Simulation et optimisation du contôle et de la mesure du qubit supraconducteur

par

Ross Shillito

Thèse présentée au département de physique

en vue de l'obtention du grade de docteur ès sciences (Ph.D.)

FACULTÉ des SCIENCES

UNIVERSITÉ de SHERBROOKE

Sherbrooke, Québec, Canada, 18 septembre 2023

Le Lundi, 18 septembre 2023

le jury a accepté la thèse de M Ross Shillito dans sa version finale.

Membres du jury

Professeur Alexandre Blais
Directeur de recherche
Département de physique

Professeur Michel Pioro-Ladrière
Membre interne
Département de physique

Professeur Ioan M Pop
Membre externe
University of Stuttgart

Professeur Ion Garate
Président rapporteur
Département de physique

To my family and friends in New Zealand, who I love and miss dearly;
À mes amis Québécois, les personnes les plus chaleureuses et les plus accueillantes;
Y al amor de mi vida.

Sommaire

Le domaine des circuits supraconducteurs a connu une croissance spectaculaire au cours de la dernière décennie et constitue une architecture de premier plan pour l'informatique quantique. Le succès de cette architecture tient en partie à la simplicité du contrôle, à l'évolutivité et à la durée de vie de plus en plus longue du qubit transmon, qui est devenu dominant dans la plupart des processeurs. Bien que la fidélité de la plupart des opérations quantiques sur ces dispositifs dépasse généralement $98 - 99\%$, des améliorations significatives des performances seront nécessaires pour parvenir à un calcul tolérant aux fautes. L'une des principales limites actuelles de ces opérations sur les qubits est due aux "fuites", c'est-à-dire au peuplement des états d'énergie supérieure du qubit. Ces effets sont difficiles à appréhender d'un point de vue numérique et théorique, car de nombreuses transitions résultent de processus hors résonance qui sont généralement ignorés afin de rendre le problème traitable. En outre, beaucoup de ces effets sont causés par les niveaux d'énergie plus élevés du transmon, présents hors de son potentiel de confinement. Ceux-ci sont également négligés dans les simulations et les études théoriques pour des raisons de simplicité, ce qui entraîne une mauvaise caractérisation de la dynamique.

Dans cette thèse, j'étudie comment nous pouvons comprendre, simuler avec précision et supprimer les processus à l'origine d'erreurs résultant de mesures et de contrôles forts sur les qubits transmon. Dans le cadre de ma recherche, je présente une méthode permettant de simuler efficacement des systèmes en présence de commandes à oscillation rapide, et comment celle-ci nous permet d'optimiser les commandes afin d'améliorer la fidélité des opérations. Cette méthode permet une accélération significative par rapport aux intégrateurs numériques traditionnels, et une version de ce solveur est maintenant incluse dans Qiskit Dynamics, un logiciel d'IBM.

Je démontre en outre comment un pilote de mesure suffisamment forts pendant la lecture peut non seulement entraîner des fuites, mais aussi la population d'états en dehors du potentiel de confinement du transmon, ce que nous appelons l'"ionisation". La compréhension

de l'impact de ces effets parasites - et de la manière dont ils peuvent être évités - sera cruciale pour l'optimisation et le fonctionnement de la lecture dans les futurs dispositifs.

Je conclurai par mes travaux en collaboration avec l'ETH Zurich sur les résultats de lecture de pointe utilisant des transmons et des filtres de Purcell réglables en fonction du flux, et j'expliquerai comment le réglage de la fréquence du transmon peut améliorer la fidélité de la lecture. Cela ouvre une autre voie pour améliorer la lecture des qubits supraconducteurs et nous rapproche potentiellement des applications utiles de ces dispositifs.

The field of superconducting circuits has grown dramatically in the last decade and is a leading architecture for quantum computation. Part of this architecture's success is the simple control, scalability, and increasingly long lifetime of the transmon qubit, which became dominant in most processors. Whilst the fidelity of most quantum operations on these devices generally exceeds $98-99\%$, there will need to be significant improvements in performance to reach fault-tolerant computation. A major limitation of these qubit operations today is from 'leakage', or from populating the higher energy states of the qubit. These effects are difficult to capture from both a numerical and a theoretical point of view, as many of the transitions result from off-resonant processes that are typically disregarded in order to make the problem tractable. Furthermore, many of these effects result from the higher energy levels of the transmon outside of its confining potential, which are also neglected from simulations and theoretical studies for simplicity, resulting in a mischaracterization of the dynamics.

In this thesis, I investigate how we can understand, accurately simulate, and suppress error-causing processes resulting from strong measurement and control drives on transmon qubits. As part of my research, I present a method of efficiently simulating systems in the presence of fast-oscillating drives, and how this allows us to optimize the controls to improve the fidelity of operations. This method provides a significant speed-up over traditional numerical integrators, and a version of this solver is now included in Qiskit Dynamics, a software package by IBM.

I additionally demonstrate how sufficiently strong measurement drives during readout can not only result in leakage but also in the population of states outside of the transmon's confining potential, which we refer to as 'ionization'. Understanding the impacts of these spurious effects – and how they can be avoided – will be crucial in the optimization and operation of readout in future devices.

I conclude with my work in collaboration with ETH Zurich on state-of-the-art readout

results using flux-tunable transmons and Purcell Filters, and provide an understanding of how tuning the transmon frequency can improve readout fidelities. This opens up another avenue to improving superconducting qubit readout and potentially brings us closer to useful applications of these devices.

