



Gate-Error Modelling of Microwave-over-Fiber Drive Links for Superconducting Qubits Under Cryogenic Heat Constraints

Suhail Najm Abdullah *

Department of Electrical Engineering, College of Engineering, Al-Iraqia University, Baghdad, Iraq

*Corresponding Author: suhail_najm@aliraqia.edu.iq

Citation: Suhail Najm Abdullah. (2026). Gate-Error Modelling of Microwave-over-Fiber Drive Links for Superconducting Qubits Under Cryogenic Heat Constraints. Iraqi Journal of Communications and Computer Networks, 2(3), 11–24.

ARTICLE INFO

Received: Jul. 26, 2026

Revised: Sep. 14, 2026

Accepted: Oct. 02, 2026

Published: Oct. 11, 2026



ABSTRACT

Analog microwave-over-fiber (MoF) links can deliver qubit-control signals to a cryogenic photodiode with a much lower passive heat load than coaxial wiring. However, the required photocurrent dissipates heat at the cold stage, and its shot noise perturbs the qubit. This paper develops a physics-structured model that maps the Mach–Zehnder modulator (MZM) bias, the modulation depth, the photocurrent scale, and the gate duration of an analog MoF link to the link-induced average infidelity of a transmon $X_{\pi/2}$ gate. The model is then used to design links that minimize idle optical heat. The optical link was simulated in OptiSystem software, and the received photocurrents drove a three-level transmon model with Lindblad decoherence in MATLAB. Across 1344 simulation runs covering 15–45 dB of optical attenuation, bias angles of 90° – 170° , modulation indices of 0.2–1.3, and gate durations of 20 and 40 ns, deterministic link distortion contributed an infidelity of at most 2.7×10^{-8} , the link-induced error scaled as $I_0^{-1.004 \pm 0.018}$, and leakage accounted for about 60% of it. The closed-form model, which includes pulse-induced shot noise and Bessel-function drive compression and has a single calibrated coefficient, predicted held-out low-power and low-bias configurations with a logarithmic root-mean-square error of 0.040 decades, compared with 0.155–0.412 decades for three reference models. Six operating points designed with the frozen model met their pre-registered error targets to within 8.5% in new simulations and required 9.1–11.1 times less idle optical power at the photodiode than the best quadrature-biased design.

Keywords: Microwave photonics, Analog optical link, Superconducting qubit, Shot noise, Cryogenic heat load

INTRODUCTION

Superconducting quantum processors are built from qubits, the quantum counterpart of the bit, and each qubit is operated by short microwave pulses known as gates. Every gate has an error, which is the probability that the pulse fails to produce the intended change of the qubit state, and this error must be kept very small. The pulses are generated at room temperature and carried by coaxial cables into a dilution refrigerator, where the processor operates at a few tens of millikelvin. Each cable also conducts heat from the warm environment to the cold stages, and the limited cooling power of the refrigerator restricts a processor to a few thousand qubits [1].

An optical fiber offers a way around this limit, because it conducts very little heat at low temperature and has a very large bandwidth. Lecocq et al. [1] replaced the coaxial drive line with a fiber that guided modulated laser light from room temperature to a commercial InGaAs photodiode operated at 20 mK. The photodiode converted the light back into a microwave current, which was used to control and read out a transmon, a widely used type of superconducting qubit. They estimated a passive heat load of about 3 pW for the fiber, compared with a measured 14 nW for a coaxial cable [1]. Such a link is an analog microwave-over-fiber (MoF) system of the kind studied in microwave photonics, the field that generates and processes microwave signals in the optical domain and supplies the modulators and detectors that these links require [2].

A photonic link does not remove the heat load; it moves it. The optical power absorbed by the cryogenic photodiode is dissipated at the cold stage, and the photocurrent carries shot noise, the random fluctuation caused by the discrete arrival of photons, which can excite the qubit and cause gate errors [1], [3]. More light therefore gives a stronger and relatively cleaner drive signal but adds heat, whereas less light saves heat but increases the relative noise. Recent work has approached this trade-off from several directions. Nakamura et al. [3] showed that short-pulse illumination improves the shot-noise-limited signal-to-noise ratio, and that generating several microwave carriers with short pulses needs less optical power than generating one carrier with modulated continuous-wave light. Coherent optical control of a transmon through a microwave–optical transducer [4] and all-optical qubit readout [5] have also been demonstrated.

For an analog link driven by continuous-wave light, however, the designer must still choose the modulator bias, the modulation depth, and the optical power. Each choice affects both the gate error and the heat delivered to the cold stage, yet conventional link metrics describe the classical signal and do not predict the error of a quantum gate, and, to the best of the author's knowledge, no compact model connects these settings to gate error and cryogenic heat. This paper therefore treats the choice of link settings as an inverse design problem: given a gate-error target and a photocurrent budget, find the modulator bias that minimizes the idle optical heat, meaning the optical power absorbed by the photodiode while no pulse is applied. The contributions are as follows. First, an OptiSystem and MATLAB co-simulation of an analog MoF drive link is built, in which OptiSystem computes the optical link, MATLAB computes the qubit dynamics. Second, a closed-form gate-error model is derived that accounts for the dependence of shot noise on the bias, the noise generated by the pulse itself, and the compression of the drive described by Bessel functions, and that contains a single calibrated coefficient. Third, on configurations withheld from fitting, the model predicts the link-induced error with a logarithmic root-mean-square error of 0.040 decades, compared with 0.155 to 0.412 decades for three reference models. Fourth, the model is turned into a design rule for the bias, and six operating points designed with the frozen model were then simulated against error targets and predictions recorded beforehand. They met the targets within 8.5% and required 9.1 to 11.1 times less idle optical power at the photodiode than the best quadrature-biased design.

The remainder of the paper is organized as follows. The literature review summarizes the relevant models and states the gap. The methodology section describes the simulated system, the physical model, the dataset construction, and the statistical evaluation. The results and discussion section reports the link physics, the model comparison, the design application and its verification, and the limitations of the study.

LITERATURE REVIEW

Photonic Drive of Superconducting Qubits

Lecocq et al. [1] reported photodiode-generated drive and readout of a transmon qubit with qubit frequencies of 5.052 GHz (readout) and 5.088 GHz (control), an anharmonicity of 210 MHz, and relaxation and coherence times of $T_1 \approx 43\mu\text{s}$ and $T_2 = 37\mu\text{s}$, respectively. Rabi drive was achieved with an average photocurrent of 4 μA , the measured photocurrent noise was consistent with shot noise, and the authors inferred that shot noise has a negligible impact on gate fidelity at this photocurrent. They further showed that, once the passive heat load is negligible, the active heat load during gates limits the number of qubits that can be supported, and that this number is inversely proportional to the duty cycle at which each qubit is driven [1]. Nakamura et al. [3] demonstrated high responsivity of an extended-InGaAs photodiode at 1550 nm and 20 mK and showed that short-pulse illumination improves the shot-noise-limited signal-to-noise ratio relative to modulated continuous-wave light, by 20 dB in the phase quadrature and 3 dB in the amplitude quadrature at 40 μA and 4 K. They also demonstrated simultaneous microwave pulses at two carrier frequencies, which indicates a path toward multiplexed drive. Neither study, however, optimizes the modulator bias against gate error and cryogenic heat.

