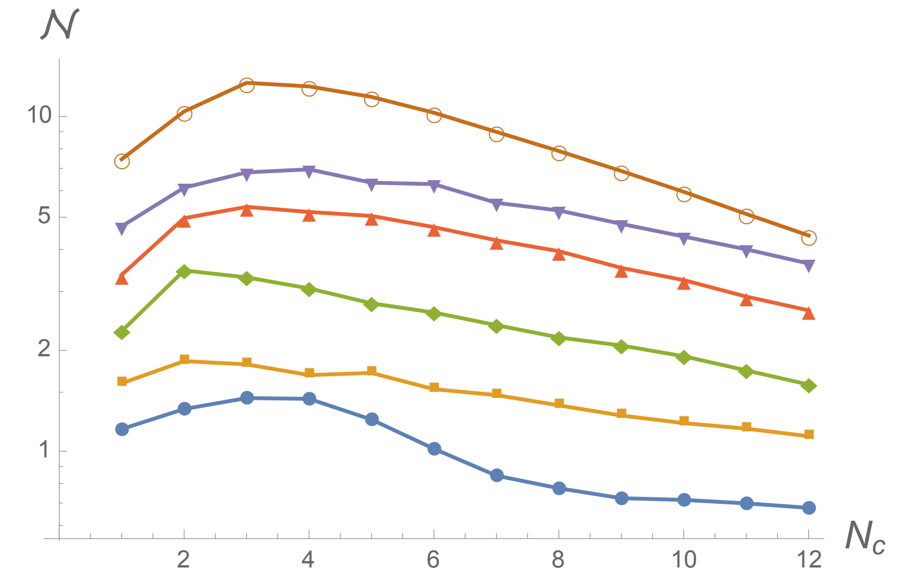

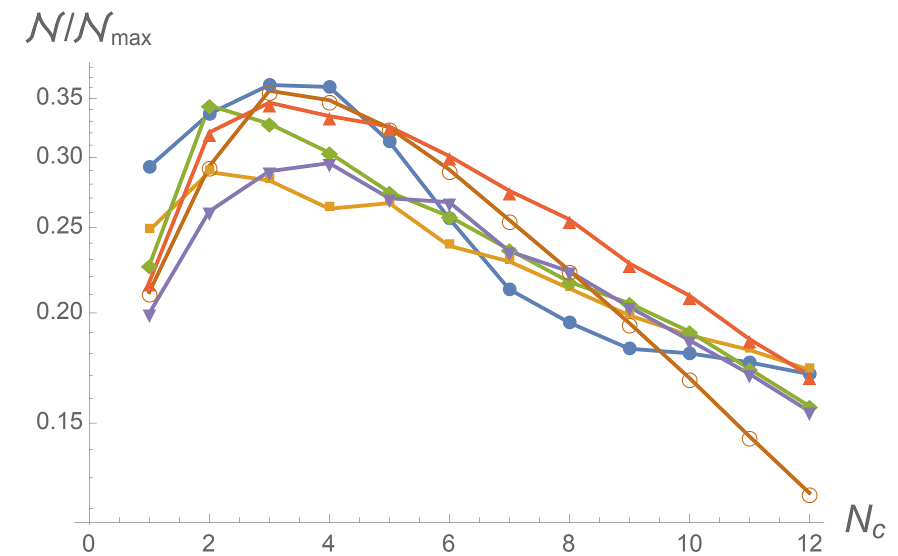

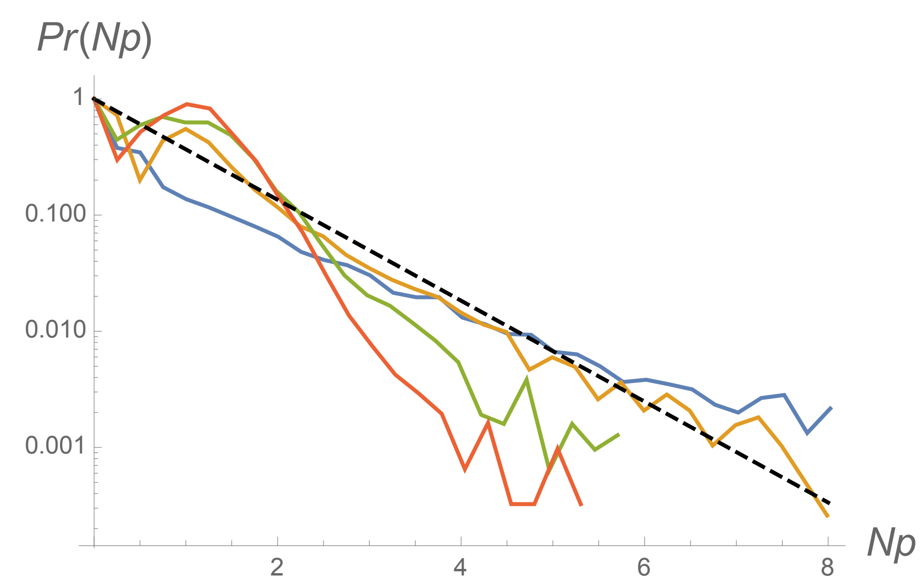

SIMULATION DETAILS

- Numerically simulate evolution up through 12 coupler pulse cycles, using full wfn evolution with quantum trajectories (see Daley 2014 for a review)
- Simulate only blue sidebands, for simplicity and experimental relevance— other choices to be discussed in paper
- Experimentally realistic parameters (all in MHz unless otherwise noted):

$$g_{max} = 2\pi \times 40, \quad \delta = -2\pi \times 200, \quad \Delta = 2\pi \times 5$$

$$h_{i,max} = \pm 2\pi \times 20, \quad \Omega_{QC,max} = 2\pi \times 3, \quad h_C = 0$$

$$\Gamma_C = 10, \quad t_{cycle} \in \{20, 30\}\ \mathrm{ns}, \quad N_{cycle} = 12$$

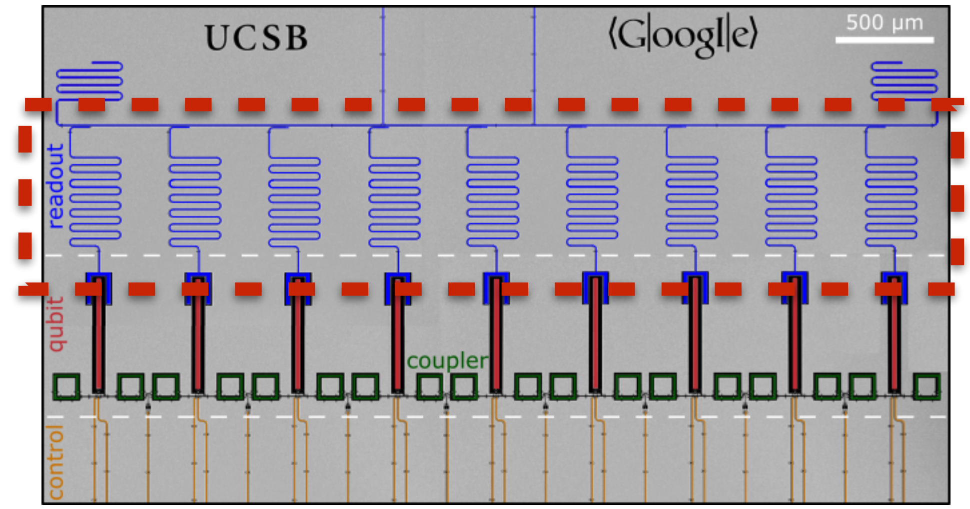

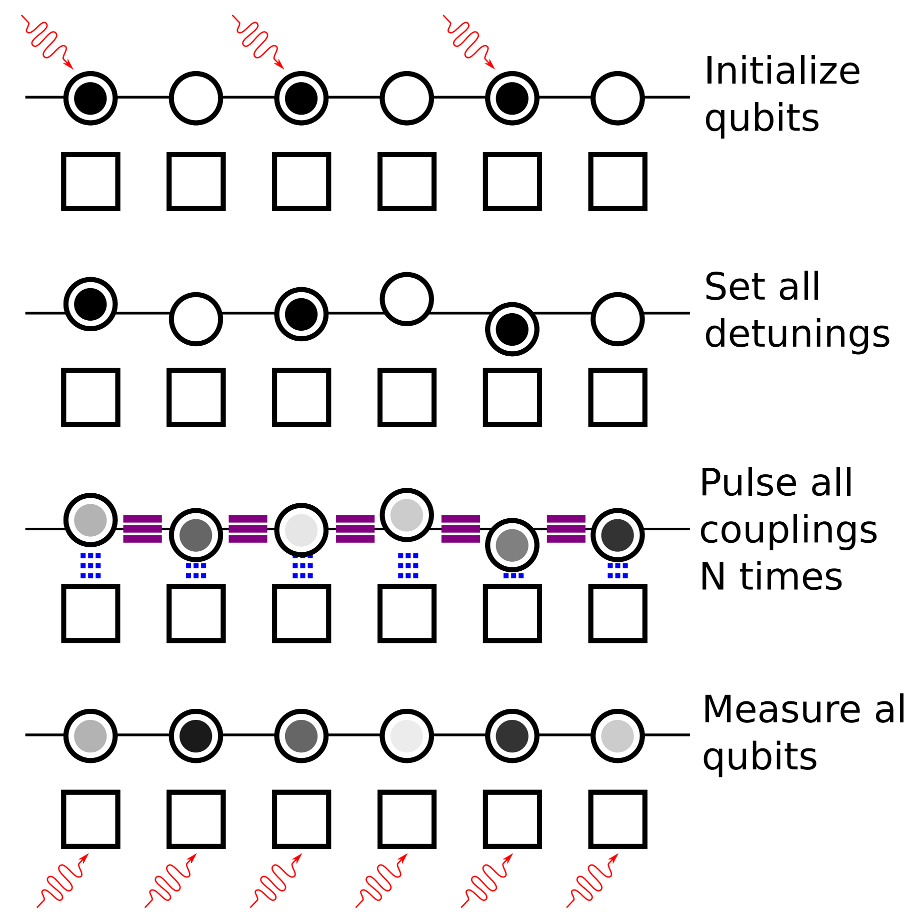

Image credits: Neill et al 2017

WHAT WE COMPUTED

- **Number fluctuations:** extensive scaling shows we can extrapolate these results to larger L
- **IPR:** shows that an O(1) fraction of Hilbert space is explored
- **Output statistics:** information scrambling, suggests intermediate-time chaotic behavior
- **Negativity:** demonstrates volume-law entanglement scaling
- In sum, these measures suggest an efficient classical method is extremely unlikely for this problem