Thanks

The last four years have been marked by numerous challenges, and my personal journey has been no exception. It all began on August 16th, 2019, when I arrived in Québec, excited and eager to embark on a new chapter in my life, and blissfully unaware of the impending global pandemic. Just over a year later, I found myself returning to my home country of New Zealand. The onset of the pandemic had forced many of us to adapt to a new way of life, with remote work and virtual interactions becoming the norm. As I settled back into New Zealand, I transformed into a semi-nocturnal virtual rectangle, constantly participating in blurry Zoom meetings, struggling to maintain a semblance of normalcy amidst the chaos.

For the next year and a half, my daily routine consisted of navigating time zones, engaging in endless Slack conversations, and relying on the ever-amusing giphys to inject some much-needed humour into our interactions. It was a challenging period, but it was made possible by the unwavering support, understanding and patience of my supervisor, Alexandre Blais. My immeasurable thanks for everything you have done over the last four years.

If I were to list everyone who were invaluable to my progression over the last few years, this section would be as long as the rest of the thesis. As such, I've decided to restrict this list to close family and academic friends.

- A big thank you to Agustin Di Paolo, who taught me how to use Slack /giphys, alongside being an amazing mentor at the start of the PhD. You were truly inspirational.

- As the Dysolve project was developing, we quickly uncovered the need to batch and benchmark the code on the university servers. Not only did Elie Genois tackle this part of the project, navigating my one-too-many-for-loops and variable names that brought up several questions (hmm...), he also played a huge role in developing and optimizing the code well past its original standing. This work wouldn't have been

possible without him.

- Comme indiqué précédemment, une grande partie du doctorat a été passée à distance. Heureusement, il y avait un autre membre nocturne du groupe : Catherine Leroux. Pendant cette période en Nouvelle-Zélande, Kat m'a vraiment fait sentir que je faisais partie du groupe, et elle m'a beaucoup soutenue à mon retour au Canada. Kat a une intelligence spectaculaire et j'ai appris énormément en travaillant sur plusieurs projets avec elle;

- Many thanks to Francois Swiadek, for coping with my early morning coffee-induced rants about inconsistencies in SNR formulae and for writing a fantastic paper;

- A huge thank you to my sister, Dr. Georgina Shillito. Whilst I did my PhD for the love of physics, the sibling rivalry never failed to provide extra motivation when I needed it. You've been an incredible support and inspiration.

- Thank you so much to my novia, Palmira, for everything you've done over the last year. You've been my home away from home, and your love and support has helped me immeasurably. Te amo.

- Finally, to my mum and dad, who've been an amazing support throughout the PhD and the pandemic. I love you with all my heart.

Contents

List of Figures

Chapter 1

Introduction

In 1981, Richard Feynman proposed a radical idea: To build a new type of computer which harnesses quantum mechanics at its core. His argument for this was rather simple [4]:

Nature isn't classical, dammit, and if you want to make a simulation of Nature, you'd better make it quantum mechanical, and by golly it's a wonderful problem because it doesn't look so easy . . . the full description of quantum mechanics for a large system with R particles is given by a function which we call the amplitude to find the particles at x_1, x_2, . . . ,x_R, *and therefore because it has too many variables, it cannot be simulated with a normal computer.*

Today, more than 40 years later, our understanding of how to construct a quantum computer has grown considerably, with the story of this journey nicely summarized in Refs. [5, 6]. As of today, there are many candidates, such as spin qubits [7, 8, 9, 10], trapped ions [11, 12, 13, 14] and superconducting qubits [15, 16, 17, 18], which are all the subjects of intensive study. In particular, superconducting qubits are a strong candidate, known for their scalability and versatility and have already been shown to demonstrate useful quantum speedup [19] and error correction [20, 21]. This architecture relies on Josephson junctions, a circuit element [22, 23] which provides the nonlinearity required for useful computation and control [24]. Furthermore, the two-qubit gate performance is excellent, with errors of 3×10^{-3} reported [25].

Whilst the error rates of these qubits seem small, useful quantum computation will require dramatically lower error rates. Given that future algorithms will rely on billions of operations, the anticipated tolerable error rate for gates between logical qubits on a quantum computer is as low as 1×10^{-15} [26], orders of magnitude lower than the current transmon error rates. Quantum error correction (or QEC for short) allows us to circumvent

this issue by using many physical qubits, such as the transmon, to encode each of these logical qubits. Unfortunately, current predictions for the surface code are in the thousands of superconducting qubits per logical qubit to achieve fault tolerance [27]. The number of required qubits per logical qubit will largely depend on the highest error rates for the system.

Many of the current gate operations on superconducting qubits are currently limited by the qubit coherence time. Improvements to these coherence times are anticipated as fabrication and materials research develops – for example, the use of tantalum has greatly increased qubit lifetimes [28]. Nevertheless, it is desirable to perform gate operations as fast as possible to limit these unavoidable decoherence effects. As the lifetime of these qubits increases, coherent processes – spurious qubit rotations caused by miscalibration of the drive, crosstalk and more – will contribute significantly to the total error of the quantum operation. The control of these qubits will need to become more sophisticated which necessitates an increased understanding and improved simulation requirements.

