Approximate autonomous quantum error correction with Reinforcement Learning

Clemens Gneiting

Y. Zeng, Z.Y. Zhou, E. Rinaldi, CG, F. Nori, *PRL 131, 050601 (2023)*

QTML 2023, CERN, Nov 24, 2023

Outline

- Quantum computing, noise, qantum error correction
- *Autonomous* quantum error correction
- Code space optimization using reinforcement learning
- Surpassing break even with the RL code
- An experimental proposal

Quantum computing, noise, qantum error correction

Quantum computation

- Quantum computation promises drastic speedup and scaling advantages for specific tasks: factorization, search, quantum simulation, **QML**, etc.

- basic building blocks: **qubits:** either natural or artificial two-level quantum systems

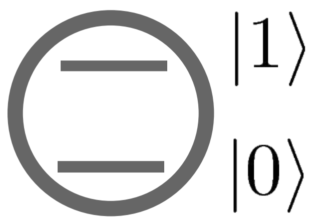

$$|\psi_{\theta\phi}\rangle = \cos\frac{\theta}{2}|0\rangle + e^{i\phi}\sin\frac{\theta}{2}|1\rangle$$

- promising platforms: superconducting qubits, trapped ions, quantum dots, photons, etc.

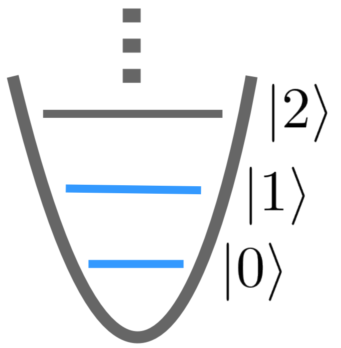

bosonic

- different computational paradigms to implement algorithms: unitary gate-based, measurement-based, adiabatic, dissipative, etc.

- **in this talk:** bosonic qubits, unitary gate-based quantum information processing

adressability requires nondegenerate levels

Noise

- coherent (parameter drifts etc.) and incoherent (uncontrolled entanglement with environment) noise deteriorates the proximity to the target state

  ➡ quantum advantage is rapidly lost!

- time-continuous description

$$\dot{\rho} = -\frac{i}{\hbar}[\hat{H}(t), \rho] + \gamma \sum_{j=1}^{N} \left(\hat{L}_j \rho \hat{L}_j^\dagger - \frac{1}{2}\{\hat{L}_j^\dagger \hat{L}_j, \rho\} \right)$$ ..Lindblad master equation

$\hat{H}(t)$ ..quantum algorithm

$\hat{L}_j$ ..error channels

generic qubit error channels: spin flip, dephasing, ... $\hat{L}_j = \hat{\sigma}_{x/y/z}$

dominant **bosonic** error channel: photon loss, $\hat{L} = \hat{a}$

Quantum error correction (QEC)

Define "logical qubit" $|0_\mathrm{L}\rangle, |1_\mathrm{L}\rangle$ in a larger Hilbert space, such that errors neither erase nor distort the stored quantum information

$$|\psi_{\theta\phi}\rangle = \cos\frac{\theta}{2}\,|0_\mathrm{L}\rangle + e^{i\phi}\sin\frac{\theta}{2}\,|1_\mathrm{L}\rangle$$

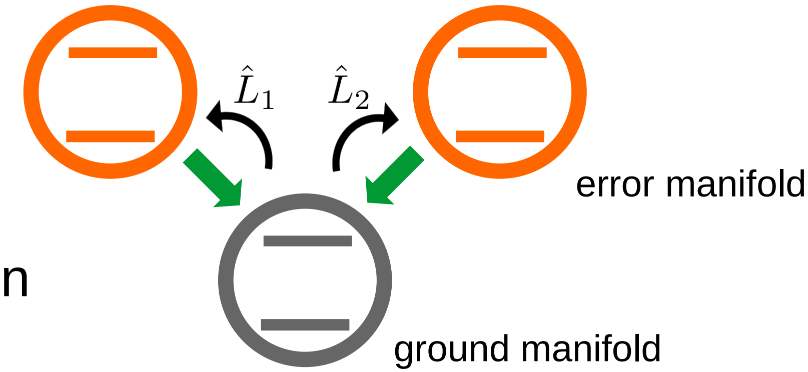

- guaranteed by Knill-Laflamme (KL) condition

$$\langle u_L|\hat{E}_i^\dagger \hat{E}_j|v_L\rangle = \alpha_{ij}\delta_{uv} \qquad u, v \in \{0, 1\}$$

$$\hat{E}_i \in \{\hat{I}, \hat{L}_1, \ldots, \hat{L}_N\}$$

implies for bosonic codes: $\langle 0_L|\hat{a}^\dagger\hat{a}|0_L\rangle = \langle 1_L|\hat{a}^\dagger\hat{a}|1_L\rangle$

requires that at least one code state is a superposition of Fock states!

