

Preliminary Results: Temperature-Dependent Noise Modelling for IBM Backends

Jinzhao Sun,¹ Hamza Pereira,¹ and Jan Mol¹

¹*School of Physical and Chemical Sciences, Queen Mary University of London, London E1 4NS, United Kingdom*

(Dated: March 14, 2026)


This document presents preliminary simulation results tracking quantum algorithm performance as cryostat temperature is varied. We describe the modular simulation architecture, the assumptions underlying the noise models, the theoretical basis for temperature-dependent parameter extrapolation, and the case study results obtained across multiple IBM backends and benchmarks.

CONTENTS

I. Design Architecture	1
II. Preface of Result Validity	2
III. Case Studies	3
A. QV over qubit count	3
B. QV on different modern IBM processors	4
C. QV on different Falcon r5.11 backends	5
D. QV on older backends	7
E. VQE	7
F. Stim test	8
G. Mirror RB over temperature and depth	9
IV. Assumptions	10
A. Hardware parameters	10
1. Hardcoded parameters	10
2. Per-qubit parameters	11
3. Per-gate parameters	12
V. Theory behind Assumptions	12
A. Environment temperature estimation	12
B. Zero-temperature relaxation rate	12
C. Temperature-dependent predictions	12
References	13

I. DESIGN ARCHITECTURE

The simulation framework is fully modular: all parameters can be altered and all models can be swapped in and out, providing complete flexibility over how the performance value is computed.



?¹: Choose between an IBM Fake Backend or a basic noise model class

?²: A Pauli noise model is generated if the chosen metric is a Clifford metric. If a full noise model is generated, Pauli twirling can optionally be turned off.

FIG. 1. Dataflow for simulated performance on a quantum algorithm, from control parameters to noise model to performance. All nodes noted with a * are modular and customisable.

We computed simulations of both:

- Full quantum circuits, tracking full density matrix representations, allowing us to employ full Qiskit noise models.
- Circuits using only Clifford gates, tracking qubit state using stabiliser states and only employing simpler Pauli noise models. This allowed us to simulate much larger numbers of qubits.

II. PREFACE OF RESULT VALIDITY

The validity of these results hinges primarily upon the correctness of the code used to generate them. We have generated an extensive suite of unit tests to ensure the framework is operating as expected.

Moreover, the validity and applicability of these results further hinges upon the use of correct and properly catalogued assumptions. We employ the following to mitigate the risks of compromising these qualities:

- The use of a variety of different methods of producing results, including a variety of different metrics and methods of generating noise models. These methods vary in complexity and specificity, avoiding how many assumptions are universally relied upon.
- Comprehensively cataloguing the assumptions used at every stage of generating the results, including how they change for each method of producing results.

We treat each set of results as a case study. Because there is significant overlap between the assumptions of each case study, we include the full set of assumptions in Section IV, and specify which ones apply for each.

There are a few assumptions that, due to the inherent limitations of the simulation architecture, universally apply:

- All noise models are inherently gate-local (and thus also qubit-local).
- All noise models are state independent and unchanging over time throughout the whole metric simulation, and every shot and trial of the metric.
- All noise models are Markovian.

III. CASE STUDIES

In any case study, to get more information on the processor type, development date and retirement date of each backend used, refer to Section IV A 1. In all case studies we analyse one extra independent variable as we vary cryostat temperature.

A. QV over qubit count

We measured how the heavy probability changed as cryostat temperature was increased, running full quantum circuit simulations on progressively increased numbers of qubits.




FIG. 2. Sherbrooke QV over cryostat temperature on different qubit counts.




FIG. 3. Sherbrooke T_1 and T_2 times over cryostat temperature.

B. QV on different modern IBM processors

We measured how heavy probability changed as cryostat temperature was increased for three modern IBM backends: Sherbrooke, Geneva and Perth.




FIG. 4. Different modern IBM backends: QV over cryostat temperature.




FIG. 5. Geneva T_1 and T_2 times over cryostat temperature.




FIG. 6. Perth T_1 and T_2 times over cryostat temperature.

C. QV on different Falcon r5.11 backends

We measured how heavy probability changed as cryostat temperature was increased for three Falcon r5.11 backends: Cairo, Hanoi and Auckland.