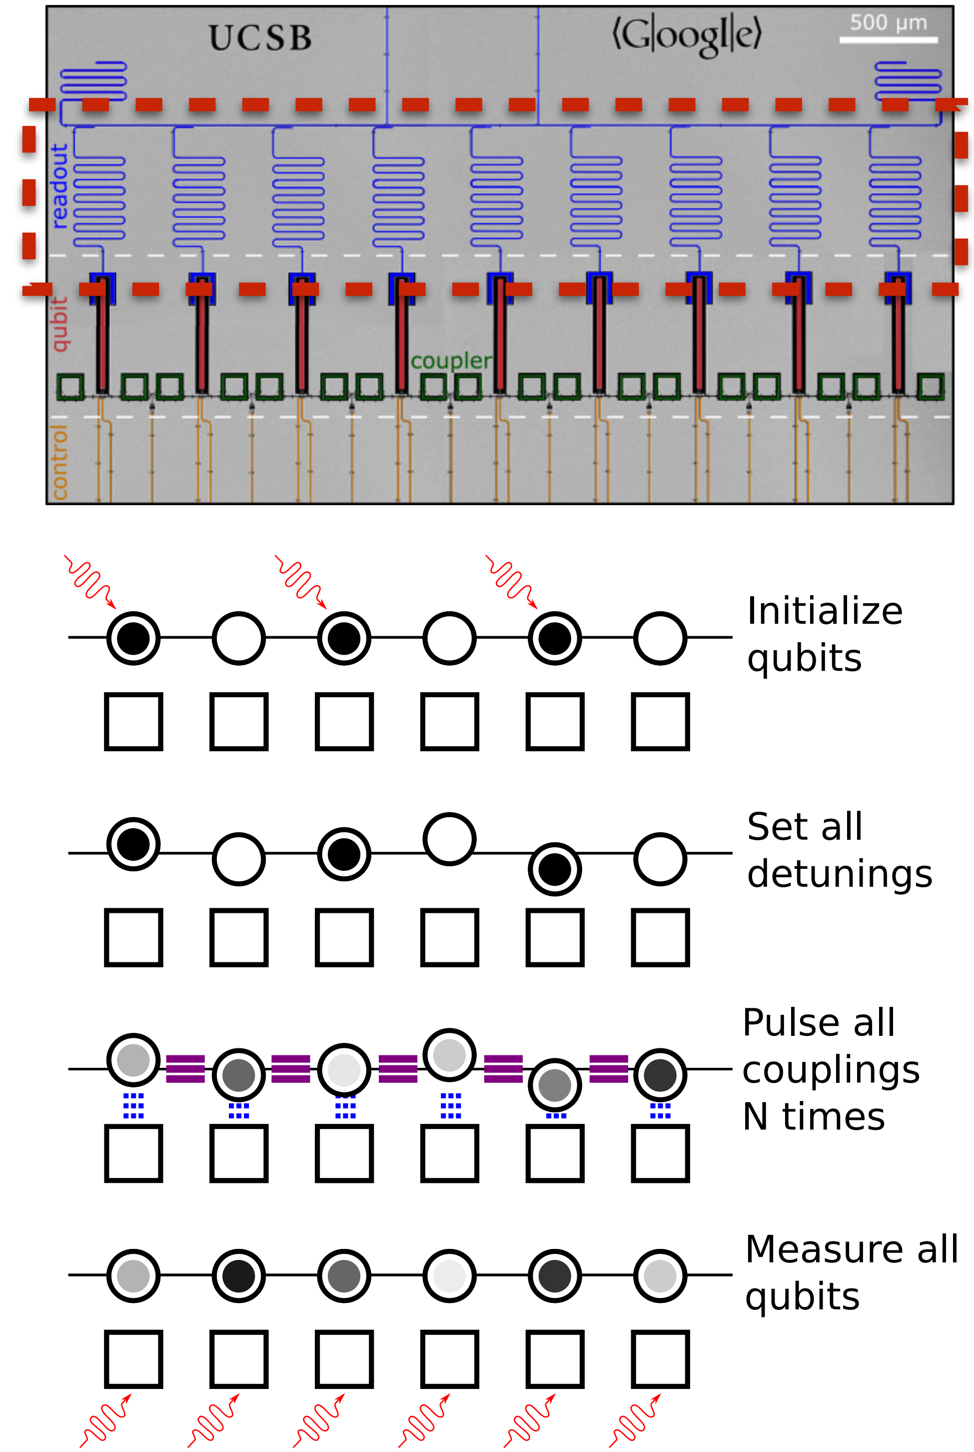

Image credits: Neill et al 2017

OUTPUT DISTRIBUTION: NUMBER FLUCTUATIONS AND CAVITY PHOTON LOSS

- Initialize the system with L/2 -1 photons, rounded down, in Neel-like state
- Top to bottom (averages): # photons added to qubits, # photons (total) in cavities, # photons lost from cavities (cumulative)
- **All increase extensively**, though with noticeable even-odd effects
- Hypothesis: even-odd effect primarily due to existence of zero energy hopping mode for odd L
- **Key: L=4 (blue), 5 (gold), 6 (green), 7(red), 8 (purple), 9 (brown), 10 (light blue), 11 (yellow)**

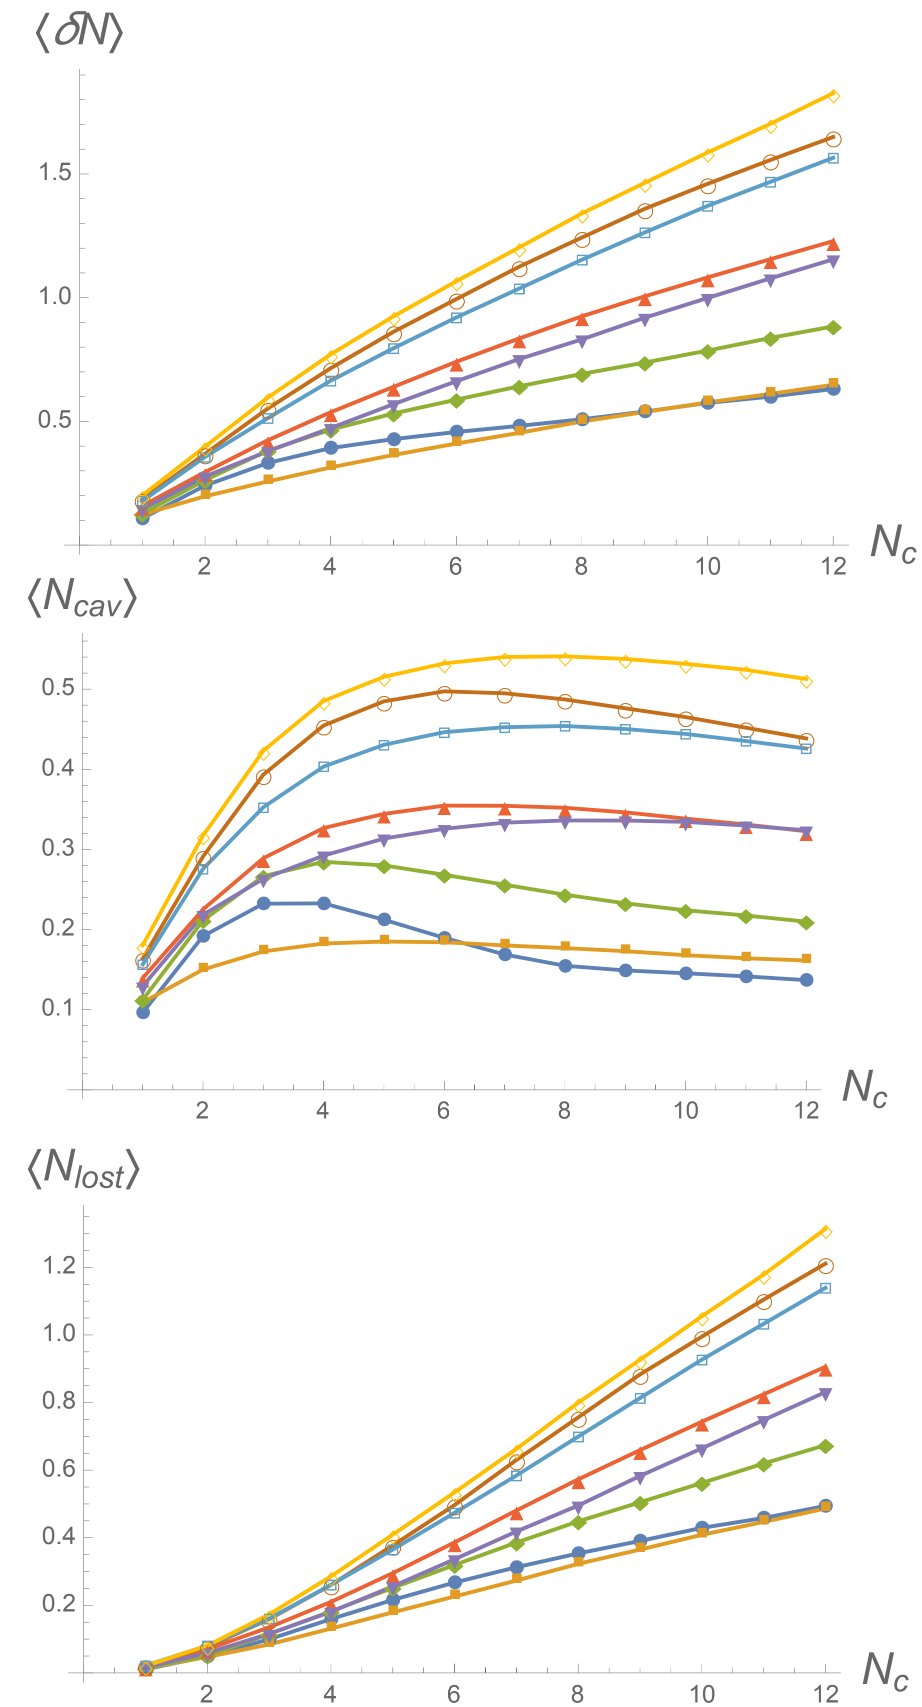

OUTPUT DISTRIBUTION: INVERSE PARTICIPATION RATIO AND OUTPUT HEAVINESS

- System explores O(1) fraction of qubit Hilbert space!
- Right: inverse participation ratio, output heaviness, and example probability distribution for L=9 for Nc=3 (blue) cycles, 6 (gold), 9 (green), and12 (red)
- In 2016, Aaronson and Chen proposed "Heavy Output" as a quantum hardness criteria; likely exponentially hard to sample a quantum distribution (from an RQC) and produce outputs with greater than median probability more than 2/3 of the time. For a P-T dist approximately 85% of outputs satisfy this. Satisfied for all studied cases here as well.
- **Key: L=4 (blue), 5 (gold), 6 (green), 7(red), 8 (purple), 9 (brown), 10 (light blue), 11 (yellow)**

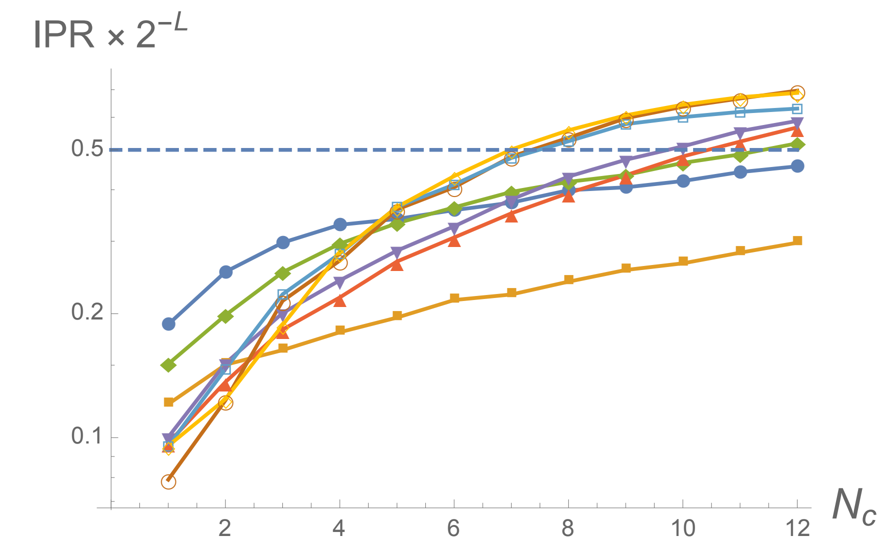

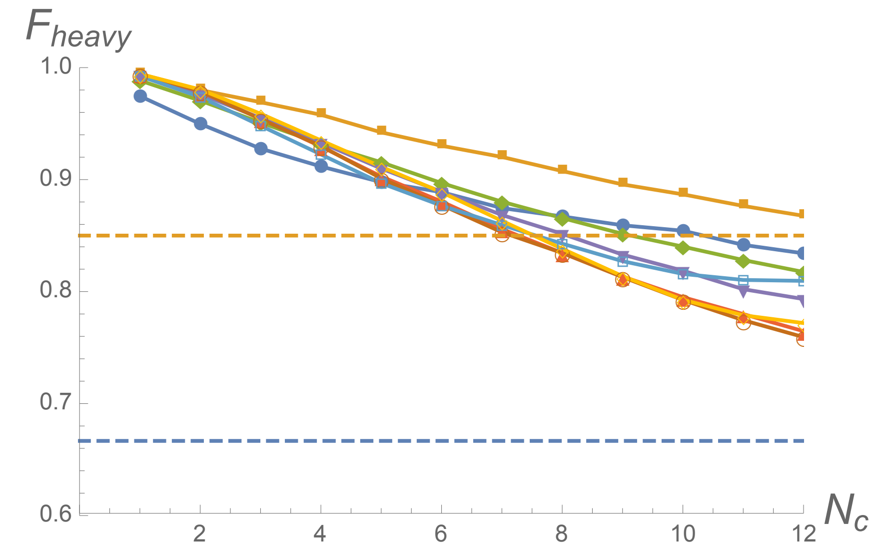

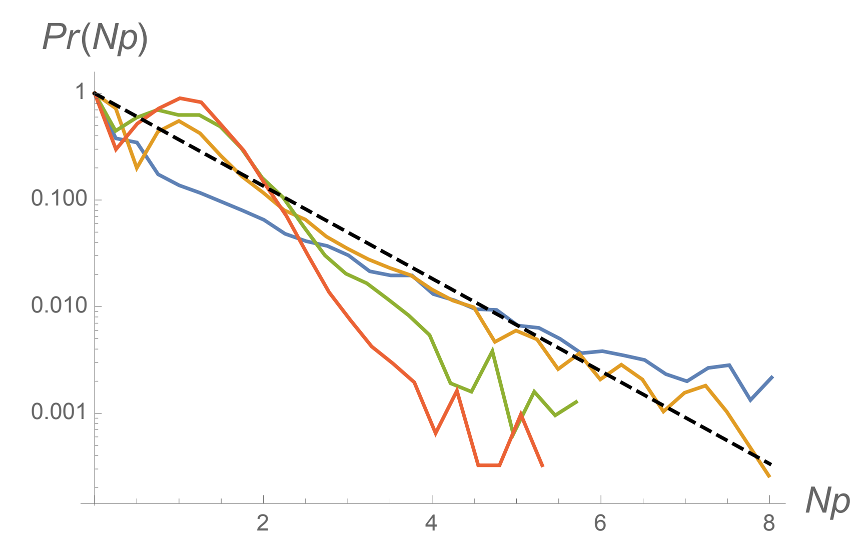

OUTPUT DISTRIBUTION: SCRAMBLING

- Evidence of intermediate-time quantum chaotic behavior, with a likely trivial final state at very long times
- System comes very close to an exponential (Porter-Thomas) distribution over full qubit Hilbert space before slowly moving away from it. **Note:** cavity photon loss is O(1) by the time this point is reached!
- Not well-approximated by incoherent uniform randomness, or a reweighted version (poisson distributed random photon addition)
- Right: K-L divergence from Porter-Thomas, IUR, and reweighed IUR
- **Key: L=4 (blue), 5 (gold), 6 (green), 7(red), 8 (purple), 9 (brown), 10 (light blue), 11 (yellow)**