Control Electronics and Gate Fidelity

Errors in classical control electronics have also been related analytically to qubit fidelity. Van Dijk et al. [6] presented an analytical study of how static, dynamic, systematic, and random errors in the control electronics affect the fidelity of single-qubit rotations, two-qubit gates, and readout, focusing on single-electron spin qubits and deriving electrical specifications for a target fidelity. Ball et al. [7] connected the phase-noise metrics of the master oscillator to qubit dephasing spectral densities and used them to calculate fidelity bounds for trapped-ion and superconducting qubits. These results translate electronic specifications into fidelity, but they do not include the modulator transfer function or a photocurrent-dependent noise source.

Analog Photonic Links

On the photonic side, analog links have long been characterized by gain, noise figure, and dynamic range, and

the limits that these metrics impose on device design are well documented [8]. Such figures of merit describe the classical signal, however, and do not predict the error of a quantum gate driven by that signal.

Research Gap

Existing studies relate gate error to photocurrent at a fixed operating point [1], improve the signal-to-noise ratio of the link through the detection scheme [3], or relate generic electronic impairments to fidelity without modeling the optical link [6], [7]. A targeted search of Nature-family, APS, Optica, IEEE, and arXiv sources on photonic qubit control, cryogenic photodiodes, multiplexed optical drive, and control-electronics error models, conducted in September 2026, did not identify a model that jointly relates the MZM bias, the modulation depth, and the photocurrent to transmon gate error, leakage, and cryogenic heat in a form suitable for inverse design. This gap motivates the physics-structured gate-error model and the heat-load-based inverse design developed in this research.

METHODOLOGY

The method has four parts. First, the optical link is simulated in OptiSystem and the qubit response is computed in MATLAB. Second, the drive waveform, the photocurrent, the shot noise, and the transmon are modeled. Third, a closed-form model of the link-induced gate error is derived. Fourth, the model is fitted, tested on withheld data, and used for design. For a reader who is not familiar with quantum computing, the qubit can be regarded as a device whose state is changed by the drive pulse, and the gate error as the probability that this change is not the intended one.

Simulated System and Its Components

Fig. 1 shows the OptiSystem layout used for this research. Signals travel from left to right: a MATLAB component generates the drive voltage, the modulator imposes it on a laser carrier, the optical signal passes through fiber and an attenuator, two photodiodes convert it back to current, and a second MATLAB component evaluates the qubit. Green connections carry optical signals and blue connections carry electrical signals. The role of each component is included as follows.




Fig. 1. OptiSystem layout of the analog microwave-over-fiber drive link. Green lines are optical connections and blue lines are electrical connections.

A sine generator operating at 5.052 GHz with zero phase provides the time grid. OptiSystem exchanges signals between components as sampled data on a common time grid; the generator supplies the MATLAB transmitter with an electrical signal carrying this grid, and the transmitter retains the grid while replacing the samples. The generator frequency equals the qubit frequency, so the probe waveform used during commissioning was physically meaningful.

The MATLAB transmitter has one electrical input and one output. For each sweep iteration, it reads the next pending row of the design table and generates the drive voltage $V(t) = V_b + V_{pk} s(t)$, as given in (1), where $s(t)$ is the normalized DRAG pulse defined in (2) and (3). Generating the waveform in MATLAB allows the bias, amplitude, and pulse duration to change between iterations and allows the run context to be recorded for the receiver.

An electrical 1-to-2 fork sends one copy of $V(t)$ to the modulator and a second copy to input 3 of the MATLAB receiver. The receiver uses this reference to verify that the transmitter and the receiver refer to the same run and

to fit the actual modulator operating point from the received current.

A continuous-wave laser with a frequency of 201.204 THz and a power of 0 dBm serves as the optical source. The carrier frequency corresponds to the 1490 nm wavelength used in [1]. The laser power is fixed, and the received power is set by the attenuator.

A lithium niobate Mach–Zehnder modulator converts the drive voltage into optical intensity according to the cosine transfer law in (4). The switching voltage was set to 4 V and the extinction ratio to 30 dB. The bias is applied through the DC component V_b of the drive, so the bias angle can be varied for each run. Because the modulator has two electrical inputs and OptiSystem computes a component only when every input carries a signal, the second input is connected to a zero-amplitude source.

The zero-amplitude source is a sine generator with zero amplitude and zero bias. It supplies 0 V to the second modulator arm, so the modulator is driven on a single arm. This leaves the intensity transfer law unchanged for a photodiode receiver.

A single-mode fiber of 0.02 km (20 m) represents the run from the room-temperature modulator to the cryostat. Over this length, dispersion and loss are negligible, as confirmed by the fitted attenuation offset of 0.08 dB.

An optical attenuator sets the received optical power and, therefore, the photocurrent scale I_0 . Stepping the attenuation from 15 to 45 dB varies I_0 from 12.4 μ A to 12.4 nA, which covers the photocurrents of 20 nA for readout and 4 μ A for drive reported in [1].

An optical 1-to-2 fork duplicates the optical signal without splitting loss, so that both photodiodes receive identical inputs and differ only in their noise settings.

Two PIN photodiodes with a responsivity of 0.8 A/W, designated noisy and clean, convert the optical signals to current. The clean photodiode has all noise sources disabled and yields the deterministic current, which is used to calibrate the gate and to measure deterministic distortion. The noisy photodiode has shot noise enabled. Thermal noise and dark current are set to zero in both, because the photodiode in [1] operates at 20 mK into a cold load and was found to be shot-noise limited.

Each photodiode is followed by a fourth-order Bessel low-pass filter with a 20 GHz cutoff, which represents the 20 GHz nominal bandwidth of the photodiode in [1]. A Bessel response was chosen for its nearly linear phase, which preserves the pulse shape.

The MATLAB receiver has three electrical inputs: input 1 receives the noisy current, input 2 the noise-free current, and input 3 the drive reference. The receiver estimates the link parameters, forms the shot noise, computes the transmon response, and writes all results to disk.

An oscilloscope visualizer displays the receiver output. It was used only for inspection during commissioning and does not affect the data.

Fig. 2 summarizes how the work is shared between the two programs. OptiSystem computes the deterministic optical link; MATLAB generates the waveform, realizes the shot noise and computes the qubit dynamics and the metrics.




Fig. 2. Functional partition of the co-simulation. OptiSystem computes the optical link; MATLAB generates the drive, estimates the link parameters, realizes the shot noise and evaluates the transmon.

The global simulation settings were a reference bit rate of 1 Gb/s, a sequence length of 128 bits, and 64 samples per bit. These settings give a sampling rate $f_s = 64$ GS/s, $N = 8192$ samples, and a time window of 128 ns.