- errors are corrected via syndrome **measurements** and adapted **feedback** of (unitary) correction operators

Quantum error correction (QEC)

How to achieve redundancy in state space

many physical qubits

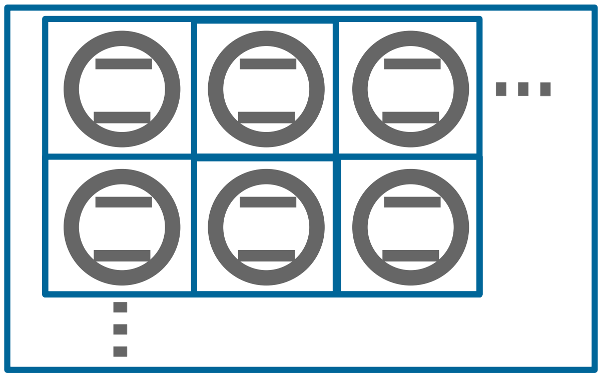

set of n physical qubits provides redundant Hilbert space with dimension $N = 2^n$

Examples: Steane code, surface code, color code, etc.

bosonic system

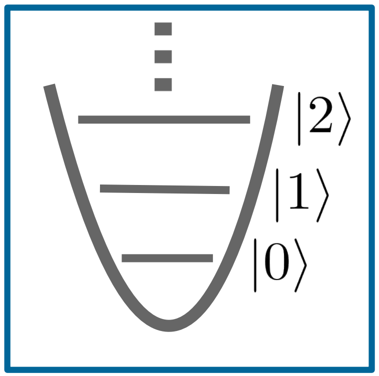

oscillator states up to $|N\rangle$ provide redundant Hilbert space equivalent to a set of $\log_2 N$ qubits

Examples: Gottesman-Kitaev-Preskill code, cat code, binomial codes, etc.

QEC in *bosonic* systems minimizes overhead, as only one nonlinear element needed for operations

- challenge: larger Hilbert space introduces more errors!

How to measure the effectiveness of quantum error correction

- average fidelity $$\bar{F}(t) = \frac{1}{4\pi} \int_\Omega F(\theta, \phi, t)\ \mathrm{d}\Omega$$

$$F(\theta, \phi, t) = \mathrm{Tr}[\rho_{t_0}(\theta, \phi)\rho_t(\theta, \phi)]$$

$$|\psi_{t_0}\rangle = \cos\frac{\theta}{2}\,|0_\mathrm{L}\rangle + e^{i\phi}\sin\frac{\theta}{2}\,|1_\mathrm{L}\rangle$$

short-time expansion: $$\overline{F}(\delta t) = 1 - \frac{1}{2}\gamma_{\mathrm{err}}\delta t$$

overall decay rate of the average fidelity gives the error rate

- ratio of error rate with and without error correction: gain G

without error correction refers to the physical or “natural” qubits

the larger G the better; **break even:** $G = 1$

Bosonic codes

dominant error channel: photon loss $\hat{L} = \hat{a}$

Example: binomial code

$$|0_L\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |4\rangle)$$

$$|1_L\rangle = |2\rangle$$

KL: $\hat{a}|0_L\rangle \propto |3\rangle$ $\hat{a}|1_L\rangle \propto |1\rangle$

$$\langle 0_L|\hat{a}^\dagger\hat{a}|0_L\rangle = \langle 1_L|\hat{a}^\dagger\hat{a}|1_L\rangle = 2$$

parity measurement detects errors

M. Michael et al, PRX 6, 031006 (2016)

Example: Gottesman-Kitaev-Preskill (GKP) code

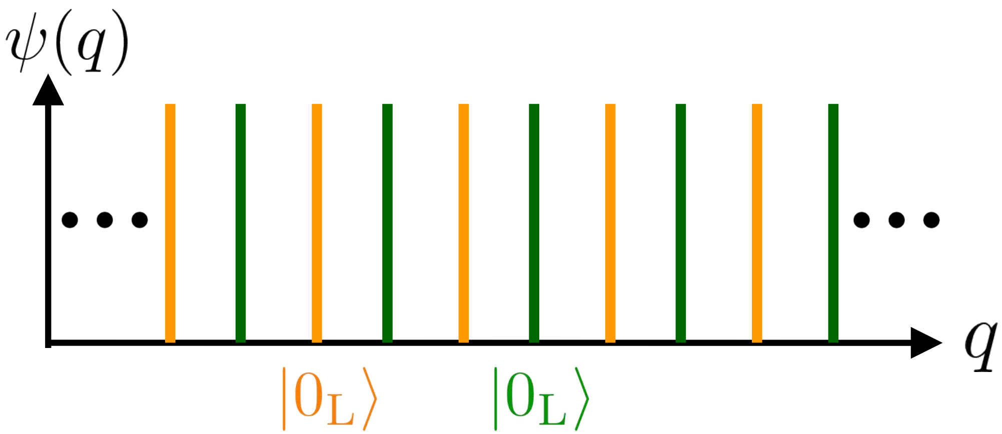

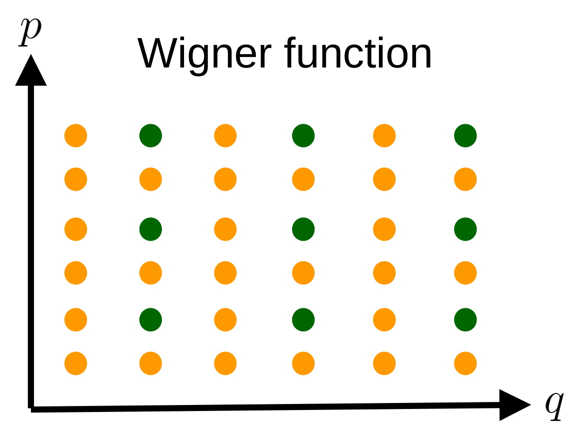

designed to protect against (small) phase-space displacements,
but works also well for photon loss!