In this thesis, I investigate how we can understand, accurately simulate, and suppress some of these error-causing processes. One such major performance limitation of superconducting qubit operations today is from 'leakage', or from populating the higher energy states of the qubit. Leakage is generally incurred from *off-resonant* processes, which are hard to capture from both a numerical and theoretical viewpoint, and our understanding of such processes must be improved to continue improving gate fidelities. In Chapter 2, I detail how the transmon and its readout resonator can be accurately simulated, the approximations that are made to simply understand this system, and where such approximations fail. I then proceed to discuss a new method of efficiently simulating systems in the presence of a fast-oscillating drive in Chapter 3, and how this method allows us to optimize the controls to improve the fidelity of operations. Then in Chapter 4, I extend the same formalism to account for the aforementioned loss effects and compare it to another ubiquitous method of simulating such systems.

One of the greatest problems in the field of superconducting qubits is qubit readout, where the logical state $|0\rangle$ or $|1\rangle$ is measured. Whilst this process was predicted to be straightforward and highly accurate [15, 29], the fidelity of readout rarely exceeds 99.5% [30, 31]. This is in part due to long operation times, but also due to spurious effects and transitions invoked by driving the system [32, 33, 34]. Understanding such effects during these readout protocols is thus critical to improving our understanding. In Chapter 5, I perform an investigation into the strongly-driven transmon-readout system and observe dramatic leakage effects such that states outside of the transmon's confining well become populated. We refer to this process as 'ionization'.

A crucial component of most superconducting circuits is the Purcell filter – an additional resonator which helps to prevent the qubit from decaying into the readout circuitry [35, 36, 37]. When simulating the readout of a transmon qubit, one generally neglects the presence of this resonator due to the great associated increase in complexity. In many cases, simulating this component is vital to accurately describe the system dynamics. In addition, the parameters of this filter can be critical to the readout performance. In collaboration with ETH Zürich, I conclude in Chapter 6 with how the combined transmon-resonator-Purcell system can be efficiently modelled, and use the results to describe some state-of-the-art experimental readout fidelities. I provide additional information regarding the optimal parameters of this system in the appendix.

Chapter 2

Key Concepts for Circuit QED Simulation

"If all you have is a hammer, then everything looks like a nail."

– Abraham Maslow, *The Psychology of Science*

In this chapter, I lay the foundations needed to perform accurate simulations of the transmon qubit and provide some context for the succeeding chapters.

2.1 Diagonalizing the Transmon Qubit

The transmon [38] is one of the most ubiquitous qubits used in current superconducting architectures, and is a central focus of this thesis. Part of the transmon's appeal is its simplicity, consisting only of a shunt capacitance and a Josephson junction [22], see the circuit in Fig. 2.1(a). By performing a standard circuit quantization procedure [39], we readily arrive at the Hamiltonian for this qubit,

$$\hat{H} = 4E_c(\hat{n} - n_g)^2 - E_J \cos(\hat{\varphi}), \tag{2.1}$$

where $\hat{n}$ and $\hat{\varphi}$ are the charge and phase operators respectively, with the commutation relation $[\hat{\varphi}, \hat{n}] = i$. Further, E_c, E_J are the charging and Josephson energies respectively, and n_g is an offset gate charge. This is the same Hamiltonian as the 'Cooper Pair Box' qubit [40, 41], but operated in a very different regime, with $E_J/E_C \gg 1$. It has been demonstrated

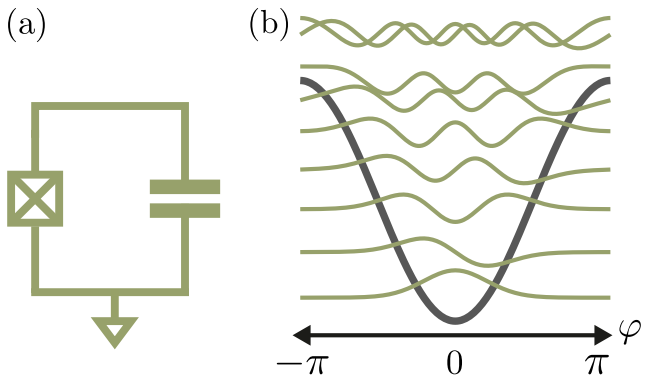

Figure 2.1 (a) Circuit of the transmon qubit. The square box represents the nonlinear Josephson junction element. (b) Representative energy spectrum and first few eigenstates, with the cosine confining potential. Adapted from Ref. [2].

in Ref. [38] that the sensitivity of the transmon's eigenenergies to the charge offset n_g diminishes exponentially as the ratio E_J/E_c increases. This exponential suppression greatly reduces the effects of charge noise on the qubit and increases the qubit's coherence time.

This simple Hamiltonian has a rich and complex energy structure, but it can be non-trivial to simulate. In the following sections, I consider three different methods for diagonalizing this Hamiltonian – the Fock, phase, and charge bases – and detail why the charge basis is the most reliable.