The Nyquist frequency of 32 GHz exceeds both the 20 GHz filter bandwidth and the second harmonic of the qubit frequency, and the window holds a 40 ns pulse with ample margin. Table I lists all parameters and their origin.

TABLE I. SIMULATION PARAMETERS AND THEIR ORIGIN

Quantity	Symbol	Value	Origin
Qubit frequency	$\omega_q/2\pi$	5.052 GHz	Reported in [1]
Anharmonicity	$\alpha/2\pi$	-210 MHz	Magnitude from [1]; negative sign of a transmon [9]
Relaxation and coherence times	T_1, T_2	43 μ s, 37 μ s	Reported in [1]
Laser wavelength and power	λ, P_L	1490 nm, 0 dBm	Wavelength from [1]; power chosen
Modulator switching voltage	V_π	4 V set, 3.998 V measured	Set; measured in the pilot
Extinction ratio	ER	30 dB	Chosen
Photodiode responsivity	$\mathcal{R}$	0.8 A/W	Chosen
Photodiode bandwidth	f_c	20 GHz, 4th-order Bessel	Nominal value from [1]
Sampling rate, window	f_s, T	64 GS/s, 128 ns	Calculated
Pulse shape	σ, λ_D	$\tau/4, 1$	DRAG [10]
Drive-line high-pass edge	f_{HP}	1 GHz	Chosen
Target rotation	θ	$\pi/2$	Chosen

Drive Waveform and Link Model

The MATLAB transmitter writes the modulator voltage as

$$V(t) = V_b + V_{pk} s(t) \quad (1)$$

where V_b is the bias voltage, V_{pk} the peak drive voltage, and $s(t)$ the normalized pulse. The pulse is a first-order DRAG Gaussian [10] centered at $t_c = 64$ ns. DRAG is a standard pulse shape that adds a second, quadrature component to reduce leakage of the qubit into higher energy levels. The pulse is written as a quadrature-modulated carrier, as in [10]:

$$s(t) = X(t)\cos(\omega_d t) + Y(t)\sin(\omega_d t) \quad (2)$$

$$X(t) = \frac{e^{-\frac{(t-t_c)^2}{2\sigma^2}} - e^{-\frac{\tau^2}{8\sigma^2}}}{1 - e^{-\frac{\tau^2}{8\sigma^2}}}, \quad Y(t) = -\frac{\lambda_D \dot{X}(t)}{\alpha} \quad (3)$$

for $|t - t_c| \leq \tau/2$ and zero elsewhere, where $\omega_d = \omega_q$ is the drive frequency, τ the gate duration, $\sigma = \tau/4$, $\lambda_D = 1$ the DRAG coefficient, and α the anharmonicity in rad/s. Both envelopes are scaled so that the envelope magnitude $a(t) = [X^2(t) + Y^2(t)]^{1/2}$ has a maximum of one. The in-phase component $X(t)$ performs the rotation, and the quadrature component $Y(t)$ suppresses leakage to the second excited state [10]. The envelope pair in (3) is the truncated Gaussian with the first-order derivative quadrature of DRAG [10], written with the normalization used here.

For an ideal photodiode the modulator law gives the photocurrent

$$i(t) = I_0 \{1 + \cos[\varphi_B + \beta s(t)]\}, \quad \varphi_B = \varphi_0 + kV_b, \quad \beta = kV_{pk}, \quad k = \frac{\pi}{V_\pi} \quad (4)$$

where $I_0 = \mathcal{R}P_{\max}/2$ is the current at half transmission, $P_{\max}$ the received optical power at maximum transmission, $\mathcal{R}$ the responsivity, φ_0 the intrinsic bias phase of the modulator, φ_B the bias angle, and β the modulation index. With this convention, $\varphi_B = \pi/2$ is quadrature (the point of most linear response) and $\varphi_B = \pi$ is the null (the point of minimum transmission). Writing $s = a\cos(\omega_d t + \chi)$ with $\chi(t) = \arctan[Y(t)/X(t)]$ and applying the Jacobi–Anger expansion gives the amplitude of the current at the drive frequency and the carrier-averaged current:

$$i_1(t) = 2I_0 \sin\varphi_B J_1(\beta a(t)), \quad \langle i(t) \rangle = I_0 [1 + \cos\varphi_B J_0(\beta a(t))] \quad (5)$$

where J_0 and J_1 are Bessel functions of the first kind. Equation (5) contains the two effects that the design

exploits. The useful drive i_1 grows with $\sin\varphi_B$, while the average current $\langle i \rangle$ falls toward the null. During a pulse at low bias, however, $\langle i \rangle$ rises because $J_0(\beta a) < 1$, so the pulse generates its own photocurrent.

The photocurrent carries shot noise with the one-sided power spectral density of the form given in [1], evaluated here at the instantaneous carrier-averaged current:

$$S_I(t) = 2q \langle i(t) \rangle \quad (6)$$

where q is the electron charge. The heat considered in this work is the optical power absorbed by the photodiode when no pulse is applied:

$$Q_{\text{idle}} = \frac{I_0}{\mathcal{R}} [(1 - \delta)(1 + \cos\varphi_B) + 2\delta], \quad \delta = 10^{-\text{ER}/10} \quad (7)$$

where δ accounts for the finite extinction ratio ER. Equation (7) omits the Joule heating of the photodiode bias and the dissipation of the delivered microwave power. Both are small beside the absorbed optical power in the present regime, but an experimental heat budget should include them.

Link Calibration from the Received Waveforms

The MATLAB receiver does not rely on the nominal modulator settings. For every run, it fits the noise-free current $i_c(t)$ against the drive reference with the model below:

$$i_c(t) = c_0 + c_1 \cos[kV(t - d)] + c_2 \sin[kV(t - d)] \quad (8)$$

where d is the electrical delay, found by cross-correlation, and c_0, c_1, c_2 are least-squares coefficients. The fit of (8) returns $I_0 = c_0$ and $\varphi_0 + kV_b = \text{atan2}(-c_2, c_1)$, from which the actual φ_B and β follow. The received optical power, and hence the effective attenuation, follows from $P_{\text{max}} = 2I_0/\mathcal{R}$. A pilot block swept V_b from -4 to 4 V at 20 dB and fitted the idle current to (4) to obtain k and φ_0 . The production design then converted each target φ_B and β into V_b and V_{pk} .