- experimentally realized including error correction V.V. Sivak et al, Nature 616, 50 (2023)

gain: $G \approx 2.3$

Autonomous
quantum error correction

Autonomous quantum error correction

Idea: protect logical qubits by engineered dissipation, avoiding the necessity of frequent, error-prone measurement-feedback loops

$$\dot{\rho} = -\frac{i}{\hbar}[\hat{H}(t), \rho] + \gamma \sum_{j=1}^{N} \mathcal{D}[\hat{L}_{\mathrm{nat},j}] + M\gamma \sum_{k=1}^{M} \mathcal{D}[\hat{L}_{\mathrm{eng},k}] \qquad M \gg 1$$

$$\mathcal{D}[\hat{L}] = 2\hat{L}\rho\hat{L}^\dagger - \hat{L}^\dagger\hat{L}\rho - \rho\hat{L}^\dagger\hat{L}$$

- engineered jump operators pump the corrupted state back from the error space into the code space (before another error can happen)

$$\hat{L}_{\mathrm{eng}} = |0_{\mathrm{L}}\rangle\langle 0_{\mathrm{er}}| + |1_{\mathrm{L}}\rangle\langle 1_{\mathrm{er}}|$$

- KL condition still needs to be satisfied for full QEC
- conditioned time evolution in between error jumps requires separate correction mechanism
- experimental realization with cat qubit: J. Gertler et al, Nature 590, 243 (2021)

Example: binomial code

- code words: $|0_L\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |4\rangle)$ KL: $\hat{a}|0_L\rangle \propto |3\rangle$ $\hat{a}|1_L\rangle \propto |1\rangle$

  $|1_L\rangle = |2\rangle$ $\langle 0_L|\hat{a}^\dagger\hat{a}|0_L\rangle = \langle 1_L|\hat{a}^\dagger\hat{a}|1_L\rangle = 2$

- engineered jump operators:

$\hat{L}_{\mathrm{eng},1} = \frac{1}{\sqrt{2}}(|0\rangle + |4\rangle)\langle 3| + |2\rangle\langle 1|$ and $\hat{L}_{\mathrm{eng},2} = |0_\mathrm{L}\rangle(\langle 0| - \langle 4|)$

without $\hat{L}_{\mathrm{eng},2}$ with $\hat{L}_{\mathrm{eng},2}$

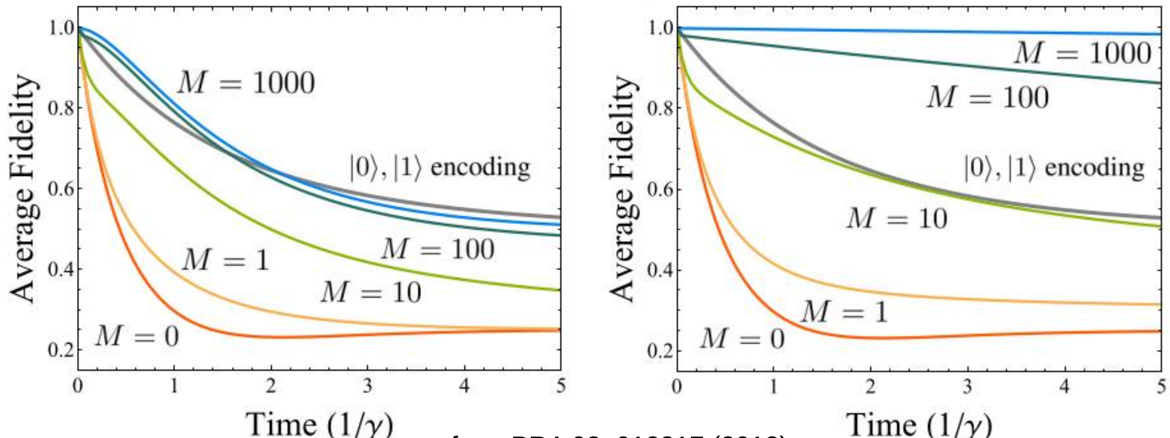

from PRA 98, 012317 (2018)

requires at least two jump operators for AQEC to work well!

J.-M. Lihm, K. Noh, U. Fischer, PRA 98, 012317 (2018)

Example: √3 code

- code words:

discovered by automatized search algorithm

$$|\psi_0\rangle = \sqrt{1-\frac{1}{\sqrt{3}}}\,|0\rangle + \frac{1}{\sqrt[4]{3}}\,|3\rangle$$

$$|\psi_1\rangle = \sqrt{\frac{2(6-\sqrt{3})}{\sqrt{3}+9}}\,|1\rangle - \sqrt{\frac{(\sqrt{3}-1)(6-\sqrt{3})}{2(\sqrt{3}+9)}}\,|4\rangle + \sqrt{\frac{3-\sqrt{3}}{2(\sqrt{3}+9)}}\,|6\rangle$$

KL: $\langle\psi_0|\hat{a}^\dagger\hat{a}|\psi_0\rangle = \langle\psi_1|\hat{a}^\dagger\hat{a}|\psi_1\rangle = \sqrt{3}$

- single engineered jump operator, realized through coupling to auxiliary qubit:

$$D[|\psi_0\rangle\langle\psi_2| + |\psi_1\rangle\langle\psi_3| + \cdots], \quad \text{where} \quad |\psi_2\rangle \propto \hat{a}\,|\psi_0\rangle$$

$$|\psi_3\rangle \propto \hat{a}\,|\psi_1\rangle$$

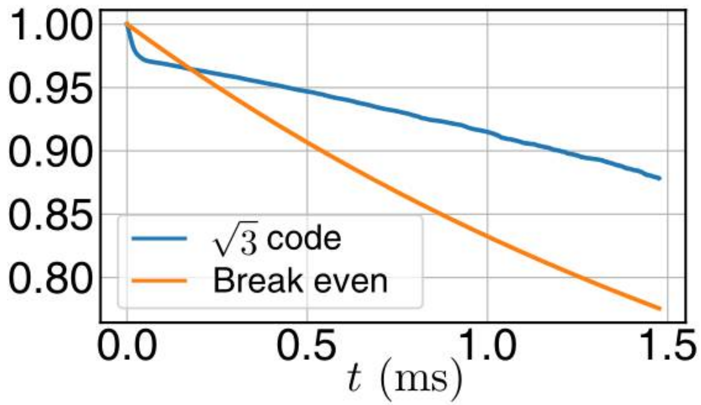

from PRX Quantum 3, 020302 (2022)

requires at least Hamiltonian distance $d = 2$, higher-order nonlinearity needs to be decomposed into 4 control fields that introduce additional noise