FIG. 7. Different Falcon r5.11 processors: QV over cryostat temperature.

FIG. 8. Cairo T_1 and T_2 times over cryostat temperature.FIG. 9. Hanoi T_1 and T_2 times over cryostat temperature.FIG. 10. Auckland T_1 and T_2 times over cryostat temperature.

D. QV on older backends

We measured how heavy probability changed as cryostat temperature was increased for four older IBM backends, each on 11 qubits: Tokyo, Kolkata, Melbourne and Montreal.




FIG. 11. Heavy probability over cryostat temperature for older backends.

E. VQE

We measured how well a noisy circuit can reconstruct the energy estimation of a 1D Fermi–Hubbard model, by taking the difference of its estimation from the calculated ground truth, as cryostat temperature was increased. Simulations were run on the modern Sherbrooke backend.




FIG. 12. VQE energy estimation error over cryostat temperature (Sherbrooke).

F. Stim test

We measured GHZ state preparation fidelity as cryostat temperature was increased, running Stim stabiliser simulations on the Sherbrooke backend.




FIG. 13. GHZ state fidelity over cryostat temperature (Sherbrooke, Stim).

G. Mirror RB over temperature and depth

We measured how initial state preservation changed over standard randomised mirror benchmarks as the number of circuit layers was increased, for increasing cryostat temperatures.




FIG. 14. Initial state preservation over depth for a range of cryostat temperatures.

As expected from [1], the state preservation can be modelled by an exponential decay function for each set cryostat temperature. The decay parameter (depolarising parameter) as a function of cryostat temperature is shown below:




FIG. 15. Depolarising parameter over cryostat temperature (Mirror RB).

IV. ASSUMPTIONS

A. Hardware parameters

This section outlines the modelling assumptions when creating noise models from IBM fake backends.

We implicitly use all the assumptions of the two-bath model outlined in [2], Section A. In summary, we assume that the qubits only thermalise with two distinct temperature environments: the mixing chamber (MXC) and the residual electromagnetic environment.

All parameters in this model can be changed. In this paper, we choose to only vary temperature as the independent parameter. The rest of the parameters are either hardcoded beforehand or fit to calibration data at runtime.

To ensure the hardcoded parameters are still representative of the specific fake backend that we are modelling, for each parameter we choose one of three values, decided by the age of the backend. We characterise each backend as either modern, intermediate or legacy.

1. Hardcoded parameters

We hardcode parameters that either can or need to be assumed, because of known manufacturing properties or insufficient information or degrees of freedom to determine them.

For some of these hardcoded parameters, we assume different values depending on the processor type used in the backend. The three groups we use are summarised below.

Modern: (Eagle and later Falcon revisions)

- *Eagle r3/r1* (127 qubits): **sherbrooke** (2022–2025), **washington** (2022–2023), **prague** (2022–2023)
- *Falcon r5.11* (27 qubits): **cairo** (2021–2024), **hanoi** (2021–2024), **auckland** (2021–2023)
- *Falcon r8* (27 qubits): **geneva** (2021–2023)
- *Falcon r5.11H* (7 qubits): **oslo** (2022–2023), **perth** (2022–2023)

Intermediate: (mid-generation Falcon sub-revisions and Hummingbird)

- *Falcon r4T* (5 qubits): **athens**, **belem**, **lima**, **quito**, **santiago**, **bogota**, **rome**
- *Falcon r5.11H* (5–7 qubits): **manila**, **casablanca**, **jakarta**, **lagos**, **nairobi**
- *Falcon r4P* (16 qubits): **guadalupe**
- *Falcon r4* (27 qubits): **toronto**, **montreal**, **sydney**, **paris**
- *Falcon r5.10/r5.11* (27 qubits): **mumbai**, **kolkata**
- *Hummingbird r2* (65 qubits): **manhattan**, **brooklyn**

Legacy: (earliest IBM Quantum backends)

- *Canary r1.3* (1 qubit): **armonk**
- *Canary r1* (5 qubits): **yorktown**, **tenerife**, **burlington**, **essex**, **vigo**, **ourense**, **valencia**, **london**
- *Canary r1.1* (15–16 qubits): **melbourne**, **rueschlikon**
- *Penguin* (20 qubits): **almaden**, **boeblingen**, **singapore**, **johannesburg**, **poughkeepsie**, **tokyo**
- *Falcon r1* (28 qubits): **cambridge**
- *Hummingbird r1* (53 qubits): **rochester**