Shot-Noise Realization

During commissioning, the noisy and noise-free outputs of OptiSystem were identical sample by sample. The PIN photodiode model of this version reports its shot noise as a separate noise spectrum attached to the signal rather than adding it to the samples. The MATLAB receiver therefore realizes the shot noise from the noise-free current according to the expression below:

$$n[m] = \sqrt{q i_c[m] f_s} w[m], \quad H_B(f) = \frac{105}{p^4 + 10p^3 + 45p^2 + 105p + 105}, \quad p = j 2.1139 \frac{f}{f_c} \quad (9)$$

where m is the sample index, $w[m]$ are independent standard normal numbers, and $H_B(f)$ is the fourth-order Bessel response normalized to a 3-dB frequency $f_c = 20$ GHz. The noise $n[m]$ is filtered by H_B and added to i_c . The per-sample variance $q i_c f_s$ corresponds to the one-sided density of (6) over the band $f_s/2$. The random seed is derived from the run identifier and stored with the run. No other noise source is added, so noise is not counted twice. To check the realization, the MATLAB receiver computes the ratio

$$r_s = \frac{\hat{S}_I(f_q)}{2q\bar{I}} \quad (10)$$

where $\hat{S}_I(f_q)$ is the periodogram of the noise averaged over $f_q \pm 0.5$ GHz and $\bar{I}$ is the window-averaged current. A correct realization gives $r_s \approx |H_B(f_q)|^2 \approx 0.96$.

Transmon Model and Gate Metric

The transmon [9], a widely used superconducting qubit, is modeled with three levels in the laboratory frame, without the rotating-wave approximation, through the Hamiltonian below. Its undriven part is the weakly anharmonic oscillator form obtained from the transmon Hamiltonian of [9], and its drive term couples the received current to the qubit in the laboratory-frame manner of [10]:

$$\frac{H(t)}{\hbar} = \omega_q \hat{b}^* \hat{b} + \frac{\alpha}{2} \hat{b}^* \hat{b}^* \hat{b} \hat{b} + g i_d(t) (\hat{b} + \hat{b}^*) \quad (11)$$

where $\hat{b}$ is the lowering operator truncated to levels $|0\rangle, |1\rangle$, and $|2\rangle$, $i_d(t)$ is the received current after an ideal

high-pass filter at $f_{\text{HP}} = 1$ GHz, which represents the capacitive drive line that blocks DC and envelope components, and g (rad s⁻¹ A⁻¹) is a coupling constant that combines the cryogenic attenuation and the qubit coupling. Because g is not known for a specific device, it is calibrated for every run on the noise-free current, together with two virtual-Z frame corrections, so that the noise-free gate is an $X_{\pi/2}$ rotation. This reproduces the amplitude calibration performed in an experiment. The unitary propagator uses a first-order Magnus step in the interaction picture [11], with the fast phases integrated exactly over each sample. An asterisk superscript on an operator or matrix denotes its conjugate transpose.

Decoherence, the loss of quantum information to the environment, is described by the Lindblad master equation below [12]:

$$\dot{\rho} = -\frac{i}{\hbar}[H(t), \rho] + \sum_k \left(L_k \rho L_k^* - \frac{1}{2} \{L_k^* L_k, \rho\} \right) \quad (12)$$

with jump operators $L_1 = \sqrt{\Gamma_1}|0\rangle\langle 1|$, $L_2 = \sqrt{2\Gamma_1}|1\rangle\langle 2|$, and $L_3 = \sqrt{2\Gamma_\phi}\hat{b}^*\hat{b}$, where ρ is the density matrix, $\Gamma_1 = 1/T_1$, and $\Gamma_\phi = 1/T_2 - 1/(2T_1)$.

The gate metric is the average gate fidelity of the computational block in the frame rotating at ω_d in the form given in [14] (see [13] for the general formula):

$$F = \frac{\text{Tr}(MM^*) + |\text{Tr}(U_t^* M)|^2}{6} \quad (13)$$

where M is the 2×2 block of the propagator on $\{|0\rangle, |1\rangle\}$ and U_t is the ideal $X_{\pi/2}$ gate. Population that leaves the computational block reduces $\text{Tr}(MM^*)$, so (13) accounts for leakage. For the open-system case, F is obtained by averaging the state fidelity over the six eigenstates of the Pauli operators. The average infidelity $\varepsilon = 1 - F$, which is the gate error used in this paper, is separated into three parts, as follows:

$$\varepsilon_{\text{total}} \approx \varepsilon_{\text{dec}} + \varepsilon_{\text{link}}, \quad \varepsilon_{\text{link}} = \varepsilon_{\text{coh}} + \varepsilon_{\text{noise}} \quad (14)$$

where ε_{coh} is the unitary error of the noise-free received waveform (deterministic distortion), $\varepsilon_{\text{noise}}$ is the additional unitary error caused by shot noise, and ε_{dec} is the contribution of T_1 and T_2 computed with the noise-free waveform. The leakage L is the average population outside the computational block after the gate. To reduce the variance of $\varepsilon_{\text{noise}}$, whose single-realization distribution is heavy-tailed, each run was evaluated with $N_{\text{MC}} = 16$ independent noise realizations under the same calibration.

Closed-Form Link-Error Model

Shot noise acts on the qubit as a white random drive. By Fermi's golden rule, it causes transitions $|0\rangle \leftrightarrow |1\rangle$ at the rate $g^2 S_I/2$ and $|1\rangle \rightarrow |2\rangle$ at twice that rate, because the matrix element of $\hat{b} + \hat{b}^*$ between $|1\rangle$ and $|2\rangle$ is $\sqrt{2}$. Averaging over input states gives

$$\varepsilon_{\text{link}} \approx C g^2 \int_w S_I(t) dt, \quad C = \frac{2}{3} \quad (15)$$

where the integral runs over the integration window w of length $T_w = \tau + 4$ ns. The calibration fixes g through the rotation angle, $\theta = g I_0 \sin\varphi_B G_{\text{rot}}(\beta)$ with $\theta = \pi/2$. Substituting this relation and (5)–(6) into (15) yields the proposed model:

$$\varepsilon_{\text{link}} = C_{\text{eff}} \frac{2q\theta^2 [T_w + \cos\varphi_B G_0(\beta)]}{I_0 \sin^2\varphi_B G_{\text{rot}}^2(\beta)} \quad (16)$$

$$G_{\text{rot}}(\beta) = \int_w 2J_1(\beta a(t)) \cos\chi(t) dt, \quad G_0(\beta) = \int_w J_0(\beta a(t)) dt \quad (17)$$

where G_{rot} (s), defined in (17), is the effective pulse area of the drive and G_0 (s) sets the pulse-induced current. Equation (16) is dimensionless, since $q/(I_0 \text{ s})$ is a ratio of charges. It contains three dependences: the $1/I_0$ scaling of shot noise, the bias factor $[T_w + \cos\varphi_B G_0]/\sin^2\varphi_B$, and the compression of the drive through J_1 in G_{rot} . The golden-rule value $C = 2/3$ gives a parameter-free model, denoted Mo. In the proposed model, denoted M, the coefficient C_{eff} is estimated once from the training data.

The model is described as physics-structured because its functional form, namely the $1/I_0$ shot-noise scaling,

the bias factor $[T_w + \cos\varphi_B G_0(\beta)]/\sin^2\varphi_B$ and the Bessel-function drive compression $G_{\text{rot}}(\beta)$, is derived from the modulator transfer law and the golden-rule transition rate rather than fitted to the data, so that only one coefficient, C_{eff} , is estimated from the simulations.