Z. Wang et al, PRX Quantum 3, 020302 (2022)

Example: √3 code

- code words:

discovered by automatized search algorithm

$$|\psi_0\rangle = \sqrt{1-\frac{1}{\sqrt{3}}}\,|0\rangle + \frac{1}{\sqrt[4]{3}}\,|3\rangle$$

$$|\psi_1\rangle = \sqrt{\frac{2(6-\sqrt{3})}{\sqrt{3}+9}}\,|1\rangle - \sqrt{\frac{(\sqrt{3}-1)(6-\sqrt{3})}{2(\sqrt{3}+9)}}\,|4\rangle + \sqrt{\frac{3-\sqrt{3}}{2(\sqrt{3}+9)}}\,|6\rangle$$

KL: $\langle\psi_0|\hat{a}^\dagger\hat{a}|\psi_0\rangle = \langle\psi_1|\hat{a}^\dagger\hat{a}|\psi_1\rangle = \sqrt{3}$

- single engineered jump operator, realized through coupling to auxiliary qubit:

$$D[|\psi_0\rangle\langle\psi_2| + |\psi_1\rangle\langle\psi_3| + \cdots], \quad \text{where} \quad |\psi_2\rangle \propto \hat{a}\,|\psi_0\rangle$$

$$|\psi_3\rangle \propto \hat{a}\,|\psi_1\rangle$$

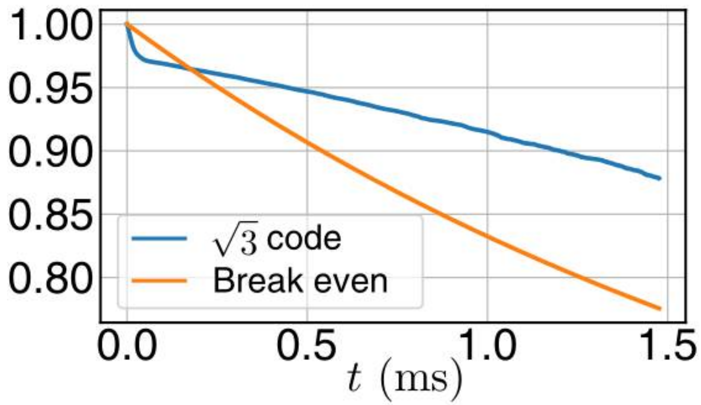

from PRX Quantum 3, 020302 (2022)

requires at least Hamiltonian distance $d = 2$, higher-order nonlinearity needs to be decomposed into 4 control fields that introduce additional noise

Z. Wang et al, PRX Quantum 3, 020302 (2022)

Is more experiment-friendly autonomous QEC possible?

Code space optimization using reinforcement learning

Our ansatz

- relax the KL condition part that requires $\langle 0_L|\hat{a}^\dagger\hat{a}|0_L\rangle = \langle 1_L|\hat{a}^\dagger\hat{a}|1_L\rangle$

$$|0_\mathrm{L}\rangle = \sum_{n=0} c_n^{(0)}|4n\rangle$$

$$|1_\mathrm{L}\rangle = \sum_{n=0} c_n^{(1)}|4n+2\rangle$$

the coefficients $c_n^{(0)}$ and $c_n^{(1)}$ are to be optimized

Fock-state superpositions not excluded!

- restrict us to a single engineered jump operator

$$L_\mathrm{eng} = L_\mathrm{o}\left\{\mathrm{Tr}\left[L_\mathrm{o}^\dagger L_\mathrm{o}\right]\right\}^{-1/2}$$

$$L_\mathrm{o} = |0_\mathrm{L}\rangle\langle 0_\mathrm{er}| + |1_\mathrm{L}\rangle\langle 1_\mathrm{er}| \qquad |u_\mathrm{er}\rangle = a|u_\mathrm{L}\rangle/\xi_u$$

- realization through coupling to auxiliary qubit:

$$\frac{d\rho}{dt} = -i[H_\mathrm{eff},\rho] + \frac{\gamma_a}{2}\mathcal{D}[a] + \frac{\gamma_b}{2}\mathcal{D}[\sigma_-]$$

$$H_\mathrm{eff} = g(L_\mathrm{eng}\sigma_+ + L_\mathrm{eng}^\dagger\sigma_-)$$

$$\gamma_a, g \ll \gamma_b$$

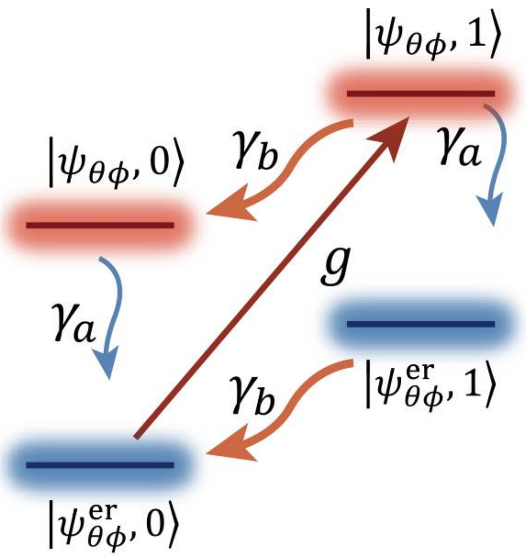

tracing out the qubit gives rise to
the desired engineered system dynamics

Code space optimization with reinforcement learning

- **goal:** find coefficients $[c_n^{(0)}, c_n^{(1)}]$ that maximize mean fidelity at $\gamma_a t = 0.6$
- simulate $\frac{d\rho}{dt} = -i[H_{\text{eff}}, \rho] + \frac{\gamma_a}{2}\mathcal{D}[a] + \frac{\gamma_b}{2}\mathcal{D}[\sigma_-]$ with $g/\gamma_a = 400$ and $\gamma_b/\gamma_a = 1750$