The following parameters are hardcoded:

1. $\gamma_{\text{MXC}}/\gamma_1^0$ (**coupling ratio**). The ratio of how much of the qubit’s relaxation couples to the mixing chamber versus the residual electromagnetic environment. Lvov et al. [2], Table III, report experimental values of 0.60–0.85 for custom devices. We assume IBM has superior equipment, and set: modern backends = 0.9, intermediate = 0.85, legacy = 0.8.
2. **Default operational cryostat temperature:** 13 mK. All calibration data is assumed to be collected at this temperature.

3. $\Delta = 2.88 \times 10^{-23} \text{ J (180 } \mu\text{eV)}$, **superconducting energy gap**. This is the bulk aluminium value; IBM transmon electrodes use aluminium thin-films that can deviate slightly, but $180 \mu\text{eV}$ is standard in the superconducting qubit literature [3].
4. ζ (**subgap transparency**). Characterises the transparency of quasiparticle tunnelling channels in the subgap regime; literature values range from 10^{-3} to 10^{-5} [4]. Treated as a uniform constant across the chip. Modern $= 5 \times 10^{-5}$, intermediate $= 1 \times 10^{-4}$, legacy $= 5 \times 10^{-4}$.

2. Per-qubit parameters

The following parameters are derived on a per-qubit basis at runtime. All per-qubit data is retrieved directly from official IBM calibration data.

5. E_c (**charging energy**). In the transmon regime ($E_J/E_c \gg 1$), the anharmonicity approaches $\alpha \rightarrow -E_c/\hbar$ [5], Eq. (2.5). Since IBM reports anharmonicity $\alpha/2\pi$ in Hz: $E_c = \hbar|\alpha|$.
6. E_J (**Josephson energy**). Obtained by inverting the transmon transition frequency $\hbar\omega_{ge} \approx \sqrt{8E_J E_c} - E_c$ [5]:

$$E_J = \frac{(\hbar\omega_{ge} + E_c)^2}{8 E_c}.$$

7. ω_p (**plasma frequency**). Sets the energy scale for quasiparticle-induced transition rates [3]:

$$\omega_p = \frac{\sqrt{8 E_J E_c}}{\hbar}.$$

8. R_n (**normal-state junction resistance**). From the Ambegaokar–Baratoff relation [6]:

$$R_n = \frac{\pi \hbar \Delta}{4 e^2 E_J}.$$

9. N_e (**effective number of junction channels**). Following [4], Eq. (15):

$$N_e = \frac{1}{\zeta} \cdot \frac{1}{2 R_n g_K},$$

where $g_K = e^2/h$ is the conductance quantum.

10. p_e (**thermal excited-state population**). The probability of measuring $|1\rangle$ after preparing $|0\rangle$ is extracted for each qubit. The residual excited state population p_e is estimated as equal to this value. This is likely a significant overestimation of p_e , as much of the readout error will not be due to excited state population, but simple measurement error. This in turn will safely overestimate T_{env} .
11. T_{env} (**effective environment temperature**). Estimated per qubit from p_e values and previously derived parameters. The median across all qubits is then taken. Derived from the two-bath model of [2], Eq. (14), in the low- T_{MXC} limit (see Section V A).
12. γ_1^0 (**zero-temperature relaxation rate**). At $T_{\text{op}} \approx 13 \text{ mK}$, quasiparticle contributions are negligible and the MXC bath occupation vanishes. This permits a closed-form inversion of the two-bath relaxation model [2] (see Section V B).
13. $\gamma_{\varphi, \text{base}}$ (**base pure dephasing rate**). At T_{op} , quasiparticle dephasing is negligible, so [7]:

$$\gamma_{\varphi, \text{base}} = \frac{1}{T_{2, \text{meas}}} - \frac{1}{2 T_{1, \text{meas}}}.$$

Clamped to $\gamma_{\varphi, \text{base}} \geq 0$. Note: IBM reports T_2 from Hahn echo measurements, so $\gamma_{\varphi, \text{base}}$ captures only high-frequency dephasing sources.