Experimental Design and Data Partitions

The production design was a full factorial of the attenuation $A \in \{15, 20, \dots, 45\}$ dB, the bias angle $\varphi_B \in \{90, 110, 130, 145, 160, 170\}^\circ$, the modulation index $\beta \in \{0.2, 0.5, 0.9, 1.3\}$ and the gate duration $\tau \in \{20, 40\}$ ns, with four noise replicates per configuration. This design gives 336 configurations and 1344 runs, each requiring about 1 s of MATLAB time. The partition was made per configuration, so the replicates of one configuration always share a partition (Table II). The test set withholds two complete families: all configurations at $A = 45$ dB, which tests extrapolation to lower photocurrent, and all configurations at $\varphi_B = 170^\circ$, which tests extrapolation toward the null. The remaining configurations were divided between training and validation by a fixed hash of the configuration number, so the division does not depend on a random seed.

TABLE II. DATASET PARTITIONS AT CONFIGURATION LEVEL

Partition	Configurations	Runs	Definition
Training	206	824	$A \leq 40$ dB, $\varphi_B \leq 160^\circ$
Validation	34	136	$A \leq 40$ dB, $\varphi_B \leq 160^\circ$, 15% by hash
Test, low power	40	160	$A = 45$ dB, $\varphi_B \leq 160^\circ$
Test, low bias	48	192	$\varphi_B = 170^\circ$, $A \leq 40$ dB
Test, both	8	32	$A = 45$ dB and $\varphi_B = 170^\circ$
Verification	8	32	model-designed points at 22 and 27 dB

An automated audit checked every stored run for completeness, duplicate records, invalid values and partition leakage. It found 1387 of 1387 runs present (pilot, production and verification), no duplicates and no configuration split across partitions. The audit also revealed a numerical defect in the first processing pass: the starting value of the calibration used a lower bound of 2.2×10^{-16} A-s, which exceeds the pulse area at nanoampere photocurrents, and 264 low-power calibrations stopped at their search limit. The qubit calculation was then repeated for all runs from the stored OptiSystem waveforms with the corrected code. The original files were kept, and every result in this paper comes from the repeated calculation, in which all calibrations converged.

Reference Models and Statistical Evaluation

The regression target is $y = \log_{10} \bar{\varepsilon}_{\text{link}}$, where $\bar{\varepsilon}_{\text{link}}$ is the mean over the four replicates, that is, over 64 noise realizations. Four predictors were compared. Model B0 is a quadrature small-signal shot-noise model of the type used in [1], [3]: it takes $\langle i \rangle = I_0$, a drive current $I_0 \beta X(t)$ and the window-averaged noise, and has one fitted offset. Model B2 is a polynomial in $(\varphi_B, \beta, \log_{10} I_0, \tau)$ fitted by least squares; degrees 1–3 were tried and degree 3 was selected. Model B3 is gradient boosting [15] on the same inputs; 200 trees of depth 3 were selected from a grid of 200 or 500 trees and depths of 2–4. The proposed model M is given below:

$$\log_{10} \varepsilon_{\text{link}} = \log_{10} \varepsilon_{M0} + \kappa_0 + \sum_k \kappa_k h_k \quad (18)$$

where ε_{M0} is (16) with $C_{\text{eff}} = 2/3$, κ_0 is a calibrated offset and h_k are optional corrections chosen from $\{\beta^2, \cos\varphi_B, \beta^2 \cos\varphi_B, 1/\tau\}$ with coefficients κ_k . A correction was accepted only if it lowered the validation error by more than 5%. For interpretation, M0 and a golden-rule model evaluated with the simulated coupling and the window-averaged noise density (B1) are also reported. All selection used the training and validation sets; the test set was used once, after selection.

For n configurations with observed y_j and predicted $\hat{y}_j$, the statistics below were computed:

$$\text{RMSE}_{\log} = \sqrt{\frac{1}{n} \sum_{j=1}^n (\hat{y}_j - y_j)^2}, \quad b = \frac{1}{n} \sum_{j=1}^n (\hat{y}_j - y_j) \quad (19)$$

$$e_j = 10^{\hat{y}_j - y_j} - 1, \quad R_{\log}^2 = 1 - \frac{\sum_j (\hat{y}_j - y_j)^2}{\sum_j (y_j - \bar{y})^2} \quad (20)$$

$\text{RMSE}_{\log}$ is the typical prediction error in decades (powers of ten); a value of 0.040 corresponds to a factor of $10^{0.040} = 1.10$. The bias b shows systematic over-prediction ($b > 0$) or under-prediction ($b < 0$). The relative error e_j is summarized by its median and maximum absolute values, and $R_{\log}^2$ is the fraction of the variance of y explained by the model. The 95% confidence interval of $\text{RMSE}_{\log}$ was obtained with the percentile bootstrap [16], using 2000 resamples of configurations. The Monte Carlo standard error of each target follows from the expression below:

$$\text{SE}_{\log} = \frac{s_{\varepsilon}}{\bar{\varepsilon}_{\text{link}} \sqrt{N_{\text{MC}} N_{\text{rep}}} \ln 10} \quad (21)$$

where s_{ε} is the standard deviation of $\varepsilon_{\text{noise}}$ across realizations, $N_{\text{MC}} = 16$ and $N_{\text{rep}} = 4$. No model can be expected to predict a target more precisely than this standard error. Finally, the power law $\varepsilon_{\text{link}} \propto I_0^s$ was tested by fitting $\log_{10} \varepsilon_{\text{link}} = a + s \log_{10} I_0$ by least squares for every combination of (φ_B, β, τ) in the training and validation sets and reporting the mean and standard deviation of s .

RESULTS AND DISCUSSION

This section first verifies the simulated link, then identifies what limits the gate error, compares the proposed model with the reference models, and finally applies the model to design and verifies the designs with new simulations.

Link Waveforms and Calibration

Fig. 3(a) shows a drive voltage for $\varphi_B = 145^\circ$, $\beta = 0.5$ and $\tau = 20$ ns: a 5.052 GHz carrier under a Gaussian envelope, offset by the bias voltage. Fig. 3(b) shows the corresponding photocurrents at 30 dB. The noise-free current from OptiSystem follows the envelope and settles to the idle level outside the pulse, and the noisy current illustrates the size of the shot noise at $I_0 = 0.39 \mu\text{A}$. Fig. 3(c) shows the pilot calibration. The idle current followed (4) with a relative RMS residual of 3.5×10^{-4} , and the fit gave $k = 0.7858 \text{ rad/V}$ ($V_{\pi} = 3.998 \text{ V}$), $\varphi_0 \approx 0$ and $I_0 = 4.00 \mu\text{A}$ at 20 dB. The minimum current at the null was 1/1000 of the maximum, as set by the 30 dB extinction ratio.