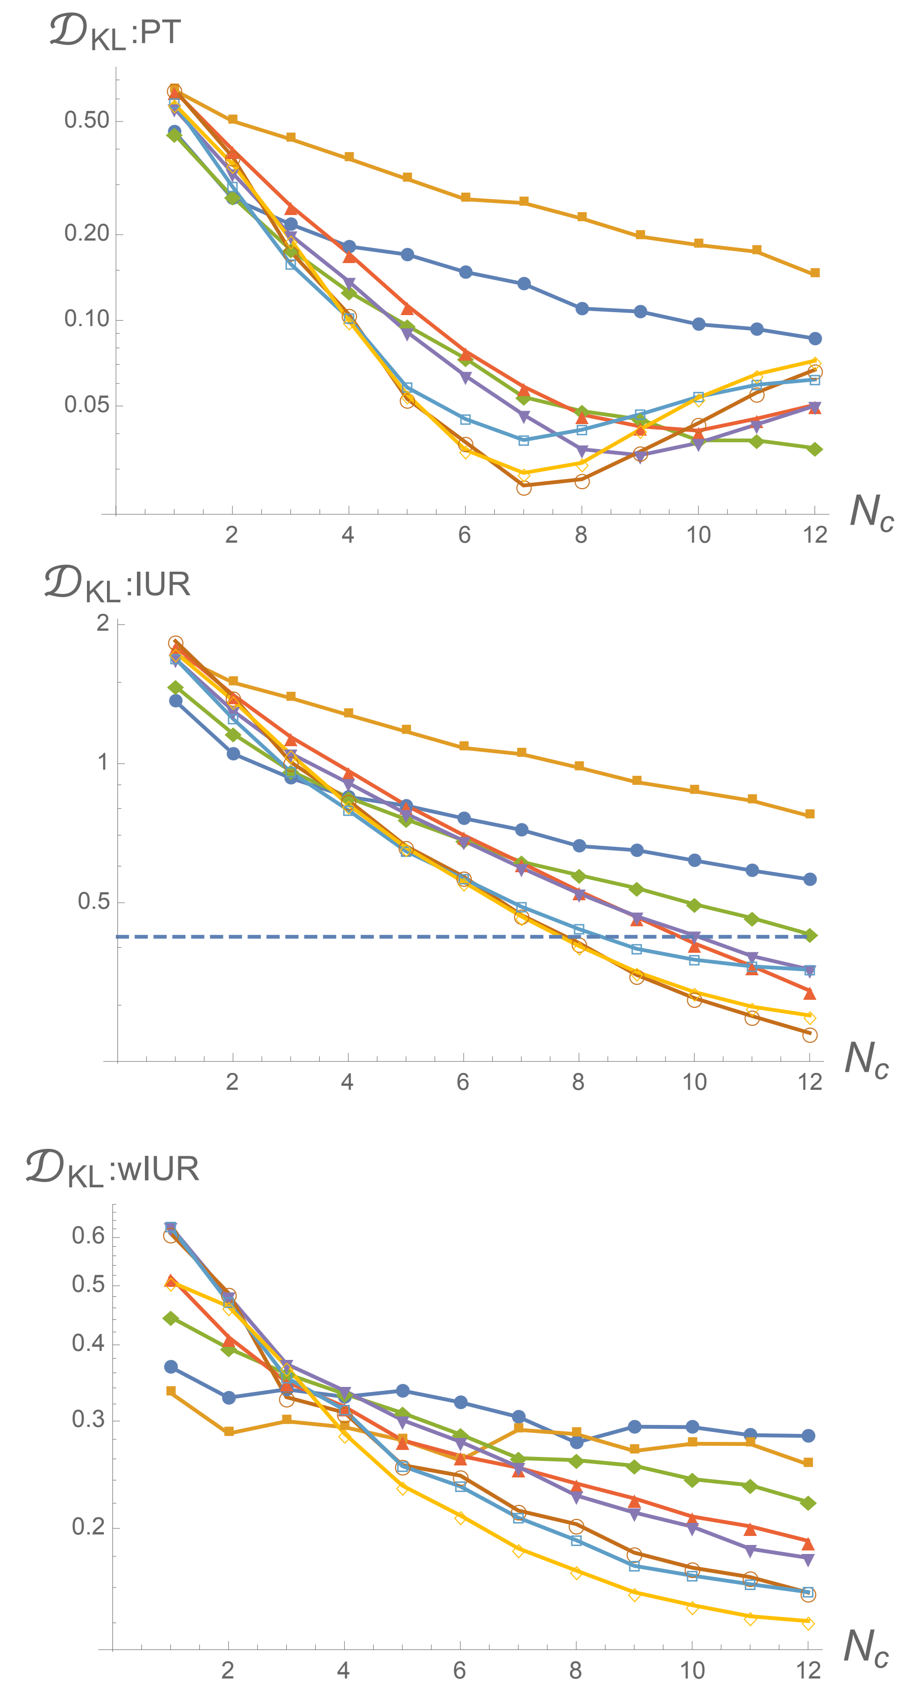

BENCHMARKS: VOLUME ENTANGLEMENT

- Open quantum system: standard Von Neumann EE does not capture entanglement
- Instead, we measure the **negativity** (*Vidal and Werner PRA 2002*):

$$\mathcal{N} = \frac{1}{2}\left(||\rho^{T_A}|| - 1\right)$$

- As with VNEE, bipartition system into subsystems A and B
- Here, ρ^{T_A} is the *partial transpose* of joint density matrix w.r.t system A. Negative EVs in matrix norm iff system is entangled. Nonzero negativity sufficient condition for entanglement.
- Unlike VNEE, equally well defined for pure states and mixed states
- For a maximally entangled state, the maximum negativity of a perfect bipartition is exp. large:

$$\mathcal{N}_{max} = \frac{1}{2}\left(\sqrt{N_H} - 1\right)$$

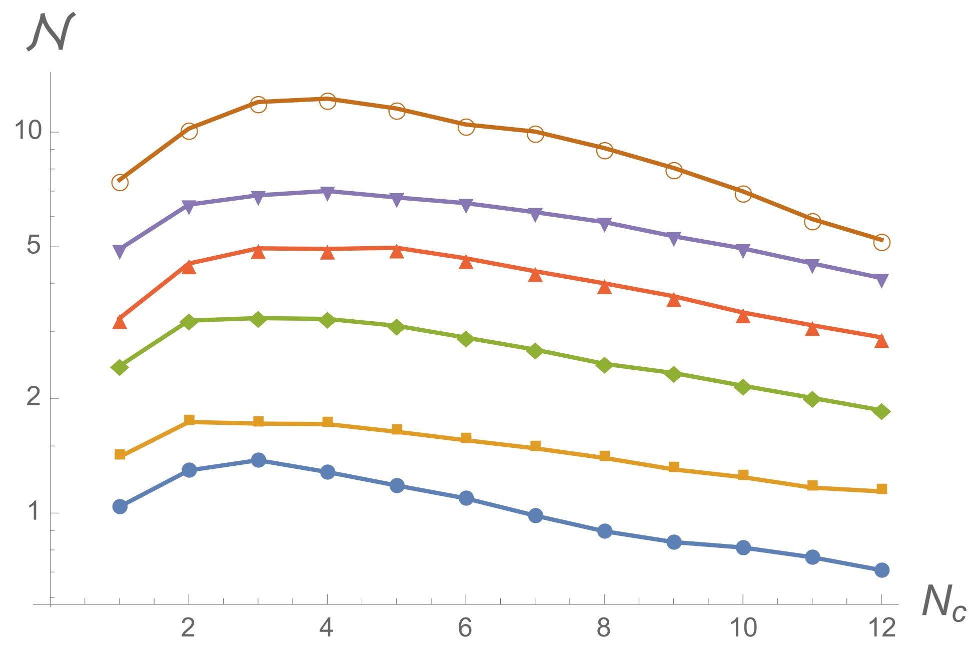

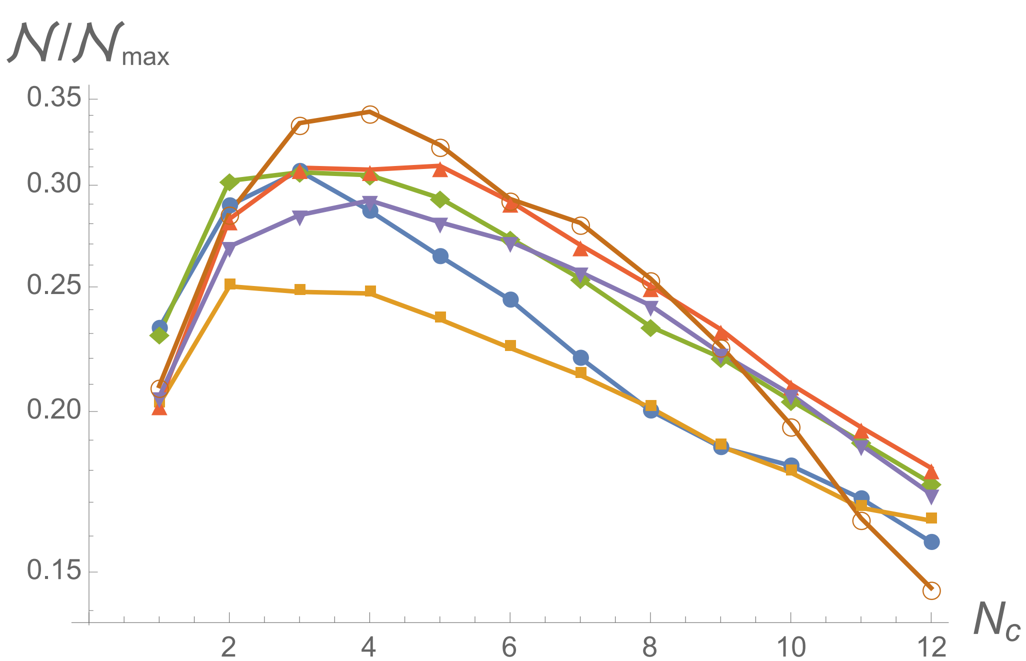

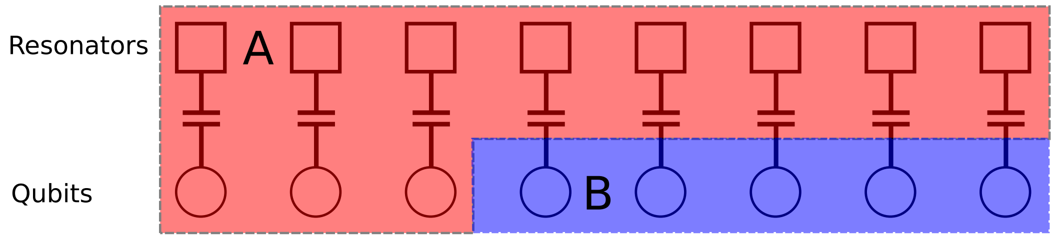

BENCHMARKS: VOLUME ENTANGLEMENT

- Bipartition system into subsystems A and B, where A includes all the cavities and somewhat less than half the qubits

- Respects truncated cavity Hilbert space and fact that photon density in cavities is low; aim for approx. equal Hilbert space size in each partition

- Note: **very** expensive to compute. For 9 qubits and up to two cavity photons (or 11 qubits, 1 cav. photon), computing negativity takes ~ 30 GB RAM.

- **Can also compute subsystem negativity:** negativity of the qubit reduced density matrix (bipartitioned) w/ cavities traced out.