3. Per-gate parameters

14. ϵ_{off} (**gate error offset**). Gate infidelity at the default operational temperature is estimated using Eq. (12), taking into account per-qubit parameters. The linear difference between this estimated infidelity and the actual calibrated infidelity is stored and added as an offset. The physical interpretation is that this represents the temperature-independent component of gate error. Gates without a calibrated gate error (e.g., reset gates) are assigned decoherence-only error.

V. THEORY BEHIND ASSUMPTIONS

A. Environment temperature estimation

At the default operational cryostat temperature, the Bose–Einstein occupation of the MXC bath at qubit frequencies is vanishingly small: $n(\omega_{ge}, T_{\text{op}}) \approx 0$. The two-bath effective occupation [2], Eq. (14), therefore simplifies to

$$n_{\text{eff}} \approx \frac{\gamma_{\text{env}}}{\gamma_1^0} n(\omega_{ge}, T_{\text{env}}). \quad (1)$$

Note that γ_1^0 will cancel out, so is not needed directly.

The effective thermal photon occupation is calculated from:

$$n_{\text{eff}} = \frac{p_e}{1 - 2p_e}. \quad (2)$$

Inverting the Bose–Einstein distribution for T_{env} :

$$T_{\text{env}} = \frac{\hbar\omega_{ge}}{k_B \ln\left(1 + \frac{\gamma_{\text{env}}/\gamma_1^0}{n_{\text{eff}}}\right)}. \quad (3)$$

Combining the two equations:

$$T_{\text{env}} = \frac{\hbar\omega_{ge}}{k_B \ln\left(1 + \frac{\gamma_{\text{env}}/\gamma_1^0}{p_e} (1 - 2p_e)\right)}. \quad (4)$$

Therefore, T_{env} depends only on the global coupling ratio, the per-qubit frequency, and p_e .

We compute T_{env} per qubit and take the median. This uses the assumption that the electromagnetic environment temperature is spatially uniform across the chip.

B. Zero-temperature relaxation rate

At $T_{\text{op}} \approx 13 \text{ mK}$, quasiparticle-induced relaxation is exponentially suppressed ($\gamma_{\text{QP}} \propto e^{-\Delta/k_B T} \approx 0$). The total relaxation rate [2], Eqs. (15) and (27), reduces to

$$\frac{1}{T_{1,\text{meas}}} = \gamma_1^0 (2n_{\text{eff}} + 1), \quad (5)$$

where n_{eff} is computed from T_{env} (Section V A). Solving:

$$\gamma_1^0 = \frac{1}{T_{1,\text{meas}} (2n_{\text{eff}} + 1)}. \quad (6)$$

C. Temperature-dependent predictions

With all per-qubit parameters, we compute the following temperature-dependent quantities at arbitrary cryostat temperature T .

a. *Quasiparticle density.*

$$x_{\text{QP}}(T) = \sqrt{\frac{2\pi k_B T}{\Delta}} \exp\left(-\frac{\Delta}{k_B T}\right). \quad (7)$$

b. *Quasiparticle relaxation rate.* From [2], Eq. (26), originally derived in [3]:

$$\gamma_{\text{QP}}(T) = \frac{\omega_p^2}{\pi \omega_{ge}} \left[x_{\text{QP}} \sqrt{\frac{2\Delta}{\hbar \omega_{ge}}} + 4 e^{-\Delta/k_B T} \cosh\left(\frac{\hbar \omega_{ge}}{2k_B T}\right) K_0\left(\frac{\hbar \omega_{ge}}{2k_B T}\right) \right], \quad (8)$$

where K_0 is the modified Bessel function of the second kind.

c. *Quasiparticle dephasing rate.* From [2], Eq. (29), following [4]:

$$\gamma_{\varphi, \text{QP}}(T) \sim 4\pi \frac{\omega_p^2}{\omega_{ge}} \sqrt{\frac{x_{\text{QP}}^A}{N_e}}, \quad (9)$$

where $x_{\text{QP}}^A = \exp(-\Delta/k_B T)$. This expression is order-of-magnitude; the numerical prefactor may differ by a factor of a few.