Fig. 3. (a) Modulator drive voltage for $\varphi_B = 145^\circ$, $\beta = 0.5$, $\tau = 20$ ns. (b) Noise-free (OptiSystem) and noisy received photocurrent at 30 dB attenuation. (c) Pilot calibration: idle photocurrent versus bias voltage with the fitted modulator law.

In the production runs, the fitted bias angle differed from its target by at most 0.36° , the fitted modulation index matched its target, and the estimated attenuation exceeded the set value by 0.080 ± 0.045 dB (mean $\pm$ standard deviation), which is the combined loss of fiber and fork. The noise ratio (10) was 0.962 ± 0.095 , in agreement with the expected 0.96. The calibrated coupling exceeded the analytic value $\theta/[I_0 \sin \varphi_B G_{\text{rot}}(\beta)]$ by 1.3–3.6%, consistent with the 0.98 magnitude response of the Bessel filter at 5.05 GHz. Halving the time step changed ε_{coh} and ε_{dec} by less than 0.1%, which confirms convergence of the propagator.

Physics of the Link-Induced Error

Deterministic distortion was negligible: $\varepsilon_{\text{coh}} \leq 2.7 \times 10^{-8}$ in all runs, although the second-harmonic current rose from about -62 dBc at quadrature to between -13 and $+3.6$ dBc at $\varphi_B = 170^\circ$. The calibration absorbs the amplitude change caused by compression, and the harmonic lies far from the qubit transitions. The link error was therefore set by shot noise.

Fig. 4(a) plots $\varepsilon_{\text{link}}$ against I_0 for $\beta = 0.9$ and $\tau = 20$ ns at the six bias angles. The points are simulation means with 95% intervals (± 1.96 SE); the lines are model M. All series are straight lines of slope -1 on the logarithmic axes, and the fitted exponent over 40 series was $s = -1.004 \pm 0.018$, as (16) predicts for shot noise. The curves for 90° – 145° nearly coincide, while those for 160° and 170° lie higher because the useful drive $\sin\varphi_B$ becomes small near the null.

A lower error at a given I_0 does not by itself imply a better link, because heat also depends on the bias. Since $\varepsilon_{\text{link}} \propto 1/I_0$ and $Q_{\text{idle}} \propto I_0$, the product $\varepsilon_{\text{link}}Q_{\text{idle}}$ does not depend on the optical power and serves as a figure of merit. Fig. 4(b) shows this product for $\tau = 20$ ns at the four modulation indices; points are averages over the seven attenuations and lines are model M. The product falls monotonically from quadrature toward the null, by factors of 12 at $\beta = 1.3$ and 95 at $\beta = 0.2$ between 90° and 170° according to model M, and it falls with increasing β . At a fixed error, a link biased toward the null therefore dissipates less idle power at the cold stage.

Fig. 4(c) breaks down the total error at $A = 25$ dB and $\beta = 0.9$ for both gate durations. The decoherence part is constant, $\varepsilon_{\text{dec}} = 3.11 \times 10^{-4}$ for 20 ns and 5.67×10^{-4} for 40 ns, because it depends only on T_1, T_2 and the gate time. The link part is split into leakage and error within the computational subspace. Leakage accounted for a median 60% of $\varepsilon_{\text{link}}$ over all configurations. At this photocurrent the link error is smaller than the decoherence error except at 170° , where it exceeds it for the 20 ns gate.




Fig. 4. (a) Link-induced error versus photocurrent scale I_0 at $\beta = 0.9$, $\tau = 20$ ns; points are simulation means with 95% intervals and lines are model M. (b) Power-independent product $\varepsilon_{\text{link}}Q_{\text{idle}}$ versus bias angle; points are means over attenuations with 95% intervals and lines are model M. (c) Error budget at $A = 25$ dB, $\beta = 0.9$: decoherence, link error within the computational subspace and leakage.

Model Comparison

Table III lists the errors of all models, and Fig. 5 compares predicted and simulated errors on the test families. Model M reached $\text{RMSE}_{\log} = 0.040$ on the test set, with a 95% bootstrap interval of 0.033–0.047, a bias of +0.002 and a median relative error of 5.7%. This error is close to the median Monte Carlo standard error of the targets, 0.034 decades, so most of the remaining scatter is simulation noise rather than model error. Its training (0.036) and validation (0.029) errors were similar to the test error, which shows that the single coefficient did not overfit. No correction term passed the 5% validation criterion; the best candidate improved the validation error from 0.0292 to 0.0283, or 3%. The calibrated coefficient was $C_{\text{eff}} = 0.835$, 1.25 times the golden-rule value. The zero-parameter model Mo already achieved 0.104 decades with a uniform negative bias, which the coefficient removes. The origin of the factor of 1.25 was not isolated; state-dependent weighting of leakage during the driven evolution is a plausible cause.

TABLE III. PREDICTION ERROR OF THE MODELS (CONFIGURATION LEVEL, $\log_{10}$ UNITS)

Model	Train RMSE	Val. RMSE	Test RMSE [95% CI]	Test, low power	Test, low bias	Test bias b	Median / max e (%)
Bo quadrature small-signal	0.138	0.195	0.408 [0.345, 0.471]	0.158	0.518	-0.195	37.2 / 104
B1 golden rule, window noise	0.164	0.159	0.333 [0.291, 0.372]	0.154	0.414	-0.267	35.0 / 81.0
B2 cubic polynomial	0.032	0.037	0.155 [0.135, 0.175]	0.051	0.203	-0.083	16.2 / 52.5
B3 gradient boosting	0.029	0.071	0.412 [0.373, 0.452]	0.435	0.298	-0.363	61.1 / 88.6
Mo golden rule, $C = 2/3$	0.104	0.094	0.104 [0.097, 0.110]	0.089	0.115	-0.096	20.9 / 33.3
M proposed	0.036	0.029	0.040 [0.033, 0.047]	0.049	0.034	+0.002	5.7 / 36.8

The reference models fail in ways that follow from their structure. The cubic polynomial fitted the training data as closely as M (0.032) but its error rose to 0.203 on the low-bias family, where it must extrapolate a steep dependence on φ_B . Gradient boosting fitted the training data best (0.029) and extrapolated worst to the unseen power level (0.435), because a tree ensemble predicts a constant beyond the range of its training data and cannot continue the $1/I_0$ law. Model Bo predicted the low-power family reasonably (0.158), since it contains the $1/I_0$ scaling, but it under-predicted the low-bias family by a factor of about 2.3 on average (mean bias -0.354 decades), because it ignores the bias dependence of the current, the pulse-induced shot noise and the Bessel compression. Model B1, which uses the simulated coupling, fails at low bias for the same reason: the window-averaged current underestimates the current flowing during the pulse. In Fig. 5, the points of M lie on the diagonal over five decades, while the other models show systematic departures for the low-bias family (squares) and, for B3, the low-power family (circles).