  for $|\psi_{t_0}\rangle = \cos\frac{\theta}{2}|0_\text{L}\rangle + e^{i\phi}\sin\frac{\theta}{2}|1_\text{L}\rangle$ until $\gamma_a t = 0.6$

realistic parameter choices that satisfy $\gamma_a, g \ll \gamma_b$

RL schematics

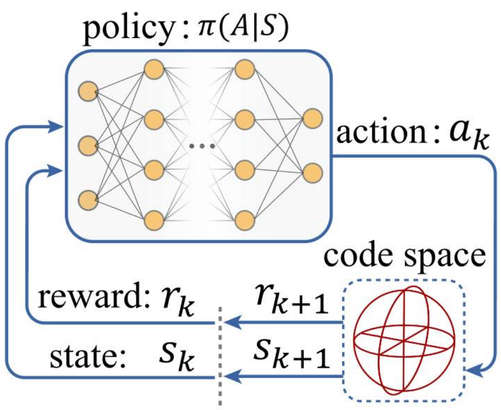

action: coefficient vector $[c_n^{(0)}, c_n^{(1)}]$

state: mean fidelity $\overline{F}(t)$

reward: difference between mean fidelity of code space and break even

policy function: decides which actions to take depending on state and reward

modeled by feedforward neural network

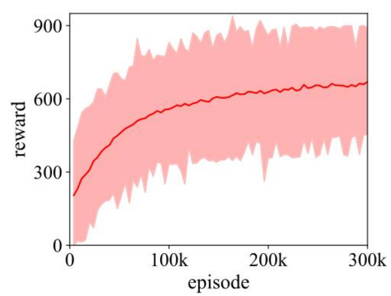

Code space optimization with reinforcement learning

- **goal:** find coefficients $[c_n^{(0)}, c_n^{(1)}]$ that maximize mean fidelity at $\gamma_a t = 0.6$
- simulate $\frac{d\rho}{dt} = -i[H_{\mathrm{eff}}, \rho] + \frac{\gamma_a}{2}\mathcal{D}[a] + \frac{\gamma_b}{2}\mathcal{D}[\sigma_-]$ with $g/\gamma_a = 400$ and $\gamma_b/\gamma_a = 1750$

  for $|\psi_{t_0}\rangle = \cos\frac{\theta}{2}|0_{\mathrm{L}}\rangle + e^{i\phi}\sin\frac{\theta}{2}|1_{\mathrm{L}}\rangle$ until $\gamma_a t = 0.6$

realistic parameter choices that satisfy $\gamma_a, g \ll \gamma_b$

some more details:

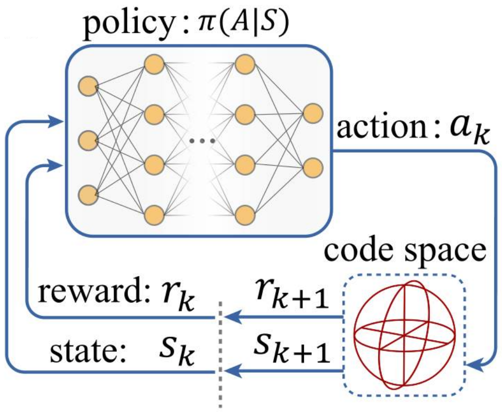

We start with vast search space up to $N = 40$ Fock states

RL allows us to gradually reduce the search space to relevant Fock states

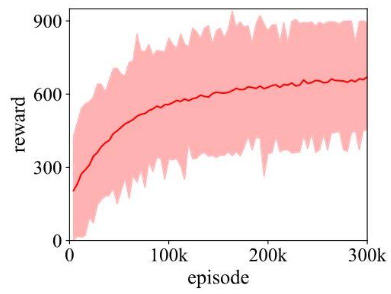

Training took 1-2 weeks on a cloud cluster

RL code

- RL finds surprisingly simple bosonic code:

  even though Fock-state superpositions were not excluded, they turn out to be not optimal

$$|0_L\rangle = |4\rangle \qquad \langle 0_L|\hat{a}^\dagger \hat{a}|0_L\rangle = 4$$

$$|1_L\rangle = |2\rangle \qquad \langle 1_L|\hat{a}^\dagger \hat{a}|1_L\rangle = 2$$

- Single engineered jump operator $L_{\text{eng}} \propto |2\rangle\langle 1| + |4\rangle\langle 3|$ realizable with Hamiltonian distance $d = 1$ (gates $d_g = 2$)

- Break even is well surpassed: the mean infidelity is about 0.17 times the break-even threshold