d. *Relaxation time.*

$$T_1(T) = \frac{1}{\gamma_{\text{QP}}(T) + \gamma_1^0 (2n_{\text{eff}}(T) + 1)}, \quad (10)$$

where $n_{\text{eff}}(T) = (\gamma_{\text{MXC}}/\gamma_1^0) n(\omega_{ge}, T) + (1 - \gamma_{\text{MXC}}/\gamma_1^0) n(\omega_{ge}, T_{\text{env}})$ follows the two-bath model [2], Eq. (14).

e. *Coherence time.*

$$T_2(T) = \frac{1}{\frac{1}{2T_1(T)} + \gamma_{\varphi, \text{base}} + \gamma_{\varphi, \text{QP}}(T)}. \quad (11)$$

f. *Gate fidelity.* For an N -qubit gate of duration t_{gate} , the decoherence-limited fidelity is [8], Eq. (6), derived from [9]:

$$F_N(T) = 1 - \frac{d t_{\text{gate}}}{2(d+1)} \sum_{i=1}^N \left(\frac{1}{T_{1,i}(T)} + \frac{1}{T_{\varphi,i}(T)} \right), \quad (12)$$

where $d = 2^N$ is the Hilbert space dimension, $T_{\varphi,i} = 1/(\gamma_{\varphi, \text{base},i} + \gamma_{\varphi, \text{QP},i}(T))$, and the sum runs over the constituent qubits. This formula assumes a local noise model acting independently on each qubit and captures only decoherence-limited infidelity.

Temperature-independent coherent error sources are accounted for via ϵ_{off} :

$$\epsilon_{\text{coh}} = \epsilon_{\text{cal}} - \epsilon_{\text{dec}}(T_{\text{op}}), \quad (13)$$

where ϵ_{cal} is the calibrated gate error from the backend properties. The total predicted gate error at temperature T is then

$$\epsilon_{\text{gate}}(T) = \min(1, \epsilon_{\text{dec}}(T) + \epsilon_{\text{coh}}). \quad (14)$$

This additive decomposition is a linear approximation, accurate in the small-error regime ($\epsilon \sim 10^{-3}$ – 10^{-2}) where product terms are negligible. By construction, the predicted gate error recovers the calibrated value at T_{op} .

-
- [1] Easwar Magesan, J. M. Gambetta, and Joseph Emerson. Scalable and robust randomized benchmarking of quantum processes. *Physical Review Letters*, 106(18):180504, 2011.
- [2] D. S. Lvov, S. A. Lemziakov, E. Ankerhold, J. T. Peltonen, and J. P. Pekola. Thermometry based on a superconducting qubit. *Physical Review Applied*, 23:054079, May 2025.
- [3] G. Catelani, Jens Koch, L. Frunzio, R. J. Schoelkopf, M. H. Devoret, and L. I. Glazman. Quasiparticle relaxation of superconducting qubits in the presence of flux trapping. *Physical Review Letters*, 106:077002, February 2011.

- [4] G. Catelani, R. J. Schoelkopf, M. H. Devoret, and L. I. Glazman. Relaxation and frequency shifts induced by quasiparticles in superconducting qubits. *Physical Review B*, 84:064517, 2011.
- [5] Jens Koch, Terri M. Yu, Jay Gambetta, A. A. Houck, D. I. Schuster, J. Majer, Alexandre Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf. Charge-insensitive qubit design derived from the Cooper pair box. *Physical Review A*, 76:042319, October 2007.
- [6] Vinay Ambegaokar and Alexis Baratoff. Tunneling between superconductors. *Physical Review Letters*, 10(11):486–489, June 1963.
- [7] Philip Krantz, Morten Kjaergaard, Fei Yan, Terry P. Orlando, Simon Gustavsson, and William D. Oliver. A quantum engineer’s guide to superconducting qubits. *Applied Physics Reviews*, 6:021318, 2019.
- [8] Slawomir Simbierowicz, Massimo Borrelli, Volodymyr Monarkha, Ville Nuutinen, and Russell E. Lake. Inherent thermal-noise problem in addressing qubits. *Physical Review X Quantum*, 5:030302, July 2024.
- [9] Tomás Abad, James Fernandez, and Iman Marvian. Noise in quantum computing. *Physical Review Letters*, 129:150504, 2022.