2.1.1 Fock Basis

The Fock basis description of the transmon qubit is excellent for pedagogical reasons and describes the lowest energy states of the system with reasonable accuracy. To motivate this expansion, we rewrite Eq. (2.1) to reveal a harmonic contribution plus a perturbation,

$$\hat{H} = 4E_c\hat{n}^2 + \frac{E_J}{2}\hat{\varphi}^2 - E_J\left[\cos(\hat{\varphi}) + \frac{1}{2}\hat{\varphi}^2\right], \tag{2.2}$$

where we have additionally dropped the gate charge dependence n_g, given that the energies are insensitive to this parameter in the transmon regime. Following the standard procedure [16, 38], we can represent the charge and phase operators in terms of annihilation and creation operators with suitable zero point fluctuations to diagonalize the harmonic contribution to Eq. (2.2):

$$\hat{\varphi} = \left(\frac{2E_c}{E_J}\right)^{\frac{1}{4}}(\hat{b}^\dagger + \hat{b}), \quad \hat{n} = \frac{i}{2}\left(\frac{E_J}{2E_c}\right)^{\frac{1}{4}}(\hat{b}^\dagger - \hat{b}). \tag{2.3}$$

We see that the zero point fluctuations in the phase variable $\hat{\varphi}$ are indeed small in comparison to $\hat{n}$, which motivates a Taylor expansion in the cosine confinement potential:

$$\hat{H} = 4E_c\hat{n}^2 + \frac{E_J}{2}\hat{\varphi}^2 - \frac{E_J}{4!}\hat{\varphi}^4 - E_J\sum_{n=3}^{\infty}\left[(-1)^n\frac{\hat{\varphi}^{2n}}{(2n)!}\right]. \tag{2.4}$$

Using the definitions of the operators in Eq. (2.3), this yields a Hamiltonian up to fourth order in $\hat{\varphi}$:

$$\hat{H}_{\text{fock}} = \sqrt{8E_JE_c}\hat{b}^\dagger\hat{b} - \frac{E_c}{12}(\hat{b}^\dagger + \hat{b})^4. \tag{2.5}$$

Finally, we can invoke the 'Rotating-Wave-Approximation' by ignoring non-energy conserving terms in Eq. (2.5), or terms that involve an unequal number of creation and annihilation operators. This yields a form of the Hamiltonian often referred to as a 'Kerr-nonlinear oscillator':

$$\begin{aligned}\hat{H}_{\text{fock,RWA}} &= (\sqrt{8E_JE_c} - E_c)\hat{b}^\dagger\hat{b} - \frac{E_c}{2}\hat{b}^{\dagger 2}\hat{b}^2 \\ &= \hat{b}^\dagger\hat{b}\left[\omega_q - \frac{E_c}{2}\left(\hat{b}^\dagger\hat{b} - 1\right)\right].\end{aligned} \tag{2.6}$$

Here, the qubit's frequency is defined by $\omega_q = (\sqrt{8E_JE_c} - E_c)$, with the anharmonicity $\alpha = -E_c$ describing the non-linearity of the oscillator [1]. This term decreases the transition frequency between the n-th and $(n-1)$-th energy state by $(n-1)E_c$, as is evident from the form of Eq. (2.6).

Whilst this Hamiltonian is excellent for aiding our understanding – it provides the frequency, anharmonicity, and a simple form for the charge and phase operators – it fails to capture the higher energy structure, which we will see in Chapter 5 has surprising and dramatic impacts on the qubit readout. Moreover, any expansion in the Fock basis eventually fails, even if an increasing number of terms in Eq. (2.4) are preserved. This is because the Hamiltonian in Eq. (2.1) is *periodic*: The eigenstates of the Hamiltonian, when expressed in the phase basis, must obey the relation $\psi(\phi + 2\pi) = \psi(\phi)$. More formally, the phase of the transmon is said to be compact, and the eigenstates strictly are defined on the interval $\phi \in [-\pi, \pi)$. On the other hand, the Fock basis is strictly *non-compact*, with wavefunctions defined over $\phi \in (-\infty, \infty)$ – consequently, it will be unable to capture the transmon's periodicity accurately.

To see this problem more explicitly, in Fig. 2.2(a) I plot the wavefunctions corresponding to the first few states of the transmon for $E_J/E_c = 50$, as calculated in the charge basis (see Sect. 2.1.3). Similarly, in Fig. 2.2(b), I plot the wavefunctions corresponding to the Fock basis in Eq. (2.4), keeping up to a 10th order Taylor expansion of the $\cos(\phi)$ term.

[1] Throughout this thesis, I use the the convention $\hbar = 1$.

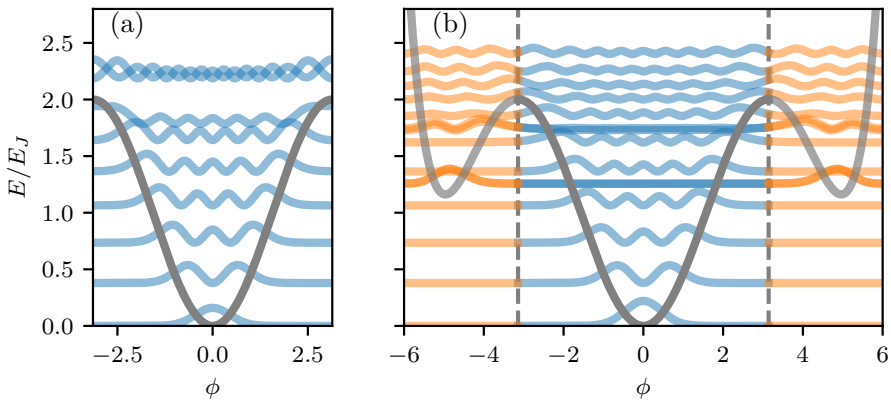

Figure 2.2 A demonstrative set eigenfunctions of the transmon qubit, plotted in the phase basis. (a) The eigenfunctions as calculated from the *charge* basis. (b) The eigenfunctions approximated in the Fock basis, where the transparent line considers a Taylor expansion of the cosine potential well up to 10th order.