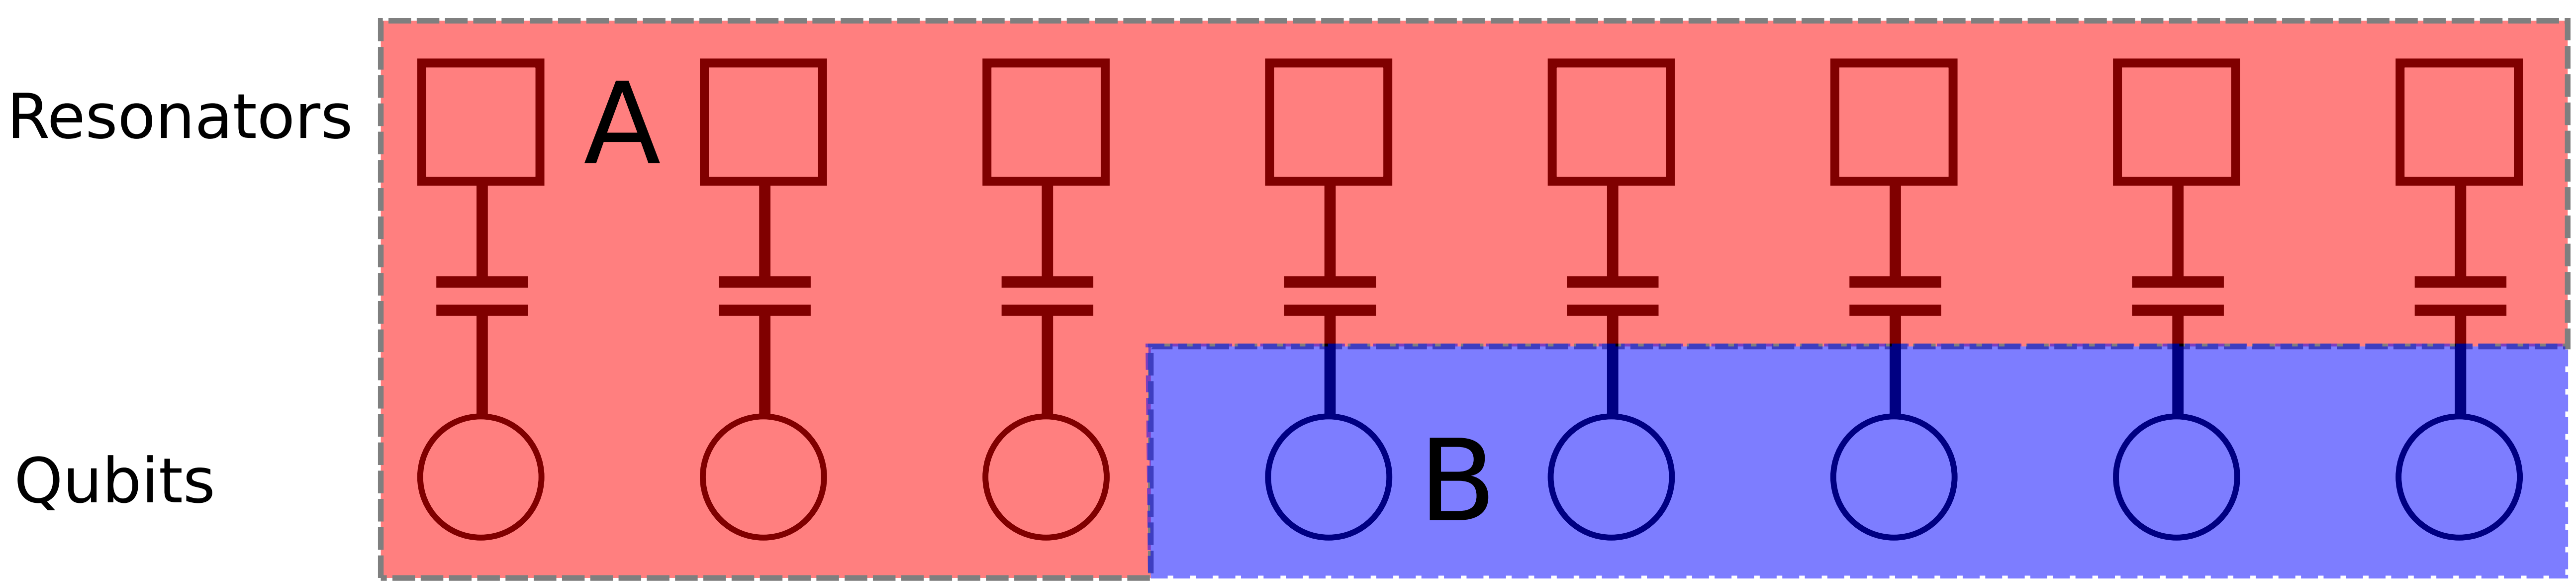

VOLUME ENTANGLEMENT POST-LOSS

- Trajectory method has another advantage: can bin trajectories to check entanglement with total cavity photon loss # fixed
- So we bin simulations to require exactly 1 photon has been lost by the end of 12 cycles
- For RQC, this gives garbage: incoherent uniform randomness, zero entanglement (when averaged over error position/time)

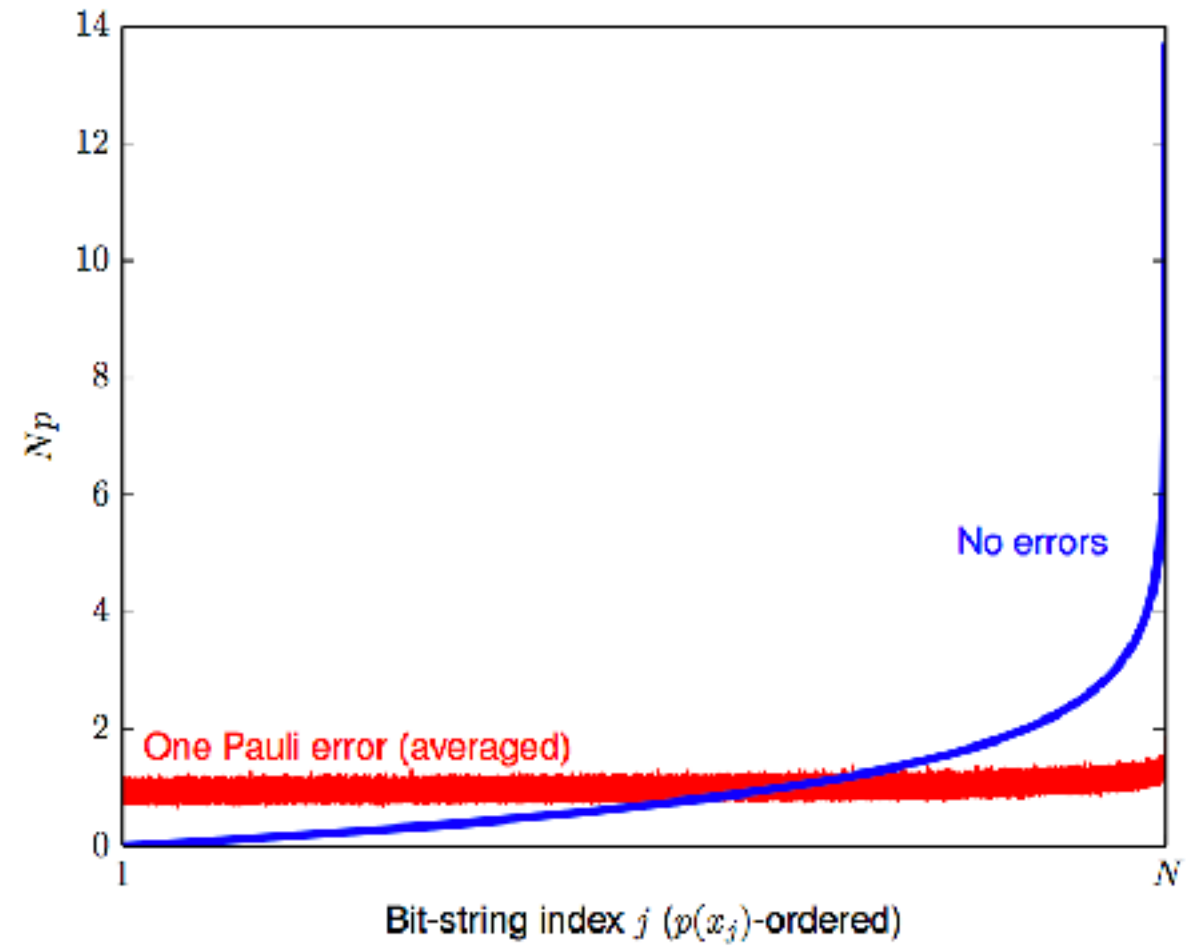

Image from Boixo et al, Nature Phys. 2018

NEGATIVITY: FULL SYSTEM

- Left: negativity of the full system (log scale), both unaltered and divided by its maximum possible value (assume cavity Hilbert space dimension L+1). Right: same things constructed only from trajectories where one photon has been lost by the end of 12 cycles. Volume entangled in both cases (w/ even-odd effects), should persist to computationally intractable L.

Full:

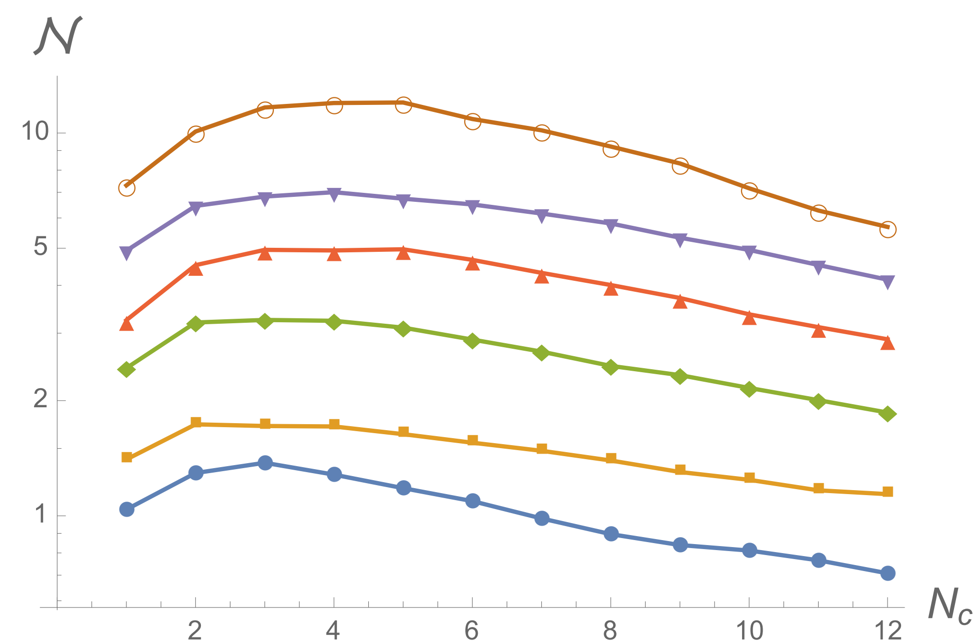

One Loss:

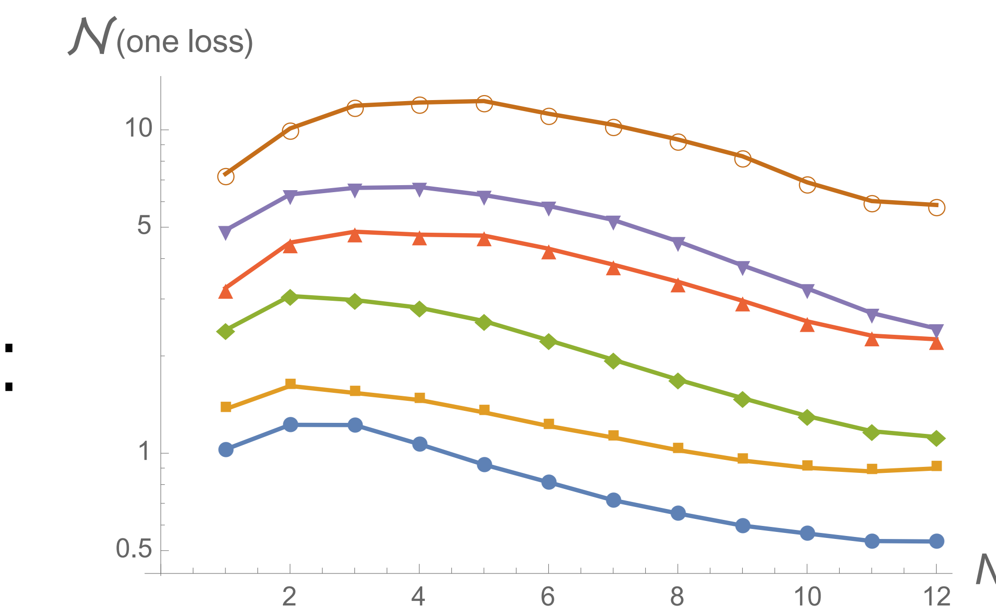

Full:

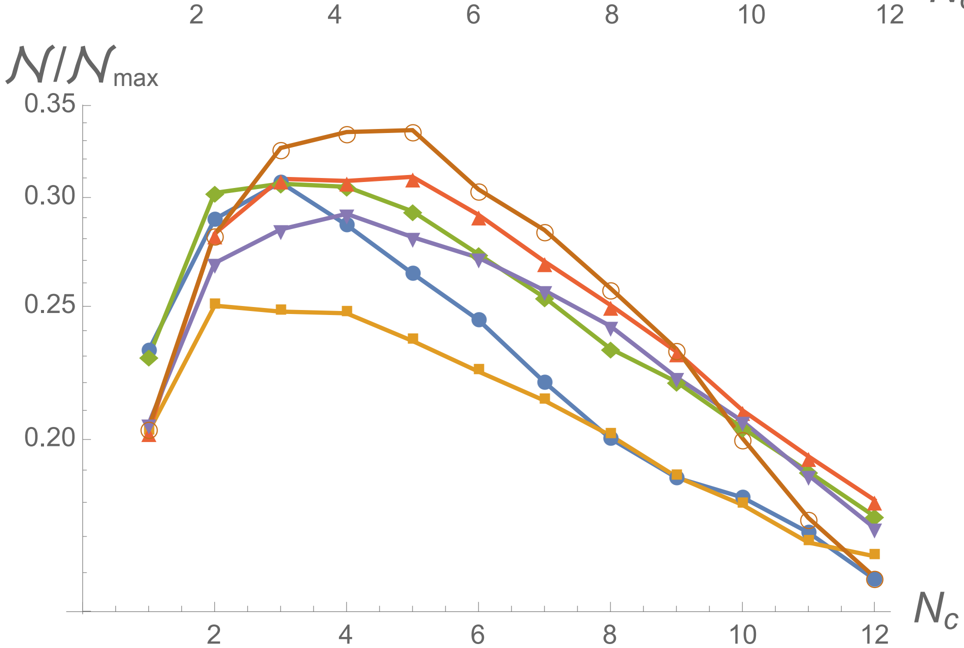

One Loss:

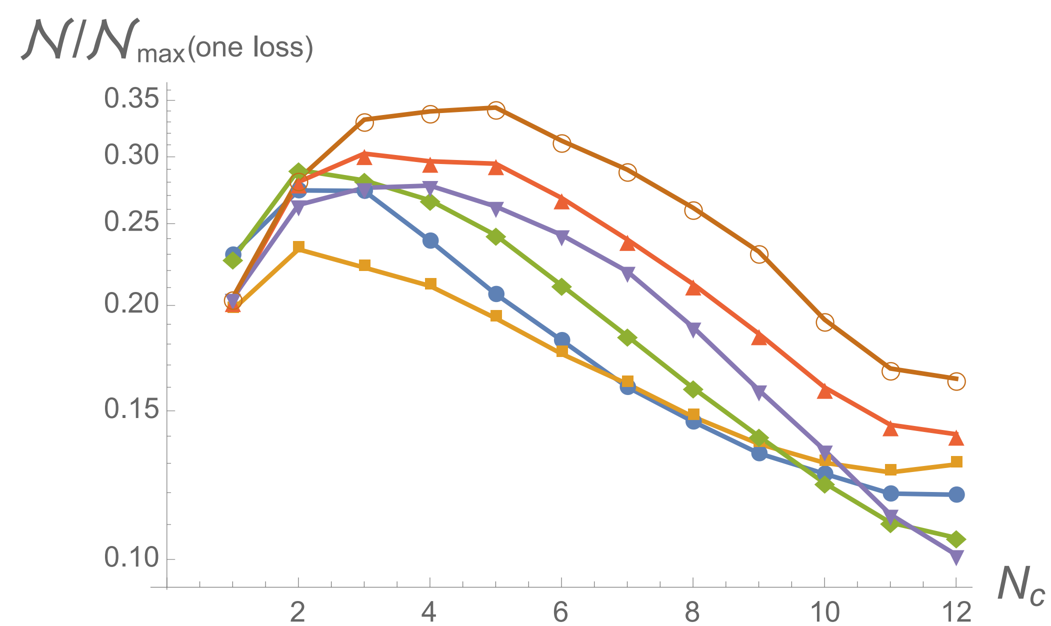

NEGATIVITY: QUBIT SUBSYSTEM

- Now, the negativity of the qubit subsystem, with cavities traced out. Obviously smaller, but still volume-scaling. Proves that **cavity photon loss reduces entanglement but does not fully disentangle state.** Values achieved at 6 cycles are approx. 1/2 negativity measured in unitarily evolving, closed-system protocol (smaller due to entanglement with cavities acting as measurement when traced out).

Full:

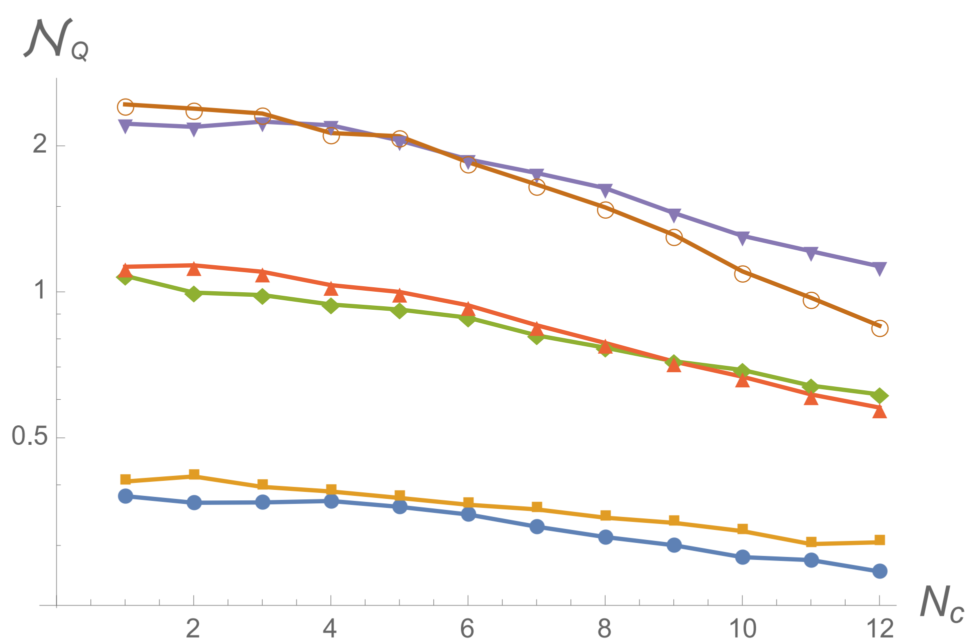

One Loss:

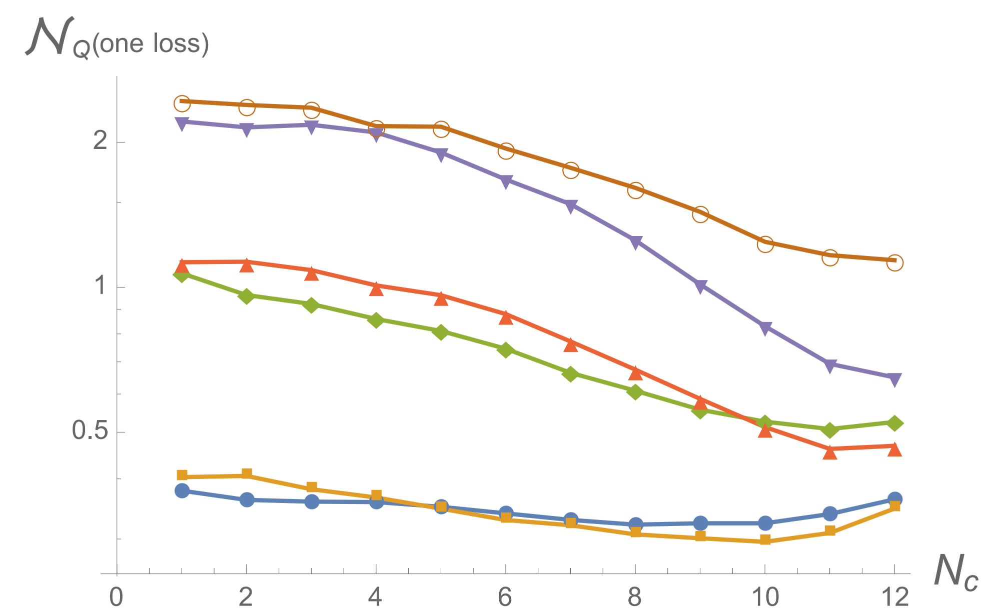

Full:

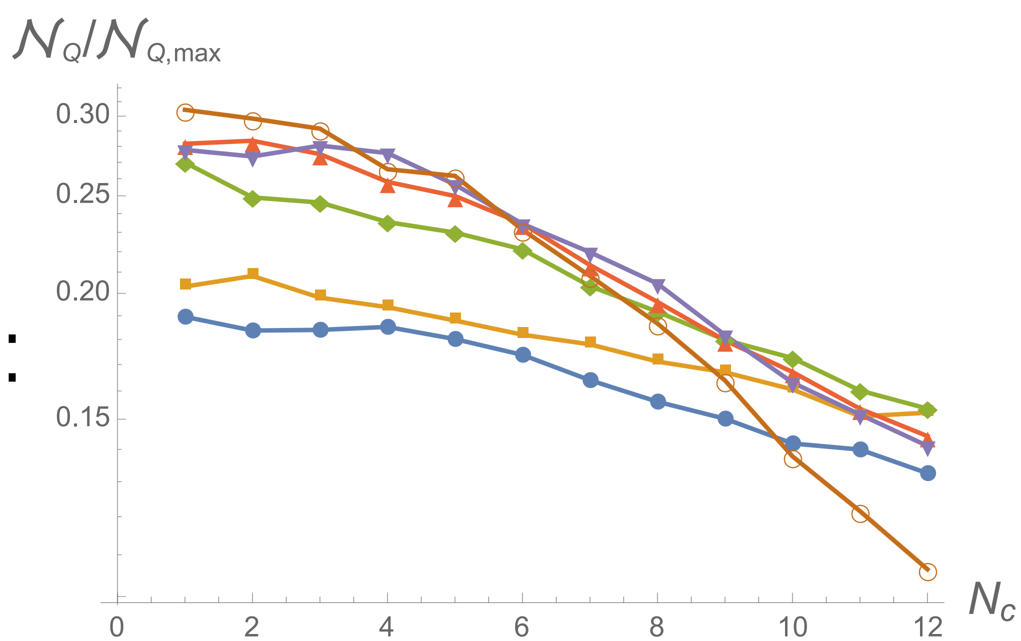

One Loss:

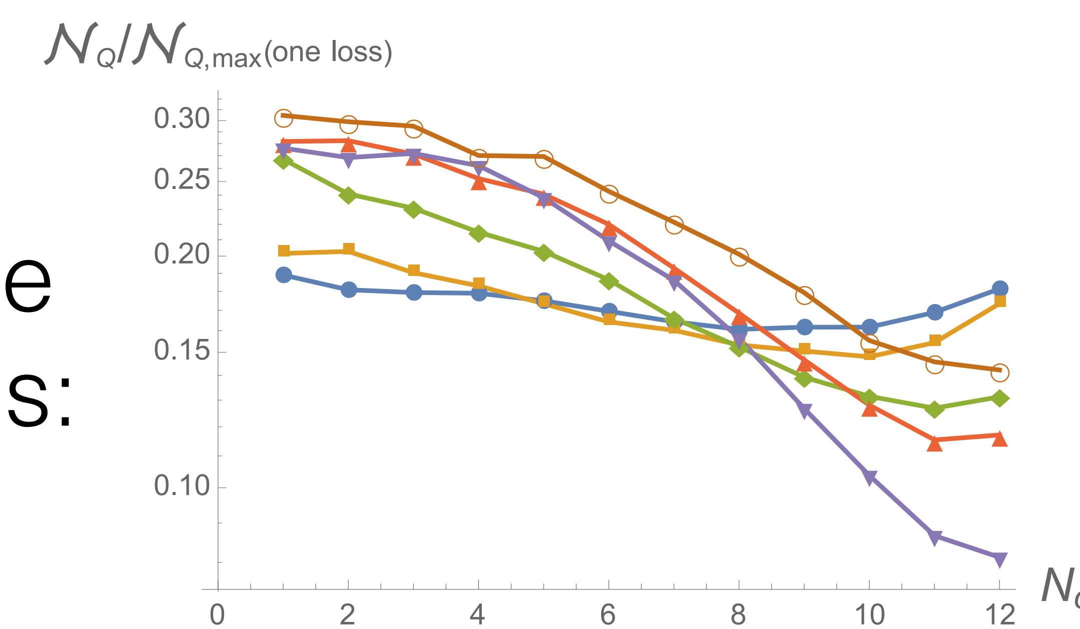

FIDELITY LOSS FROM QUBIT ERROR

- Define fidelity using the K-L divergence as in previous works:

$$\mathcal{F} \equiv 1 - \frac{\mathcal{D}_{KL}\left(\rho_{ideal}, \rho_{actual}\right)}{\mathcal{D}_{KL}\left(\rho_{ideal}, \rho_{TC}\right)}$$

- Here, ρ_{TC} is a trivial classical distribution generated by reweighting incoherent uniform randomness by total particle # (assume Poisson distributed # fluctuations).

- Considered both phase noise and photon loss. In both default protocol & modified noisy ones, single photon loss sends $\mathcal{F}$ to zero. However in default protocol, this can be post-selected out.

- Effect of phase noise is more subtle.