Fig. 5. Predicted versus simulated link-induced error of models Bo, B2, B3 and M for the held-out test families: $A = 45$ dB (circles), $\varphi_B = 170^\circ$ (squares) and both (diamonds). The diagonal marks perfect prediction; each panel reports the test $\text{RMSE}_{\log}$.

Design Application

For a link-error target ε^* , a gate duration τ and a photocurrent budget $I_{0,\max}$, which is set by the available optical power or the linear range of the photodiode, the design problem is stated as follows:

$$\min_{\varphi_B, \beta, I_0} Q_{\text{idle}} \quad \text{s. t.} \quad \varepsilon_{\text{link}}(\varphi_B, \beta, I_0, \tau) \leq \varepsilon^*, \quad I_0 \leq I_{0,\max} \quad (22)$$

within the validated domain $90^\circ \leq \varphi_B \leq 170^\circ$, $0.2 \leq \beta \leq 1.3$ and $12.4 \text{ nA} \leq I_0 \leq 12.4 \text{ } \mu\text{A}$. Solving (16) for the photocurrent that just meets the target gives

$$I_0^{\text{req}} = \frac{K[T_w + c G_0(\beta)]}{\varepsilon^*(1 - c^2) G_{\text{rot}}^2(\beta)}, \quad c = \cos \varphi_B, \quad K = 2C_{\text{eff}} q \theta^2 \quad (23)$$

Inserting (23) into (7) with $\delta \rightarrow 0$ shows that the heat is proportional to $[T_w + c G_0]/(1 - c)$, whose derivative with respect to c is $(T_w + G_0)/(1 - c)^2 > 0$. The heat therefore decreases steadily as the bias moves toward the null, while the required photocurrent grows. The budget $I_{0,\max}$ becomes the active constraint, and the design bias is obtained by setting $I_0^{\text{req}} = I_{0,\max}$, which gives a quadratic equation in c :

$$\varepsilon^* G_{\text{rot}}^2 I_{0,\text{max}} c^2 + K G_0 c + (K T_w - \varepsilon^* G_{\text{rot}}^2 I_{0,\text{max}}) = 0 \quad (24)$$

The design bias is $\varphi_B^* = \arccos c^*$, where c^* is the smaller root in $(-1,1)$, limited to the validated domain. With a finite extinction ratio, the unconstrained optimum is finite, but for ER = 30 dB it lies beyond 170° and was not used. A grid search over (φ_B, β) confirmed that $\beta = 1.3$ is optimal when the budget is active and that $\beta \approx 1.07$ – 1.09 is optimal when the bias reaches the 170° limit. Fig. 6(a) plots φ_B^* from (24) against the budget for four targets. Each curve begins at the smallest budget for which the target can be met and rises toward the null as the budget grows; a tighter target or a shorter gate requires a larger budget for the same bias.

Verification with New Simulations

After model M was frozen, six operating points were designed at 22 and 27 dB, attenuations that do not appear in the training data, for targets from 3×10^{-5} to 3×10^{-4} . Their predicted errors and heat were recorded before the simulations were run. Two quadrature points at the same attenuations were added as references. The eight configurations were simulated with four replicates and 16 noise realizations each (Table IV).

TABLE IV. VERIFICATION AT PRE-REGISTERED DESIGN POINTS

A (dB)	τ (ns)	ε^*	φ_B ($^\circ$)	β	Predicted $\varepsilon_{\text{link}}$ (10^{-4})	Simulated $\varepsilon_{\text{link}}$ (10^{-4}) $\pm$ SE	Q_{idle} (nW)	Heat reduction
22	20	1×10^{-4}	164.6	1.30	1.00	0.93 ± 0.08	120.5	11.1
22	40	3×10^{-5}	158.7	1.30	0.30	0.29 ± 0.03	222.9	9.1
22	40	1×10^{-4}	170.0	1.07	0.99	0.98 ± 0.08	54.9	11.1
27	20	3×10^{-4}	164.1	1.30	3.00	2.77 ± 0.19	40.5	11.0
27	40	1×10^{-4}	159.5	1.30	1.00	1.08 ± 0.08	65.5	9.3
27	40	3×10^{-4}	169.7	1.09	2.99	2.73 ± 0.23	18.3	11.1
22	20	reference	90.0	1.30	0.43	0.41 ± 0.03	3157	–
27	20	reference	90.0	1.30	1.36	1.34 ± 0.12	998	–

The prediction errors of the eight configurations ranged from -8.5% to $+8.0\%$, with $\text{RMSE}_{\log} = 0.026$ (Fig. 6(b)). Five designed points met their target in the mean. The sixth ($\varepsilon^* = 10^{-4}$, $\tau = 40$ ns, 27 dB) exceeded it by 8%, within its 95% Monte Carlo interval. A designer should therefore apply a margin of about 10%, roughly two standard deviations of the model error, by designing for $\varepsilon^*/1.1$. The simulated idle heat agreed with (7) to within 1.8–3.2%. The two quadrature references were predicted within -4.9% and -1.7% , which confirms the model at the conventional operating point.

The last column of Table IV compares each design with the best quadrature-biased link that meets the same target. That link was found from the verified model by optimizing β at $\varphi_B = 90^\circ$. The low-bias designs required 9.1–11.1 times less idle optical power at the photodiode (Fig. 6(c)). The saving comes from the idle current, which is proportional to $1 + \cos \varphi_B$. This saving has two costs: a larger I_0 , that is, more optical power from the laser and modulator at room temperature, and a strong second harmonic at $2\omega_q$, between -3.3 and $+1.8$ dBc at the designed points compared with -63 dBc at quadrature.




Fig. 6. (a) Design bias from (24) versus photocurrent budget for four targets ($\beta = 1.3$). (b) Simulated versus pre-registered predicted link error of the six designed points and the two quadrature references, with 95% intervals. (c) Idle optical heat of the designed points (simulated) and of the best quadrature design at the same target (model).

Limitations

The study is limited in the following respects. All results come from simulation, and no experiment was performed. OptiSystem supplied the deterministic link, whereas the shot noise was realized in MATLAB from the standard $2q i(t)$ model; the agreement of (16) with simulations built on the same noise model is a consistency check rather than an independent validation.

Relative intensity noise, laser phase noise, thermal noise, dark current, photodiode saturation and the frequency response of a photodiode at cryogenic temperature were not included; relative intensity noise penalizes the low-bias region less than shot noise does and could move the optimum. The coupling g was calibrated for every configuration, so the model predicts the error of a calibrated gate and does not determine the cold attenuation. The second harmonic did not affect the single-qubit gate in the three-level model, but it could drive a readout resonator (10.9 GHz in [1]) or a neighboring qubit, so a low-bias design should include a cryogenic low-pass filter.

Only one qubit, one tone and one pulse family were studied, with device parameters taken from a single experiment [1]. The heat metric counts only the absorbed optical power. The optimum lies at the edge of the validated domain, and extrapolation beyond $\phi_B = 170^\circ$ or $\beta = 1.3$, where the extinction ratio dominates, is not supported. Finally, C_{eff} is an empirical coefficient.