Whilst the Fock basis captures the low energy dynamics well, there is a dramatic departure from the results of (a) at higher energies, with the wavefunction occupying false minima generated from the Taylor expansion at $\phi \approx \pm 5$. This problem persists when increasing the order of the Taylor expansion and is therefore challenging to circumvent in this basis.

2.1.2 Phase Basis

To resolve the shortcomings of the Fock basis, let us now return to Eq. (2.1) and consider the Hamiltonian in the *phase* basis:

$$4E_C\left(-i\frac{d}{d\phi}+n_g\right)^2\psi_n(\phi)-E_J\cos(\phi)\psi_n(\phi)=E_n\psi_n(\phi). \tag{2.7}$$

Here, I use ϕ for the phase coordinate in the phase basis to differentiate it from the operator $\hat{\varphi}$ in the Fock or charge bases. In this form of the Hamiltonian, we clearly see a relation to Bloch's theorem, with the gate charge n_g functioning as a wavevector [42]. This equation can be treated analytically, yielding solutions in the form of Mathieu functions [38, 43, 44]. To treat this problem numerically, we begin by discretizing phase space into N intervals each of length $h = 2\pi/N$, consequently yielding a vector of complex values $\{\psi_{k,n}\}$ which approximate the eigenstate $\psi_n(\phi)$ at values $\phi = -\pi + 2\pi k/N$. The simplest way of calculating the derivatives uses a central difference method:

$$\psi_i' = \frac{\psi_{i+1}-\psi_{i-1}}{2h}, \quad \psi_i'' = \frac{\psi_{i+1}-2\psi_i+\psi_{i-1}}{h^2}, \tag{2.8}$$

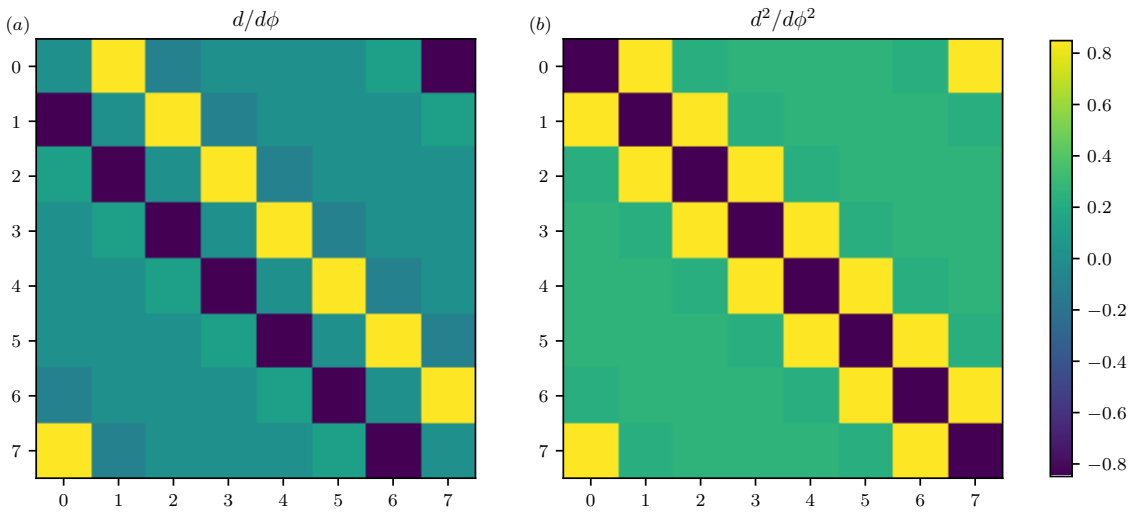

Figure 2.3 Form of the (a) first and (b) second derivative matrices for a 5 point stencil with periodic boundary conditions in phase basis with 8 divisions.

The accuracy of the phase basis diagonalization can be increased by using additional points ψ_i in the calculations of the first and second derivative, also known as a *stencil* [45], whilst preserving the sparsity of the Hamiltonian. To see this more explicitly, I plot in Fig. 2.3 an illustration of the (a) first and (b) second derivatives with 8 different phase coefficients ψ_k and a five-point stencil, although a hundred or more coefficients may be needed for accurate calculation of the eigenvalues, see Sect. 2.1.4 for more details. Note that in the case of the transmon, which is compact, the periodicity must be maintained in the structure of the derivative matrices. This can be enforced by simply setting $\psi_{i+N} = \psi_i$ – see the matrix elements connecting ψ_7 and ψ_0.

Care must be taken when expanding Eq. (2.7), as in this basis the square of the derivative matrix in Fig. 2.3 (a) will not necessarily well approximate the second derivative in (b). As such, when writing the Hamiltonian in this basis, we must expand out Eq. (2.7) explicitly:

$$\hat{H} = -4E_c\frac{d^2}{d\phi^2} - 8iE_c n_g\frac{d}{d\phi} - E_J\cos(\phi), \tag{2.9}$$

and use the definitions for each operator individually.