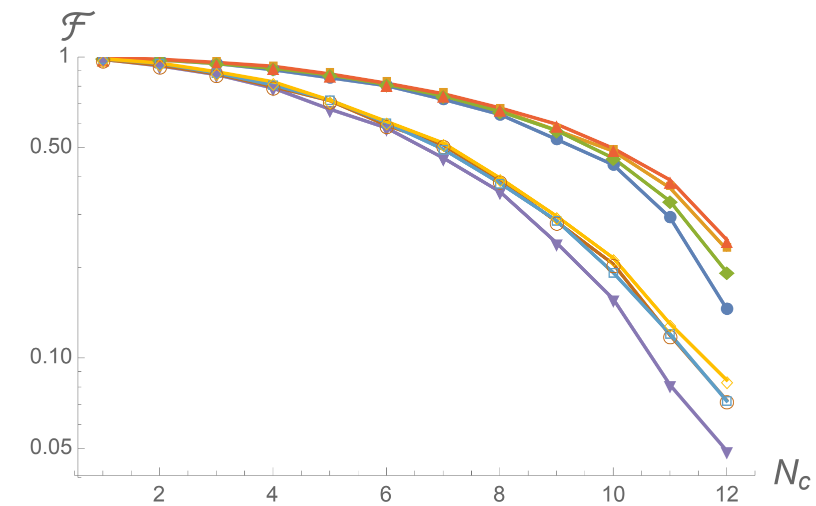

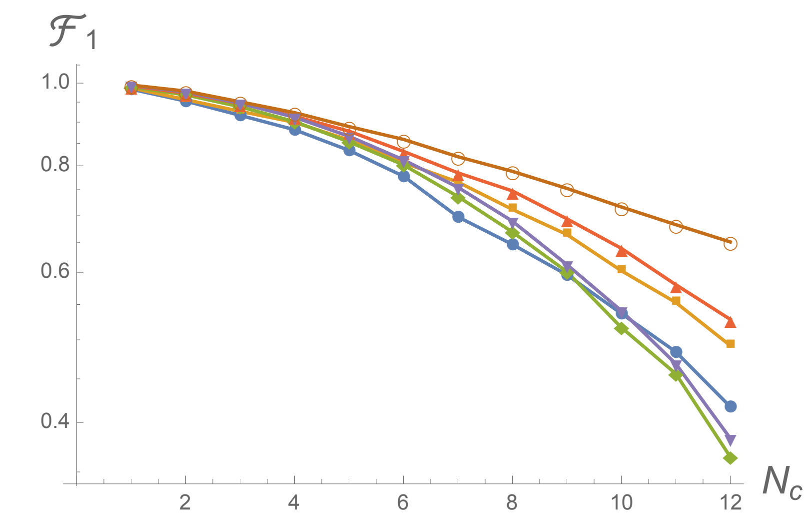

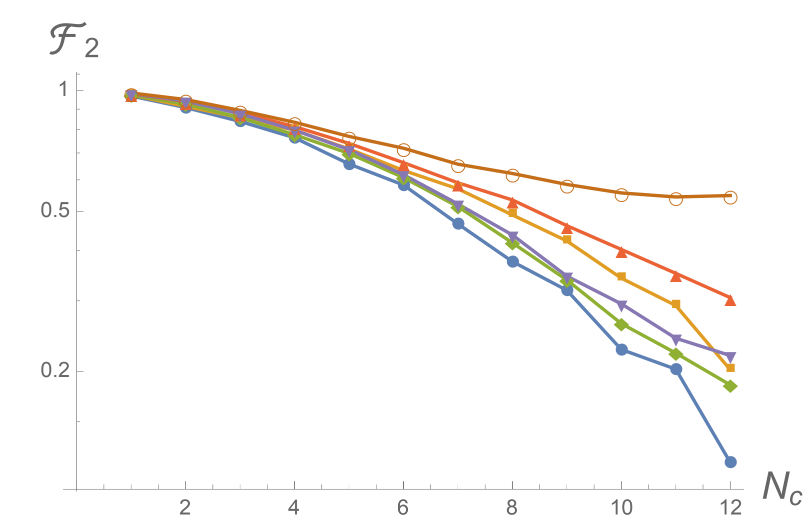

FIDELITY LOSS FROM QUBIT ERROR

- Unlike random circuits, in both unitary and lossy protocols, a single phase error does not send $\mathcal{F}$ to zero.
- However, it *does* appear to decay exponentially in the # of phase errors.
- Figures: top is 1 or 2 Z errors (by end of 12 cycles, averaged over position and time) in unitary (no cavities) protocol, middle is 1 Z error in noisy protocol, bottom is 2 Z errors in noisy protocol
- Reasons for this are not entirely clear, though I have some hypotheses— talk to me afterwards if interested.
- **Key: L=4 (blue), 5 (gold), 6 (green), 7(red), 8 (purple), 9 (brown)**

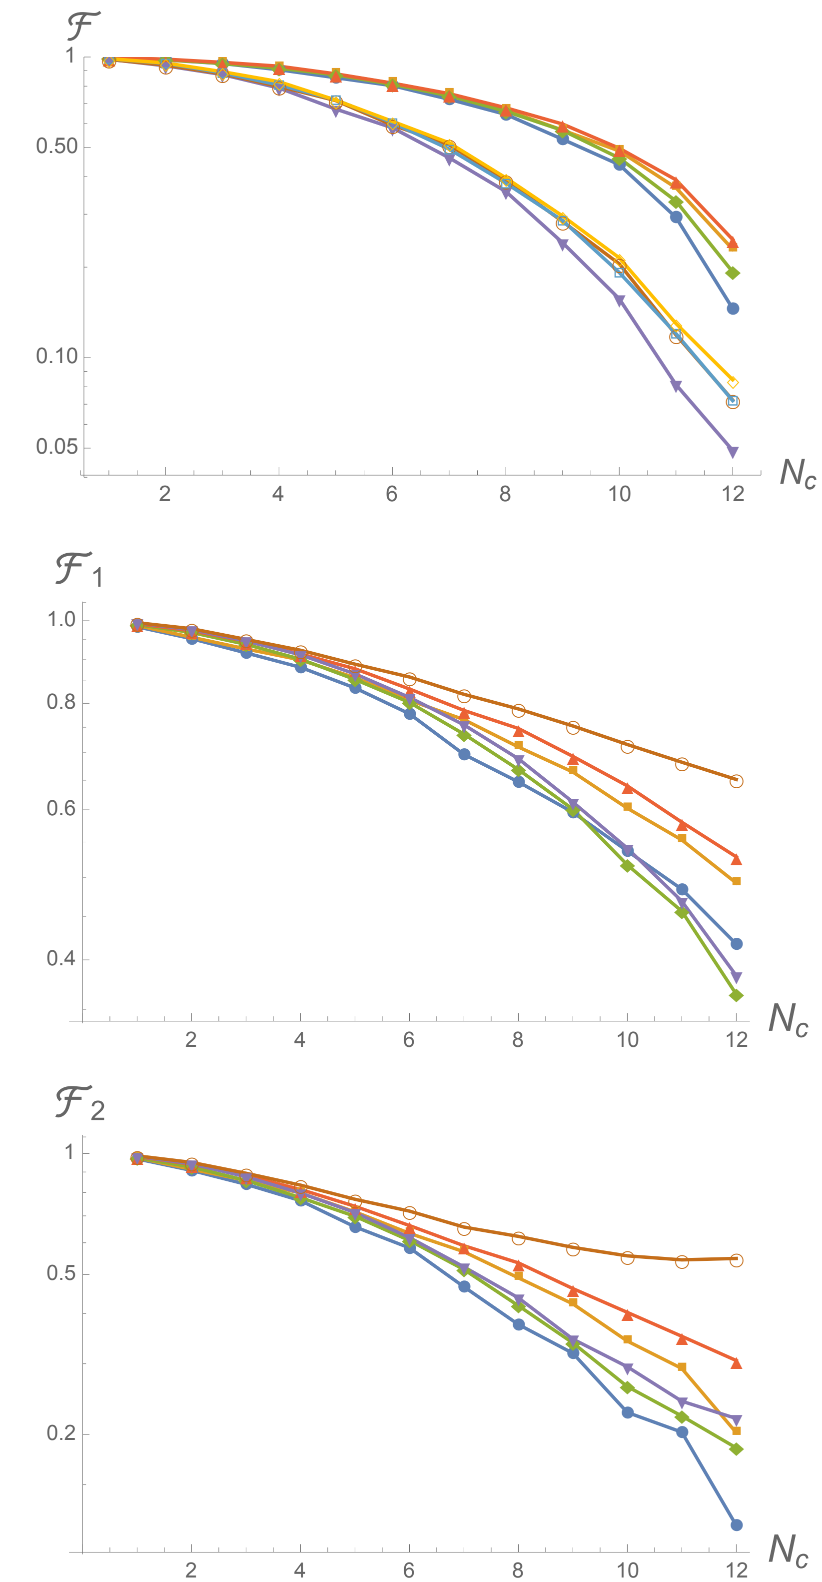

DIFFICULTY ESTIMATES

DIFFICULTY ESTIMATES

- Assumption: direct Schrodinger evolution w/ trajectory averaging is the most efficient simulation method

- Volume entanglement: matrix product state representation is extremely inefficient, with runtime scaling:

$$T_{MPS} \simeq 12 L d^3 \mathcal{N}^3$$

- Very "deep" circuit: for parameters studied, 6-8 cycles equiv to depth 42-56 in RQC. Methods which scale exp. in tree width will fail here.

- This does not rule out other methods which would scale better, but I'm not aware of any.

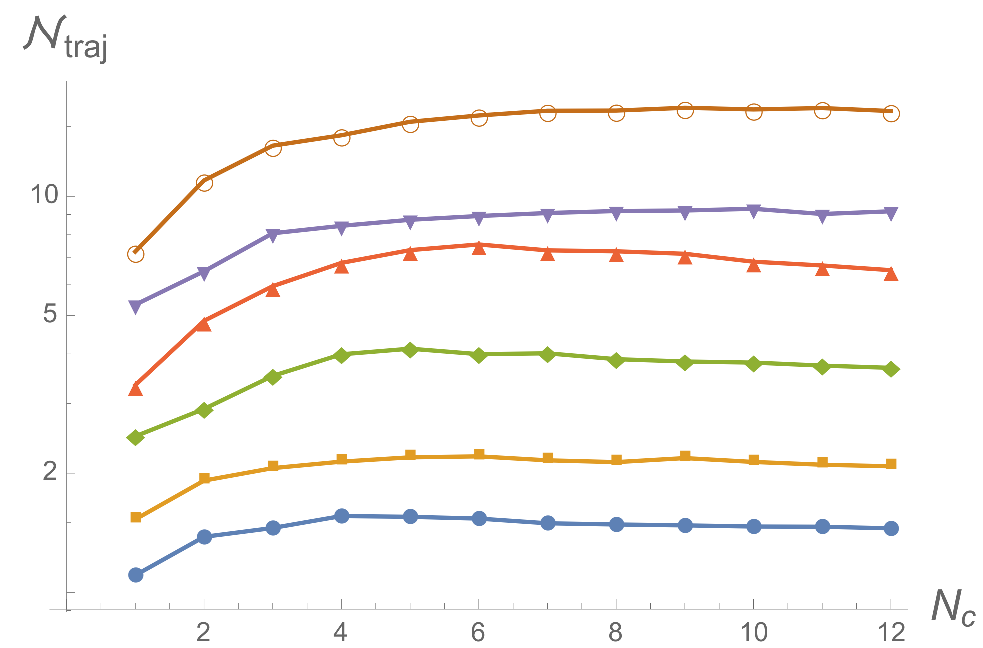

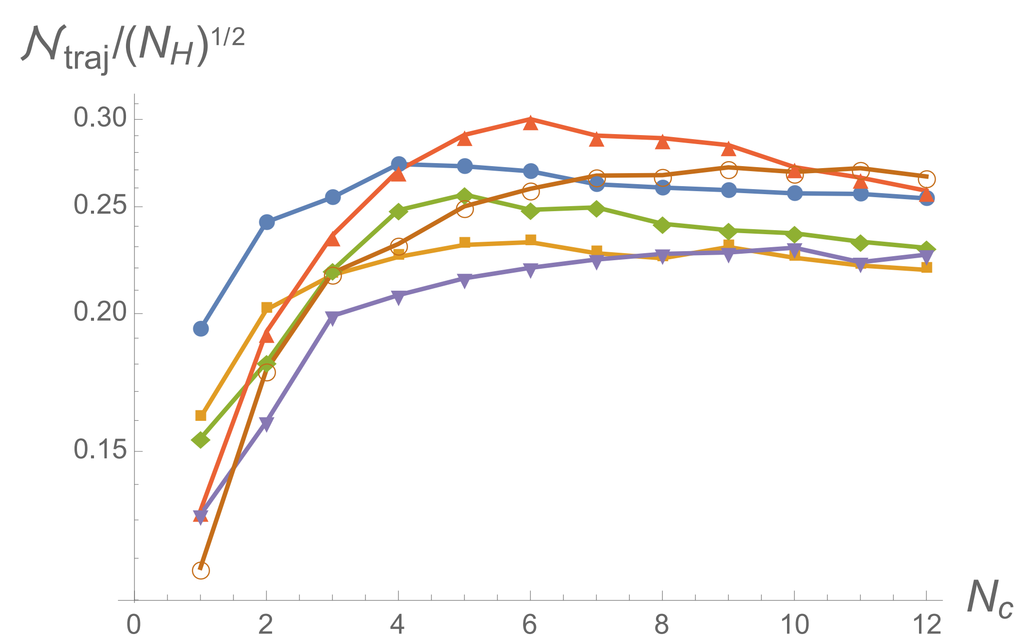

DIFFICULTY ESTIMATES

- In direct Schrodinger evolution, substantial savings are possible by truncating basis (see *Neill et al Science 2018* for estimates this draws from).
- Cap max # of double/triple occupancies in transmons
- Can throw out configurations which are far in total photon # from initial state
- Finally, cavities are lossy- exclude double occupancies in cavities, and cap max # of total cavity photons. Empirically a maximum of ~ L/6 photons seems sufficient (for these parameters, rounded to nearest int; this is protocol dependent). Right: fidelity loss from truncating from 2 max cavity photons to 1