CONCLUSION

An OptiSystem and MATLAB co-simulation and a physics-structured gate-error model with one calibrated coefficient were developed for analog microwave-over-fiber drive of a transmon. The model maps the modulator bias, the modulation index, the photocurrent and the gate duration to the link-induced gate error. In 1344 simulations the error was set by shot noise, with deterministic distortion below 2.7×10^{-8} ; it scaled as $1/I_0$, and about 60% of it was leakage. On configurations withheld from fitting, the model predicted the error with $RMSE_{log} = 0.040$ (median error 5.7%), against 0.155–0.412 for a quadrature small-signal model, a cubic polynomial and gradient boosting. A closed-form rule for the bias at a given photocurrent budget produced six designs that met pre-registered targets within 8.5% in new simulations and needed 9.1–11.1 times less idle optical power at the cold photodiode than the best quadrature design. Measurements with a cryogenic photodiode are needed to confirm these results, and extensions to intensity noise, multiplexed tones and readout-resonator constraints are planned.

ACKNOWLEDGMENT

The author would like to express his gratitude to the Labs Division at the College of Engineering, Al-Iraqia University, for providing a licensed copy of OptiSystem software and a high-performance PC to conduct this research.

ETHICAL DECLARATION

Conflict of interest: No declaration required. **Financing:** No reporting required. **Peer review:** Double anonymous peer review.

REFERENCES

- [1] F. Lecocq, F. Quinlan, K. Cicak, J. Aumentado, S. A. Diddams, and J. D. Teufel, "Control and readout of a superconducting qubit using a photonic link," *Nature*, vol. 591, no. 7851, pp. 575–579, 2021, doi: [10.1038/s41586-021-03268-x](https://doi.org/10.1038/s41586-021-03268-x).
- [2] J. Capmany and D. Novak, "Microwave photonics combines two worlds," *Nature Photonics*, vol. 1, no. 6, pp. 319–330, 2007, doi: [10.1038/nphoton.2007.89](https://doi.org/10.1038/nphoton.2007.89).
- [3] T. Nakamura, D. Lee, J. Horng, F. Lecocq, J. Teufel, and F. Quinlan, "Cryogenic photonic link using an extended-InGaAs photodiode and short pulse illumination toward high-fidelity drive of superconducting qubits," *Optica Quantum*, vol. 3, no. 3, pp. 221–227, 2025, doi: [10.1364/OPTICAQ.546795](https://doi.org/10.1364/OPTICAQ.546795).
- [4] H. K. Warner, J. Holzgrafe, B. Yankelevich, D. Barton, S. Poletto, C. J. Xin, N. Sinclair, D. Zhu, E. Sete, B. Langley, E. Batson, M. Colangelo, A. Shams-Ansari, G. Joe, K. K. Berggren, L. Jiang, M. J. Reagor, and M. Lončar, "Coherent control of a superconducting qubit using light," *Nature Physics*, vol. 21, pp. 831–838, 2025, doi: [10.1038/s41567-025-02812-0](https://doi.org/10.1038/s41567-025-02812-0).
- [5] G. Arnold, T. Werner, R. Sahu, L. N. Kapoor, L. Qiu, and J. M. Fink, "All-optical superconducting qubit readout," *Nature Physics*, vol. 21, pp. 393–400, 2025, doi: [10.1038/s41567-024-02741-4](https://doi.org/10.1038/s41567-024-02741-4).
- [6] J. P. G. van Dijk, E. Kawakami, R. N. Schouten, M. Veldhorst, L. M. K. Vandersypen, M. Babaie, E. Charbon, and F. Sebastiano, "Impact of classical control electronics on qubit fidelity," *Physical Review Applied*, vol. 12, no. 4, Art. no. 044054, 2019, doi: [10.1103/PhysRevApplied.12.044054](https://doi.org/10.1103/PhysRevApplied.12.044054).
- [7] H. Ball, W. D. Oliver, and M. J. Biercuk, "The role of master clock stability in quantum information processing," *npj Quantum Information*, vol. 2, Art. no. 16033, 2016, doi: [10.1038/npjqi.2016.33](https://doi.org/10.1038/npjqi.2016.33).
- [8] C. H. Cox, E. I. Ackerman, G. E. Betts, and J. L. Prince, "Limits on the performance of RF-over-fiber links and their impact on device design," *IEEE Transactions on Microwave Theory and Techniques*, vol. 54, no. 2, pp. 906–920, 2006, doi: [10.1109/TMTT.2005.863818](https://doi.org/10.1109/TMTT.2005.863818).
- [9] J. Koch, T. M. Yu, J. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, "Charge-insensitive qubit design derived from the Cooper pair box," *Physical Review A*, vol. 76, no. 4, Art. no. 042319, 2007, doi: [10.1103/PhysRevA.76.042319](https://doi.org/10.1103/PhysRevA.76.042319).
- [10] F. Motzoi, J. M. Gambetta, P. Rebentrost, and F. K. Wilhelm, "Simple pulses for elimination of leakage in weakly nonlinear qubits," *Physical Review Letters*, vol. 103, no. 11, Art. no. 110501, 2009, doi: [10.1103/PhysRevLett.103.110501](https://doi.org/10.1103/PhysRevLett.103.110501).
- [11] S. Blanes, F. Casas, J. A. Oteo, and J. Ros, "The Magnus expansion and some of its applications," *Physics Reports*, vol. 470, no. 5–6, pp. 151–238, 2009, doi: [10.1016/j.physrep.2008.11.001](https://doi.org/10.1016/j.physrep.2008.11.001).
- [12] G. Lindblad, "On the generators of quantum dynamical semigroups," *Communications in Mathematical Physics*, vol. 48, no. 2, pp. 119–130, 1976, doi: [10.1007/BF01608499](https://doi.org/10.1007/BF01608499).
- [13] M. A. Nielsen, "A simple formula for the average gate fidelity of a quantum dynamical operation," *Physics Letters A*, vol. 303, no. 4, pp. 249–252, 2002, doi: [10.1016/S0375-9601\(02\)01272-0](https://doi.org/10.1016/S0375-9601(02)01272-0).
- [14] L. H. Pedersen, N. M. Møller, and K. Mølmer, "Fidelity of quantum operations," *Physics Letters A*, vol. 367, no. 1–2, pp. 47–51, 2007, doi: [10.1016/j.physleta.2007.02.069](https://doi.org/10.1016/j.physleta.2007.02.069).
- [15] J. H. Friedman, "Greedy function approximation: A gradient boosting machine," *The Annals of Statistics*, vol. 29, no. 5, pp. 1189–1232, 2001, doi: [10.1214/aos/1013203451](https://doi.org/10.1214/aos/1013203451).
- [16] B. Efron, "Bootstrap methods: Another look at the jackknife," *The Annals of Statistics*, vol. 7, no. 1, pp. 1–26, 1979, doi: [10.1214/aos/1176344552](https://doi.org/10.1214/aos/1176344552).