- Not all states are equally well protected, but all well beyond break even

state-dependent fidelity at $\gamma_a t = 0.6$

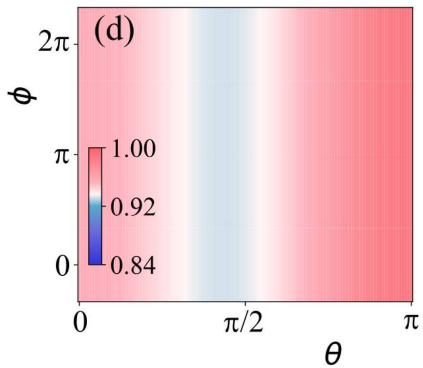

$$|\psi_{\theta\phi}\rangle = \cos\frac{\theta}{2}\,|0_\text{L}\rangle + e^{i\phi}\sin\frac{\theta}{2}\,|1_\text{L}\rangle$$

mean fidelity:

$$\bar{F}(t) = \frac{2}{3} + \frac{1}{3}\exp(-u\gamma_a t), \quad u = 3 - 2\sqrt{2} \approx 0.17$$

$$\rho_a(t) \approx \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \rho_{22}(0) & 0 & \rho_{24}(0)e^{-u\gamma_a t} \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \rho_{42}(0)e^{-u\gamma_a t} & 0 & \rho_{44}(0) \end{pmatrix}$$

effective residual dephasing

possible gain: $G \approx 5.8$

Comparison with other codes

- RL code outperforms binomial code and √3 code if there is a single corrective jump operator

all codes are corrected with
a single (optimal) jump operator

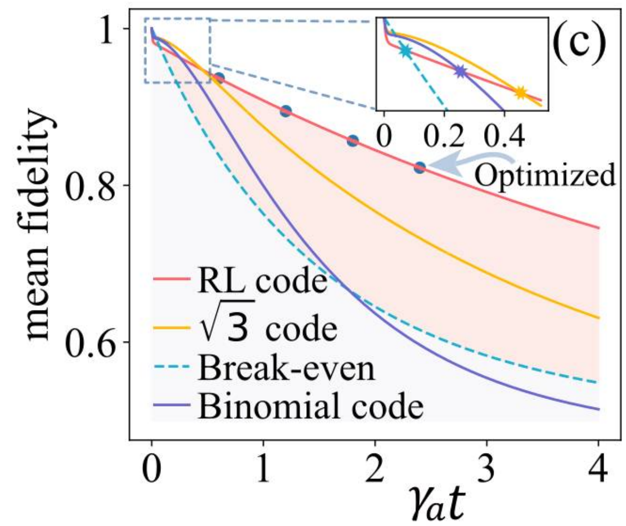

- binomial code and √3 code can perform better, however at the cost of second engineered jump operator

- RL code is optimized for *single* engineered jump operator, adding more of them yields no improvement

- short transition period where RL code performs poorer: corrective jump operator not yet effective while single-photon loss scales with mean photon number

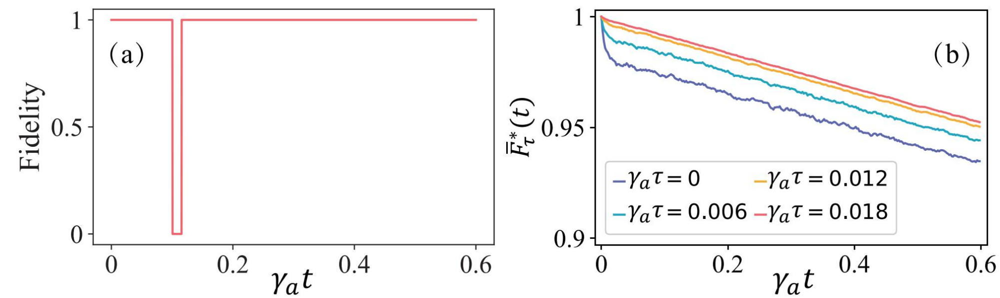

$$F_\tau(t^*) = \max_{t^* \in [t, t+\tau]} F(t^*)$$

Role of the engineered jump rate

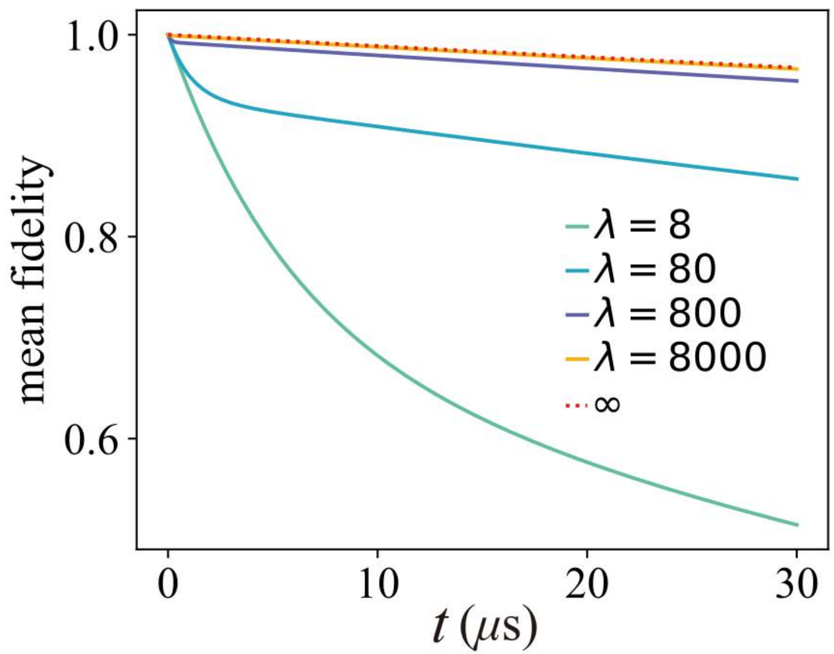

$$\frac{d\rho_a}{dt} = \frac{\gamma_a}{2}\mathcal{D}[a] + \frac{\gamma_a \lambda}{2}\mathcal{D}[L_{\text{eng}}]$$

$$\lambda = \frac{8|g|^2}{\gamma_b \gamma_a} = 8C$$

- increasing the engineered jump rate further reduces the fidelity decay
- a residual decay remains due to $\langle 0_L|\hat{a}^\dagger \hat{a}|0_L\rangle \neq \langle 1_L|\hat{a}^\dagger \hat{a}|1_L\rangle$
- initial dip can be arbitrarily mitigated