2.1.3 Charge Basis

We now finally consider the charge basis. Using the relation $\hat{n} = -i(d/d\phi)$, we are able to describe eigenstates of the charge operator $\hat{n}$ in the phase basis:

$$\hat{n}|n\rangle = n|n\rangle \longleftrightarrow -i\frac{d}{d\phi}e^{in\phi} = ne^{in\phi}. \tag{2.10}$$

The corresponding Hamiltonian in this basis is thus of the form:

$$\hat{H} = 4E_c \sum_{n=-\infty}^{\infty} (n - n_g)^2 |n\rangle\langle n| - \frac{E_J}{2}\left(|n\rangle\langle n+1| + |n+1\rangle\langle n|\right). \tag{2.11}$$

One can consider the charge basis as a Fourier series for representing wavefunctions in phase space. This makes it the natural choice for representing compact spaces, such as the aforementioned transmon, since the eigenstates $e^{in\phi}$ are natively periodic on the interval $[-\pi, \pi)$. Furthermore, the states above the transmon's cosine confinement potential are approximately charge eigenstates, meaning that this basis should well capture the high energy spectrum. This is discussed further in Chapter 5.

Fractional Charge Basis

Whilst the phase space of the transmon qubit is compact, it is important to note that more complex qubits such as the fluxonium [46] and L-shunted transmon [47] contain an inductive shunt, which breaks the periodicity. Such a Hamiltonian is usually approximated by the form:

$$H = 4E_c\hat{n}^2 - E_J\cos(\hat{\varphi}) + \frac{E_L}{2}(\hat{\varphi} - \varphi_{\text{ext}})^2, \tag{2.12}$$

where E_L is the energy corresponding to the inductive shunt – see Fig. 2.4 for the corresponding circuit and spectrum.

In the current form of the charge basis, we have no means of well approximating the $\hat{\phi}^2$ contribution outside of the interval $[-\pi, \pi)$. Consequently, the Fock and phase bases would seem like the more natural choices. Nevertheless, the charge basis can be used to approximate these non-compact spaces. Consider the transformation of the operators $\hat{n}$, $\hat{\varphi}$, which preserves the commutation relations

$$\hat{\varphi}' = \hat{\varphi}/N, \quad \hat{n}' = N\hat{n}, \tag{2.13}$$

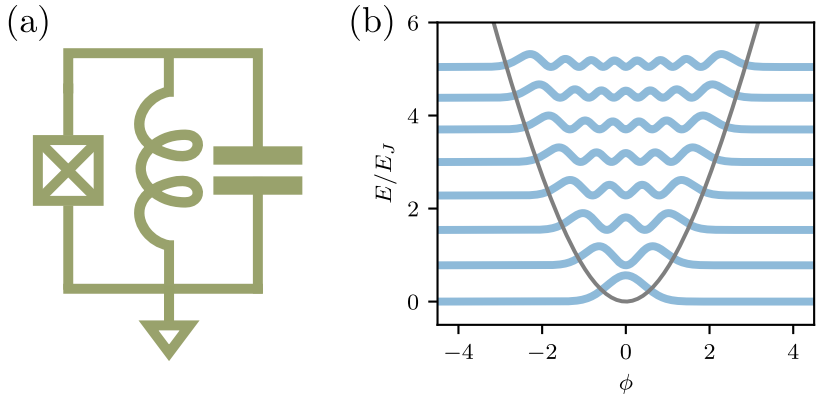

Figure 2.4 (a) Circuit Diagram for the inductively shunted transmon. (b) Spectrum of the inductively shunted transmon with $E_J/E_c = E_L/E_c = 25$.

such that the Hamiltonian becomes:

$$H = \frac{4E_c}{N^2}\hat{n}'^2 - E_J\cos(N\hat{\varphi}') + \frac{E_L}{2}(N\hat{\varphi}' - \varphi_{\text{ext}})^2. \tag{2.14}$$

We can now approximate the potential terms $\hat{\varphi}$ and $\hat{\varphi}^2$ by the corresponding lowest frequency components in this space:

$$\hat{\varphi}' \approx \sin\hat{\varphi}', \quad \hat{\varphi'}^2 \approx 2 - 2\cos\varphi', \tag{2.15}$$

such that the Hamiltonian is approximated by

$$H \approx \frac{4E_c}{N^2}\hat{n}'^2 - E_J\cos(N\hat{\varphi}) + \frac{E_L}{2}\left(N^2(2-\cos\hat{\varphi}') - 2N\sin\hat{\varphi}'\varphi_{\text{ext}}\right) \tag{2.16}$$

This Hamiltonian is still periodic over the interval $\varphi' \in [-\pi, \pi)$, but the phase space is N times larger in the *original* phase coordinate, $\varphi \in [-N\pi, N\pi)$. This approximation naturally assumes that the wavefunction is exponentially small at the boundary, else we can expect energy renormalizations arising from this enforced periodicity.

Note here that, like with the derivative matrices in the phase basis, it is important to approximate the operators $\hat{\varphi}$ and $\hat{\varphi}^2$ separately, since we would otherwise observe non-physical local minima at the boundaries. Consider:

$$(\hat{\varphi}')^2 \approx (\sin\hat{\varphi}')^2 \approx \frac{1-\cos 2\hat{\varphi}'}{2}, \tag{2.17}$$

which has minima at $\varphi' = -\pi, \pi$, unlike $\cos\hat{\varphi}'$. Consequently, the corresponding eigenstates of a Hamiltonian diagonalized with this potential would have states localized at the boundaries, not dissimilar to the issues with the Fock basis in Fig. 2.2.