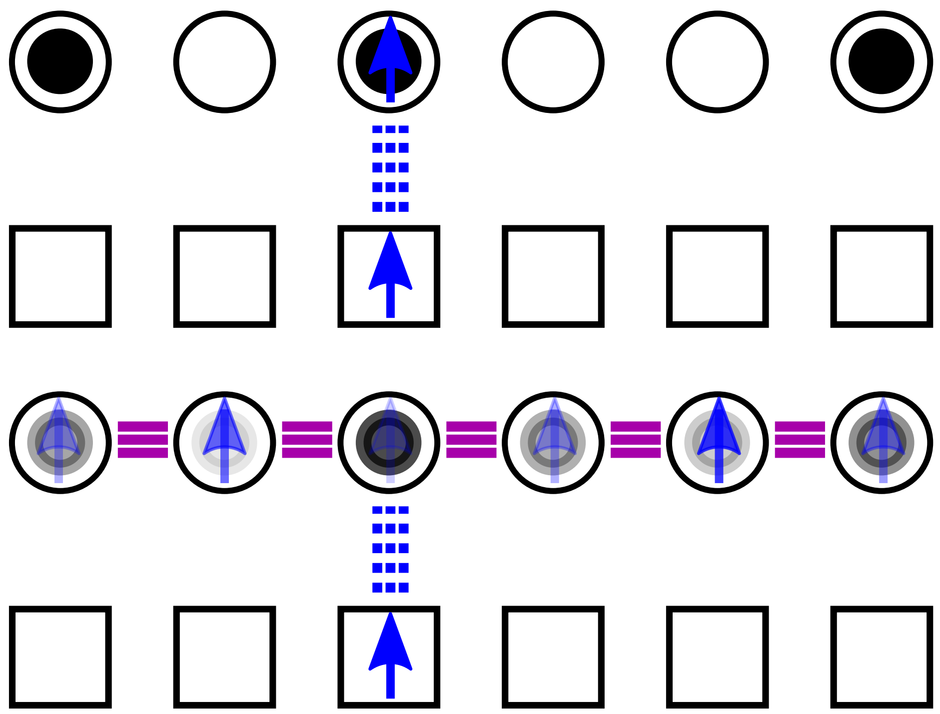

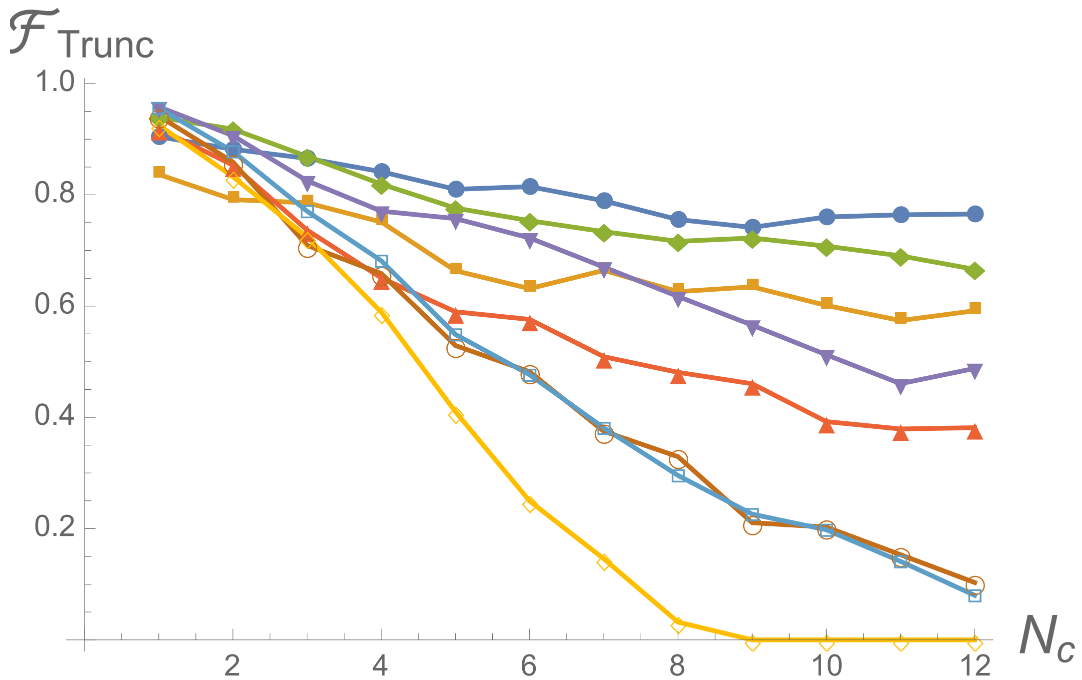

DIFFICULTY: MEMORY COSTS (FULL WFN)

- Estimated memory costs (GB) for having to keep either {2,1} or {3,2} doublons/triplons, and up to $\left\{\frac{L}{10}, \frac{L}{8}, \frac{L}{6}, \frac{L}{5}, \frac{L}{4}\right\}$ resonator photons (rounded to nearest int.)
- **Exascale** reached by L=30 in most cases.

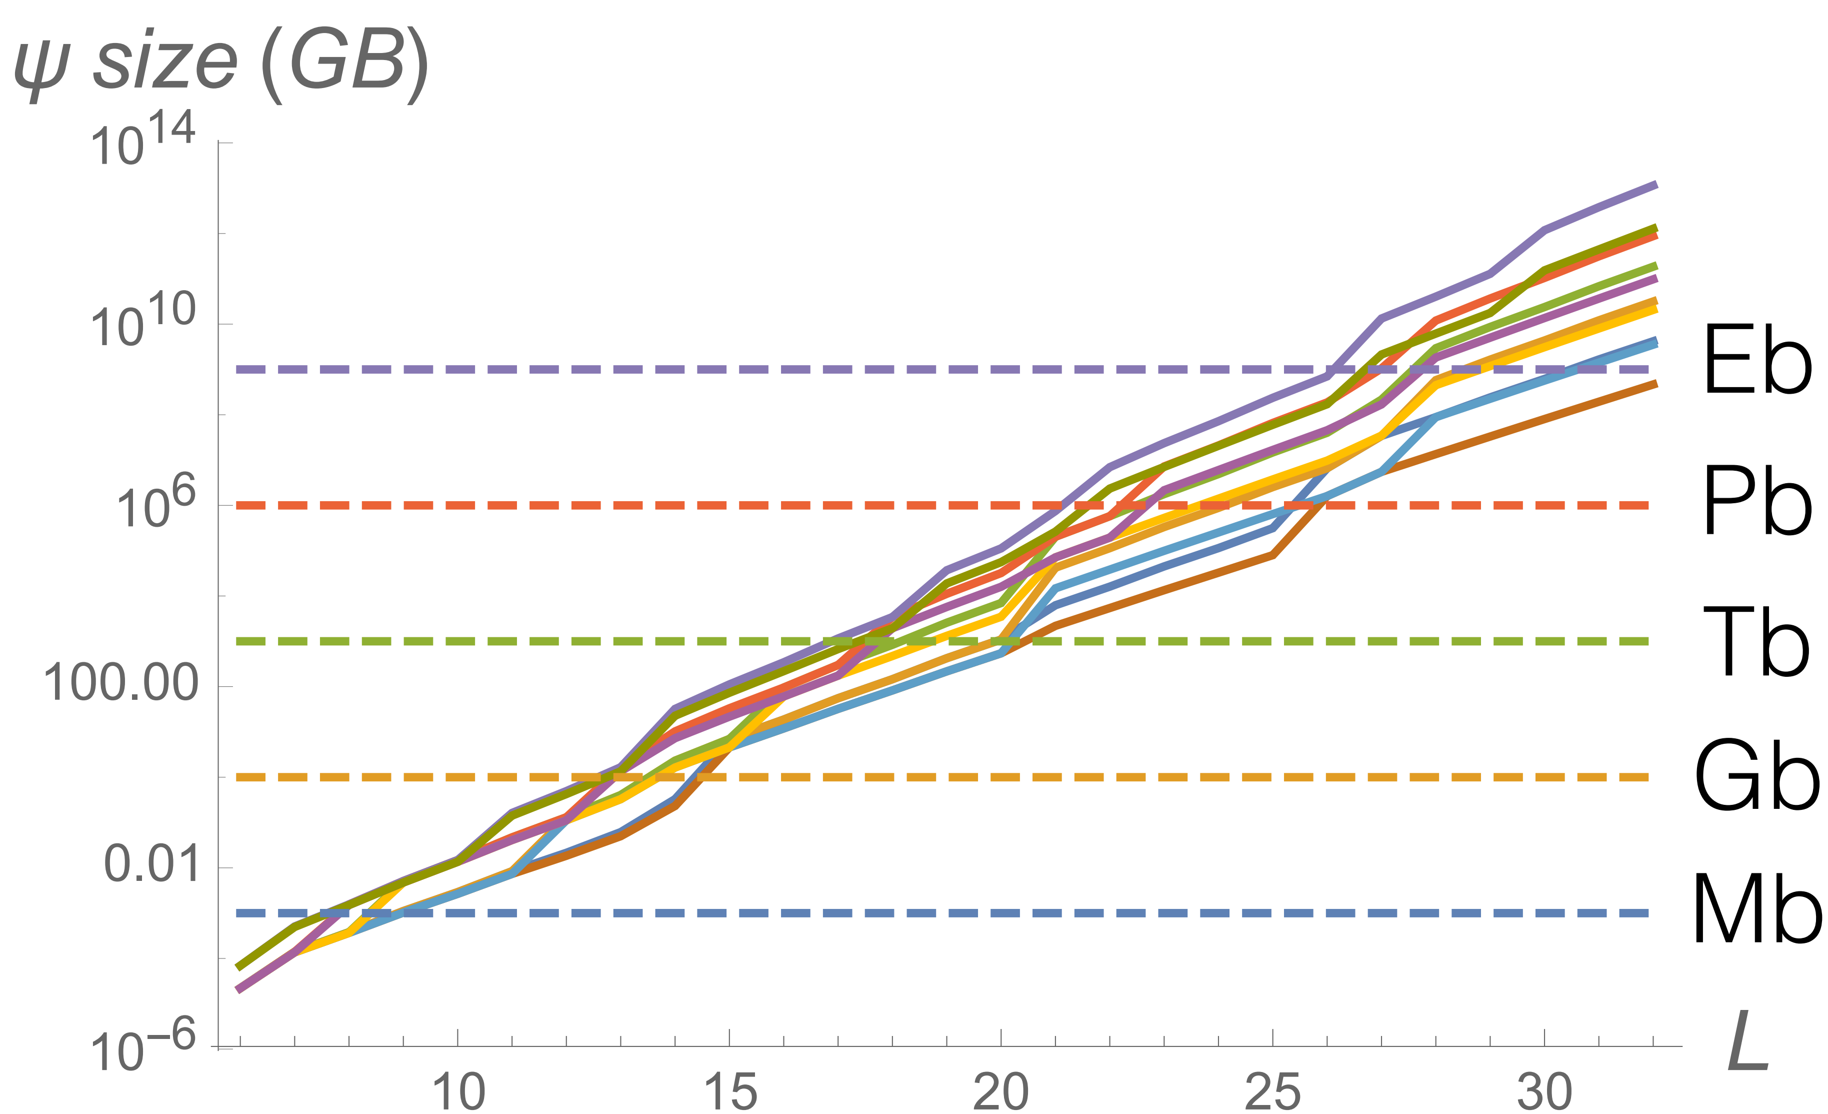

DIFFICULTY ESTIMATES: RUNTIME

- Runtime per-trajectory (assume RK4 timestep $\propto 1/L$), for a total of N_c cycles, scales as:

  $$T \propto L^2 N_H \times N_c$$

- Emprically, # of trajectories for a KL div <0.01 is $N_t \sim 3LN_c$. Error scales linearly in $1/N_t$.

- To get $\mathcal{F} \sim 0.25$, a few dozen trajectories likely needed at edge of tractability. Potentially significant!

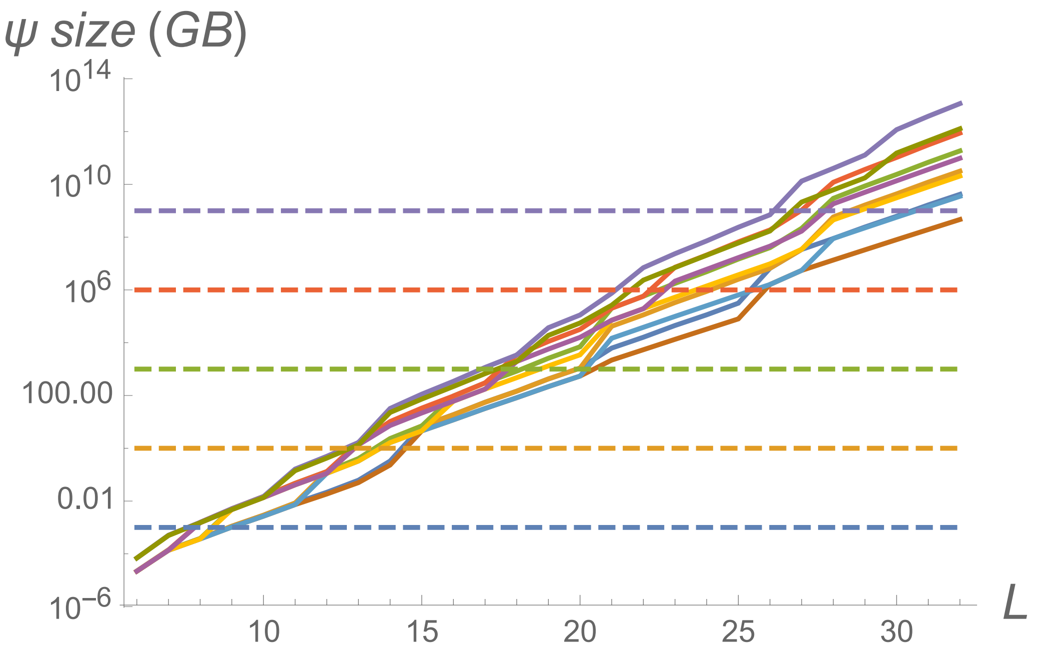

DIFFICULTY ESTIMATES: PROJECTED FIDELITY

- Projected fidelities in the supremacy regime are potentially much higher than RQCs

- Rough estimate from Google: SPAM error is ~ 0.03/qubit, and in the 9-qubit chain experiment, phase/control error was ~0.004 per qubit per cycle.

- For a 27 qubit grid w/ 6-8 cycles, using these values we find:

$$\mathcal{F} \sim 0.230 - 0.185$$

- Dominated by SPAM error (which largely can't be post-selected out). Order of magnitude higher than typical RQC targets, though this is a very rough estimate.