Optimality of the Fock-state pair

How does the performance change if we shift the code in Fock space?

code words $|m\rangle$ and $|m+2\rangle$

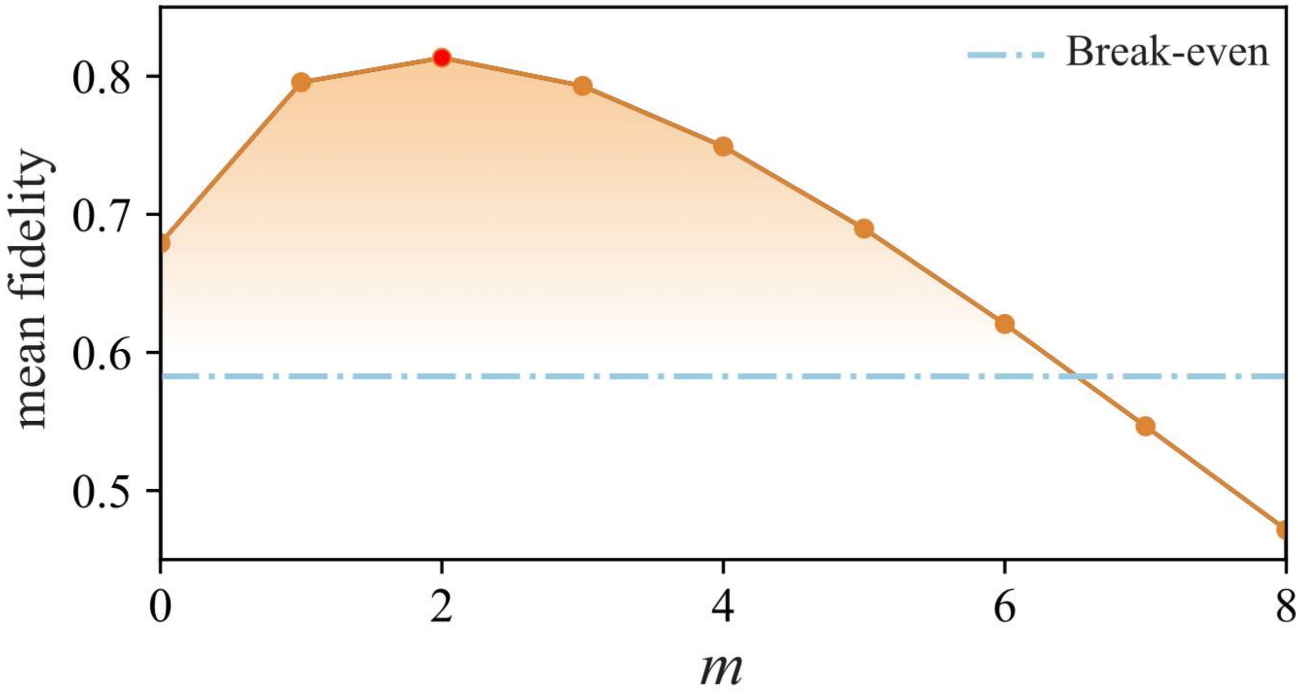

- on the one hand: the larger m the better is the KL condition satisfied
- on the other hand: the larger m the stronger is the single-photon loss

Experimental scheme
(proposal)

Experimental proposal

How to realize the effective Hamiltonian $H_{\text{eff}} = g(L_{\text{eng}}\sigma_+ + L_{\text{eng}}^\dagger \sigma_-)$

- encoding mode is complemented by a lossy auxiliary qubit and a lossy auxiliary mode

$$\frac{d\rho}{dt} = -i[H,\rho] + \frac{\gamma_{a1}}{2}\mathcal{D}[a] + \frac{\gamma_{b1}}{2}\mathcal{D}[\sigma_-] + \frac{\gamma_{c1}}{2}\mathcal{D}[c]$$

$$H = \omega_a a^\dagger a + \frac{\omega_b}{2}\sigma_z + \omega_c c^\dagger c + f(t)(a + a^\dagger)\sigma_x + g_c(t)(c^\dagger + c)\sigma_x + \frac{\chi}{2}a^\dagger a \sigma_z,$$

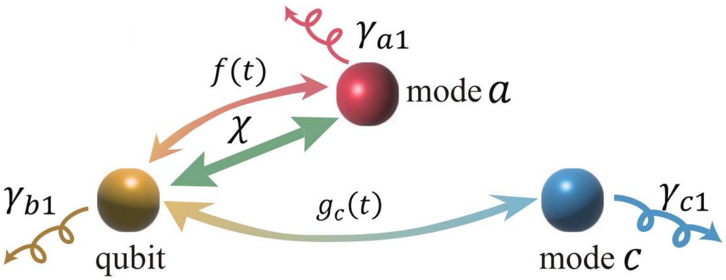

$\omega_b = \omega_c$

time-dependent control fields:

$$f(t) = \frac{2\alpha_0}{\sqrt{2}}\cos\left[(\omega_s + \frac{3\chi}{2})t\right] + \frac{2\alpha_0}{\sqrt{4}}\cos\left[(\omega_s + \frac{7\chi}{2})t\right]$$

$$g_c(t) = 2\alpha_1\cos(2\chi t) + 2\alpha_1\cos(4\chi t)$$

$\omega_s = \omega_a + \omega_b$

$\omega_a, \omega_b, \omega_c \gg \chi$ ..nonlinear coefficient

$\gamma_{c1} \gg \alpha_1 \geq \alpha_0 \gg \gamma_{a1}, \gamma_{b1}$

adiabatic elimination of high-decay mode c
results in the target Hamiltonian $H_{\text{eff}} = g(L_{\text{eng}}\sigma_+ + L_{\text{eng}}^\dagger \sigma_-)$