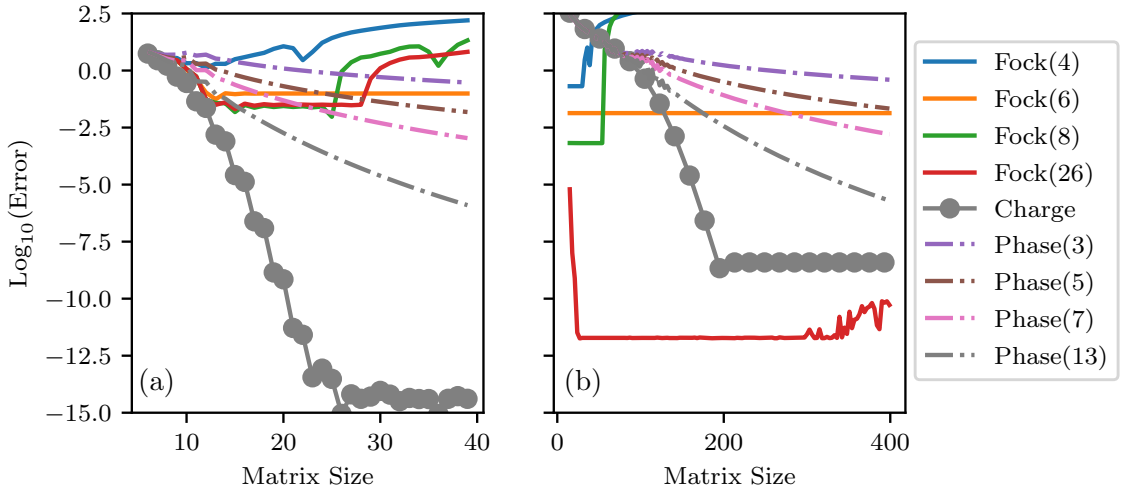

Figure 2.5 Convergence of the first 5 eigenvalues for the (a) transmon and (b) inductively shunted transmon.

2.1.4 Example – Diagonalization Convergence

To compare the different diagonalization methods, I construct the Hamiltonians in each of the different bases and perform numerical diagonalization using NumPy's `eigh` function [48]. I consider two cases: the transmon Hamiltonian in Eq. (2.1) and the inductively shunted transmon in Eq. (2.12) with $\varphi_{\text{ext}} = 0$. As a performance metric, I use the differences in energies of the first five eigenvalues:

$$\text{Error} = \sqrt{\sum_{i=0}^{4} |\lambda_i - \lambda_{i,\text{approx}}|^2}, \tag{2.18}$$

where λ_i is the 'converged' eigenvalue (in 2π GHz).

For the transmon diagonalization in Fig. 2.5(a), these converged eigenvalues are calculated with the charge basis and a total of 200 charge states. For the inductively shunted transmon, this was calculated with the fractional charge basis using $N = 25$ and 2000 basis states, as well as with the phase basis over the interval $[-20\pi, 20\pi)$ with 2000 states. In both of these bases, the eigenvalues differed by less than 10^{-15}, confirming the convergence of the two methods. The convergence properties for the transmon are shown in Fig. 2.5(a), where we consider a transmon with $E_J/E_c = 50$ and $E_c/\hbar = 250$ MHz, and for the inductively shunted transmon in (b), with $E_J/E_c = E_L/E_c = 25$, keeping the same charging energy E_c. Here, the legend Fock(n) refers to the n-th order Taylor expansion of the cosine potential, whereas Phase(n) refers to the phase basis calculated with an n point stencil. Both the

phase and charge bases were restricted to an interval $-[10\pi, 10\pi)$, corresponding to a choice of $N = 10$ in Eq. (2.14). The charge basis is optimal for the transmon without the inductive shunt, reaching numerical precision with as few as 25 charge states – in comparison, the Fock basis performs poorly, even for the 26th order expansion, due to the issues discussed in Sect. 2.1.1. The phase basis converges comparatively slowly to the charge basis, requiring many more states.

In contrast, for the shunted transmon in (b), the Fock basis performs very well for a small number of states. However, we see that Fock(4), and Fock(12) both diverge after a certain Hilbert space size is reached – this is because the leading term in the Taylor expansion is negative for these expressions – for example, $-(\varphi^8/8!)$ for Fock(8). This implies that, for sufficient Hilbert space sizes, the potential will tend to negative infinity, leading to a failure to characterize the lowest energy eigenstates of the system. However, for Fock(6) and Fock(26), the leading order of the potential is positive, resulting in no such divergence of the lowest energy eigenstates. The charge basis converges rather quickly, but not to the exact solution – this is simply because of an insufficient choice of interval size, $N = 10$. A choice of a larger interval would require more states for convergence, although would converge to a solution closer to the true eigenvalues.

2.1.5 Diagonalization summary

As we've seen in this section, accurately extracting the energy spectrum of a qubit is a non-trivial task, and careful consideration must be given to the choice of basis. In compact systems, the charge basis is generally the most efficient choice. On the other hand, for non-compact systems, the Fock basis with a Taylor expansion of the potential can be highly effective, provided the appropriate order in the Taylor series is chosen. The phase basis offers a slow but reliable method of verifying the diagonalization of either of the previous methods.

2.2 Reading out the qubit

In the following section, I provide an overview of key elements of the procedure to readout the transmon qubit which are relevant to Chapters 5 and 6 – for a complete description of readout, I encourage the reader to look at the theoretical [16] and experimental [49] reviews on the subject.