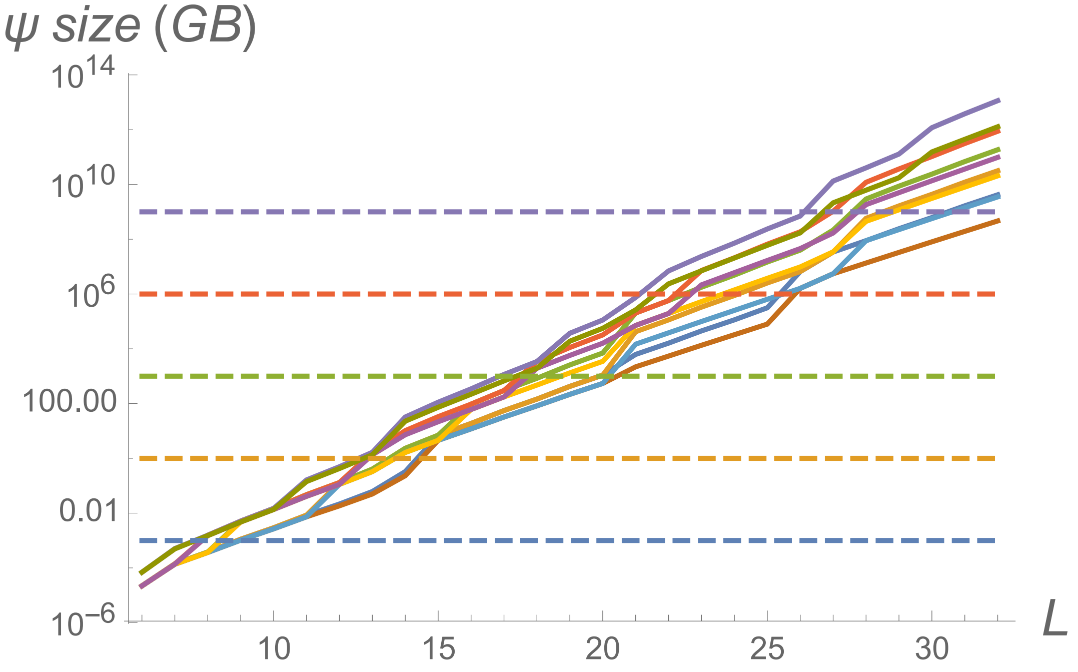

CONCLUSIONS

- Noisy sampling problems are good potential candidates for demonstrating a quantum advantage in current hardware: **classically much harder, but "quantum-easier"** due to smaller system sizes
- **Key idea:** resonantly couple lossy subsystem to propagating modes in continuous delocalized evolution
- **No complexity-theoretic proof, but: satisfies all hardness criteria computed so far** for unitary gmon chain; Volume-law entanglement should persist to classically intractable sizes
- Assuming no hidden simplifications, simulation *likely* becomes **intractable in the high 20s of qubit-cavity pairs**
- Could be adapted to other setups! "Supremacy" capable circuits (which would be too small for such a claim in other protocols) may already be in fridges as we speak…
- Thanks to the Google team, for support and many fascinating discussions: Sergio Boixo, Yu Chen, John Martinis, Charles Neill, Pedram Roushan and Vadim Smelyanskiy

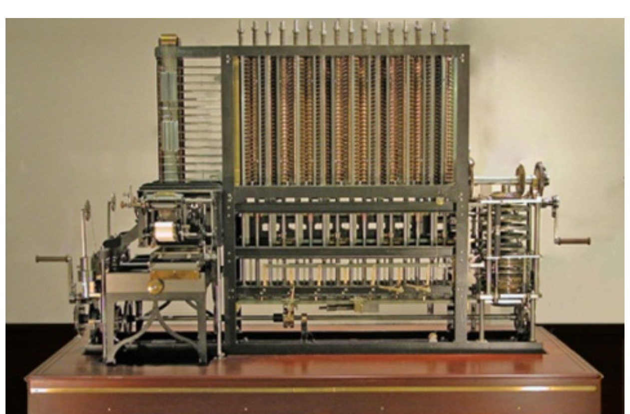

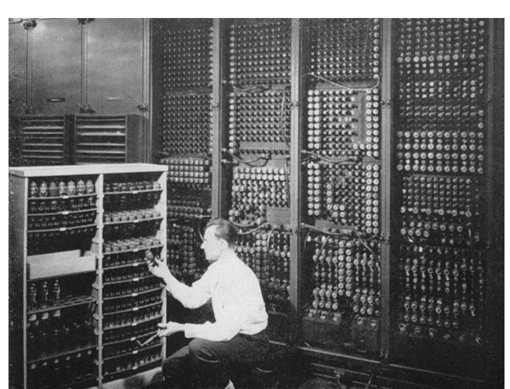

Replacing a bad tube meant checking among ENIAC's 19,000 possibilities.

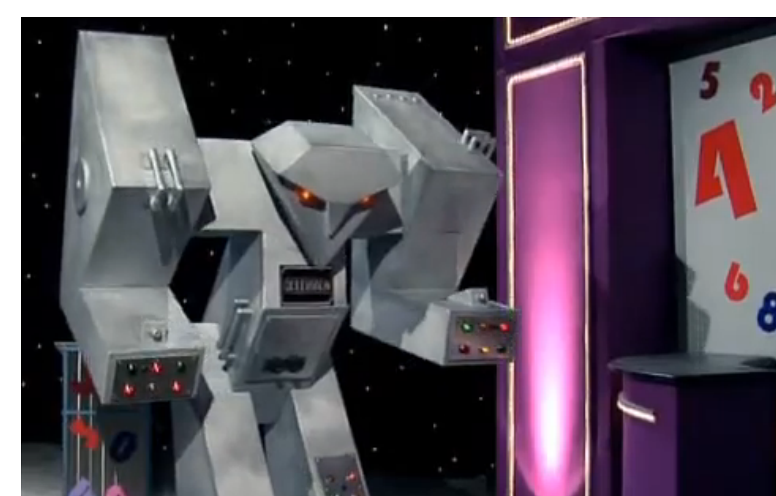

OPEN QUESTIONS

- Does an efficient classical method exist for sampling the output of an evolving delocalized system with colored noise? A general approach for all systems almost certainly does not exist, but are there special cases where the solution can be found? If so, what are they, and why?

- Given realistic qubit noise, can we simulate thermal states of matter using superconducting circuits with engineered dissipation? Proposals to do this exist (Hafezi et al, Shabani and Neven, both 2015) but it's not entirely clear how they respond to loss/dephasing.

- Engineered noise can passively correct errors, prepare/stabilize quantum states, make sampling problems harder, and more. Can noise, in analog quantum computing models, provide a quantum speedup in solving classical problems?

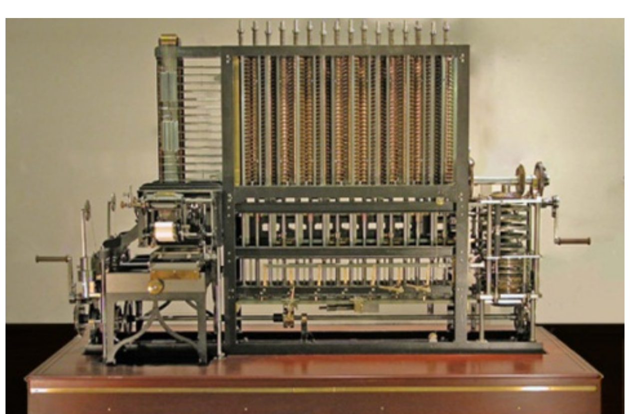

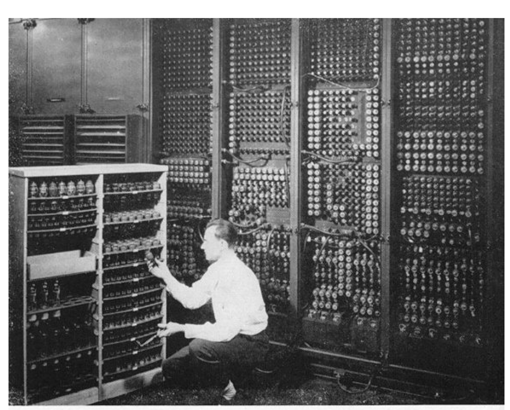

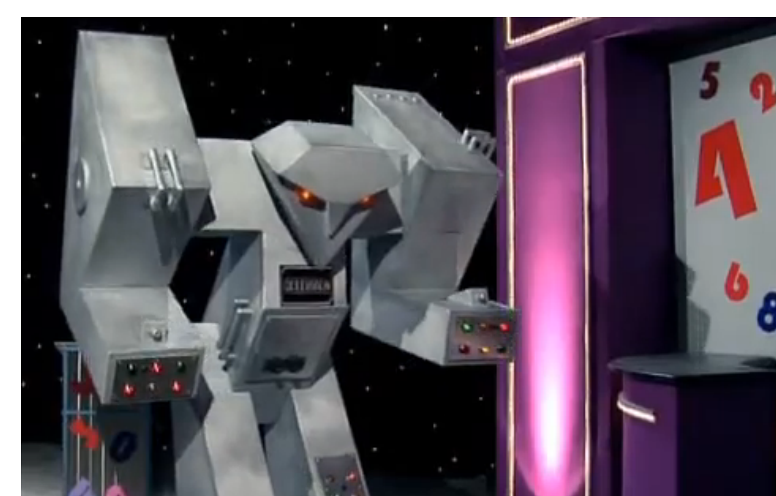

LOSSY SYSTEMS ADDENDUM

- **Objection!** Lossy systems have bounded correlations; max length scale is set by Lieb-Robinson info velocity and incoherent "error" rate (Poulin 2010; Barthel and Kliesch 2012; Huang and Guo 2017; Zhang, Huang and Chen 2018; others).

- Bounded entanglement suggests problem could be asymptotically "easy" (would permit an MPS description linear in system size, if exponential in correlation length)

- However, for realistic parameters, this length scale can still be too large for any conceivable classical computer to simulate. Increases with decreasing noise.

- What is this length scale?

LOSSY SYSTEMS ADDENDUM

- **Worst case:** consider an L-site chain, where info propagates at velocity v. Assume *uncorrelated*, Markovian errors occur at a rate Γ_R at each site. The time to entangle one end of the chain with the other is $t = L/2v$ (factor of 2 from meeting in the middle).
- The average # of errors in this time is $L\Gamma_E t$
- Intuition from gate model: assume single error totally disentangles the state
- Avg. one error or less gives $L_{max} = \sqrt{\dfrac{2v}{\Gamma_E}}$

LOSSY SYSTEMS ADDENDUM

- **Worst case:** uncorrelated error, single error disentangles state, ballistic info transport (see bound in Zhang/Huang/Chen 2018):

$$L_{max} = \sqrt{\frac{2v}{\Gamma_E}}$$

- For gmons with resonator noise? iSWAP time ~ 7ns, so v ~ 141 MHz. Error rate is $\langle n_{res} \rangle \Gamma_R; \quad \Gamma_R \simeq 10\text{MHz}$

- For protocols considered here, $\langle n_{res} \rangle \sim 0.05 - 0.1$

- Suggests $L_{max} \sim 17 - 24$

- This may actually be enough for a supremacy claim!

LOSSY SYSTEMS ADDENDUM

- **Worst case:** uncorrelated error, single error disentangles state, ballistic info transport:

$$L_{max} = \sqrt{\frac{2v}{\langle n_{res}\Gamma_R\rangle}} \sim 17 - 24$$

- **However:** Uncorrelated errors and single error disentanglement are bad assumptions here! Correlations could have much longer range

- **Key idea:** qubit-resonator coupling must be turned on simultaneously with qubit-qubit couplers

RANGE OF ENTANGLING QUBIT-RESONATOR OPERATION

- **Fundamental to this scheme:** from the point of view of the qubit chain, interaction with resonators is highly nonlocal

- **Simplified picture:** ignore interactions and disorder. Consider interaction with a single resonator at site *k*; total Hamiltonian is:

$$H = H_{Qubits}(t) + \Omega(t)\left[a_k^\dagger a_{Rk} + a_{Rk}^\dagger a_k\right] + \Delta_{Rk} a_{Rk}^\dagger a_{Rk}$$

- **Resonance condition:** photon can only be added/removed from qubits if total energy change $< \min\{\Omega, \Gamma_R\}$

- Assume throughout this talk $\Gamma_R = 10\mathrm{MHz}$

RANGE OF ENTANGLING QUBIT-RESONATOR OPERATION

$$H = H_{Qubits}(t) + \Omega(t)\left[a_k^\dagger a_{Rk} + a_{Rk}^\dagger a_k\right] + \Delta_{Rk} a_{Rk}^\dagger a_{Rk}$$

- A photon only has a significant chance of being transferred to the resonator if
- $|\delta_{EQ} - \Delta_{Rk}| < \min\{\Omega, \Gamma_R\}$
- If this is satisfied, resonator is only coupled to a **propagating mode**, with approximately equal weight over entire lattice
- Subsequent resonator photon loss is measuring a highly nonlocal operator! Will not immediately disentangle system.
- Limited by mode splitting, approx. $5.8g/L$ for a 1d chain, near center of band
- For $g_{max} = 2\pi \times 40\text{MHz}$, we get $\Gamma_R = \frac{1}{2}\frac{5.8 g_{max}}{L_{max}}; \quad L_{max} \simeq 72$