Experimental proposal

How to realize the effective Hamiltonian $H_{\text{eff}} = g(L_{\text{eng}}\sigma_+ + L_{\text{eng}}^\dagger \sigma_-)$

- encoding mode is complemented by a lossy auxiliary qubit and a lossy auxiliary mode

$$\frac{d\rho}{dt} = -i[H,\rho] + \frac{\gamma_{a1}}{2}\mathcal{D}[a] + \frac{\gamma_{b1}}{2}\mathcal{D}[\sigma_-] + \frac{\gamma_{c1}}{2}\mathcal{D}[c]$$

$$H = \omega_a a^\dagger a + \frac{\omega_b}{2}\sigma_z + \omega_c c^\dagger c + f(t)(a + a^\dagger)\sigma_x + g_c(t)(c^\dagger + c)\sigma_x + \frac{\chi}{2}a^\dagger a \sigma_z,$$

gain: $G \approx 5.8$

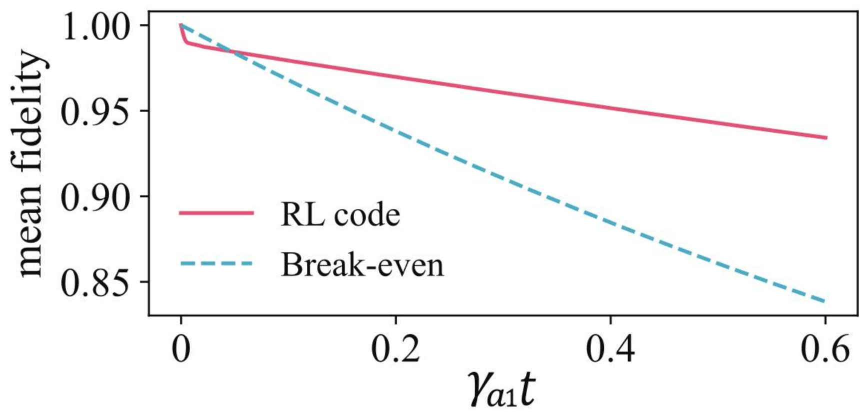

time-dependent control fields:

$$f(t) = \frac{2\alpha_0}{\sqrt{2}}\cos\left[(\omega_s + \frac{3\chi}{2})t\right] + \frac{2\alpha_0}{\sqrt{4}}\cos\left[(\omega_s + \frac{7\chi}{2})t\right]$$

$$g_c(t) = 2\alpha_1\cos(2\chi t) + 2\alpha_1\cos(4\chi t)$$

$$\omega_a/2\pi = 3.5\text{GHz}$$

$$\omega_b/2\pi = \omega_c/2\pi = 5\text{GHz}$$

$$\chi/2\pi = 3\text{MHz}$$

$$\alpha_0/2\pi = 0.05\text{MHz} \quad \alpha_1/2\pi = 0.07\text{MHz}$$

$$\gamma_{a1}/2\pi = 0.2\text{kHz} \quad \gamma_{b1}/2\pi = 2\text{kHz} \quad \gamma_{c1}/2\pi = 0.12\text{MHz}$$

Summary

Autonomously corrected RL code delivers..

- ..significantly improved bosonic qubits at the cost of moderate device overhead
- ..potential low-level element of an error correction stack
- ..neat application of how classical ML can support (inspire) the development of improved quantum technologies

Y. Zeng, Z.Y. Zhou, E. Rinaldi, CG, F. Nori, *PRL* 131, 050601 (2023)

Team

Yexiong Zeng

Zheng-Yang Zhou

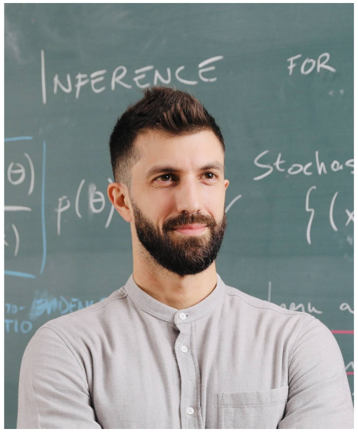

Enrico Rinaldi

Franco Nori

Team

Yexiong Zeng

Zheng-Yang Zhou

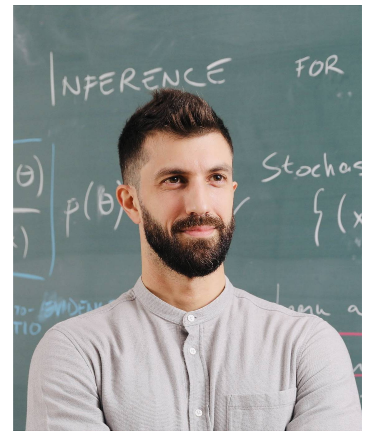

Enrico Rinaldi

Franco Nori

Postdoc positions available!

fnori@riken.jp
clemens.gneiting@riken.jp

Team

Yexiong Zeng

Zheng-Yang Zhou

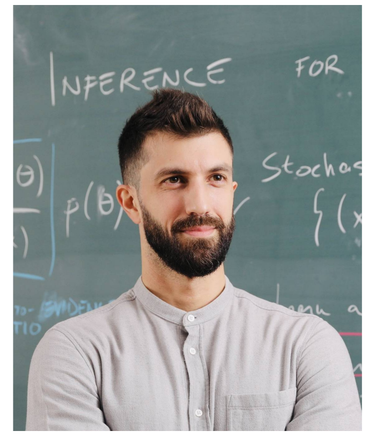

Enrico Rinaldi

Franco Nori

Postdoc positions available!

fnori@riken.jp
clemens.gneiting@riken.jp

Thank you for your attention!