To read out the state of the qubit, we need a means of coupling it to the external environment. This is usually done in the form of a 'readout resonator' – an LC circuit that is capacitively coupled to the transmon, as depicted in Fig. 2.6. This yields a Hamiltonian:

$$\hat{H}_{tr} = 4E_c(\hat{n} - n_g)^2 - E_J \cos\hat{\varphi} + \omega_r \hat{a}^\dagger \hat{a} - ig(\hat{n} - n_g)(\hat{a} - \hat{a}^\dagger), \tag{2.19}$$

where ω_r represents the resonator frequency, and g the transmon-resonator coupling strength. We note that the Fock basis is natural to describe the readout resonator, as it is harmonic – we'll revisit the accurate numerical diagonalization of this Hamiltonian in Chapter 5.

2.2.1 The Jaynes-Cummings Model and the Dispersive Approximation

As we saw in Sect. 2.1.1, while using the Fock basis for the transmon can lead to inaccuracies, it is an excellent tool for gaining intuition on the overall Hamiltonian behaviour. Expanding Eq. (2.19) in the Fock basis, we obtain:

$$\hat{H}_{tr,\text{Fock}} = \omega_r \hat{a}^\dagger \hat{a} + \omega_q \hat{b}^\dagger \hat{b} - \frac{\alpha}{2}\hat{b}^{\dagger 2}\hat{b}^2 - \tilde{g}(\hat{a}^\dagger - \hat{a})(\hat{b}^\dagger - \hat{b}), \tag{2.20}$$

where $\tilde{g} \approx (g/2)(2E_c)^{1/4}$ is the renormalized coupling strength, arising from the zero-point fluctuations of the charge operators. Here, the terms $\hat{a}^\dagger\hat{b}^\dagger$ and $\hat{a}\hat{b}$ are considered 'counter-rotating' since they do not conserve the excitation number in the system. As we performed on Eq. (2.5), we can make an addition Rotating-Wave Approximation and neglect these counter-rotating terms. This yields a 'Jaynes-Cummings'-like Hamiltonian [50]:

$$\hat{H}_{\text{JC}} = \omega_r \hat{a}^\dagger \hat{a} + \omega_q \hat{b}^\dagger \hat{b} - \frac{\alpha}{2}\hat{b}^{\dagger 2}\hat{b}^2 + g(\hat{a}^\dagger b + \hat{a}\hat{b}^\dagger). \tag{2.21}$$

We are now ready to bring this Hamiltonian to an approximately diagonal form. This can be performed with a set of unitary transformations, known as Schrieffer-Wolff transformations [51, 52]. This procedure is well documented in [16] – as such, we will simply quote the final

form of the Hamiltonian,

$$\hat{H}_{\text{Disp}} = \hat{U}\hat{H}_{\text{JC}}\hat{U}^\dagger \approx \hat{a}^\dagger\hat{a}(\tilde{\omega}_r + 2\chi\hat{b}^\dagger\hat{b}) + \tilde{\omega}_q\hat{b}^\dagger\hat{b} - \frac{\alpha}{2}\hat{b}^{\dagger 2}\hat{b}^2 + \frac{K_a}{2}\hat{a}^{\dagger 2}\hat{a}^2 + \frac{K_a'}{2}\hat{a}^{\dagger 2}\hat{a}^2\hat{b}^\dagger\hat{b}, \quad (2.22)$$

with parameters:

$$\begin{aligned} &\tilde{\omega}_r = (\omega_r + \omega_q - \sqrt{\Delta_{qr}^2 + 4g^2})/2, \quad \tilde{\omega}_q = (\omega_r + \omega_q + \sqrt{\Delta_{qr}^2 + 4g^2})/2, \\ &K_a = -E_C(g/\Delta_{qr})^4, \quad \alpha \approx E_c, \quad \chi = -g^2E_c/(\Delta_{qr}(\Delta_{qr} - E_c)), \\ &K_a' = \left(\frac{4E_cK_a}{\Delta_{qr} + E_c(1 - g^2/\Delta_{qr}^2)}\right). \end{aligned} \quad (2.23)$$

This form of the Hamiltonian is valid in the 'dispersive regime', where $|g/\Delta_{qr}| \ll 1$, with $\Delta_{qr} = \omega_q - \omega_r$ the detuning between the qubit and resonator. Similarly to the qubit anharmonicity seen in Eq. (2.6), the values K_a and K_a', referred to as the Kerr-nonlinearities, describe the degree of nonlinearity of the readout resonator conditioned on the state of the qubit.

We additionally see that Eq. (2.22) describes Jaynes-Cummings 'ladders' [53] – the set of resonator states associated with a specific qubit state. If we consider the qubit to be in either the ground or excited states ($\langle\hat{b}^\dagger\hat{b}\rangle = 0$ or 1), we find

$$\hat{H}_g = \tilde{\omega}_r\hat{a}^\dagger\hat{a} + \frac{K_a}{2}\hat{a}^{\dagger 2}\hat{a}^2, \quad \hat{H}_e = (\tilde{\omega}_r + 2\chi)\hat{a}^\dagger\hat{a} + \left(\frac{K_a + K_a'}{2}\right)\hat{a}^{\dagger 2}\hat{a}^2. \quad (2.24)$$

We can therefore consider the set of states $\{|g,0\rangle, |g,1\rangle, ..., |g,n\rangle\}$ and $\{|e,0\rangle, |e,1\rangle, ..., |e,n\rangle\}$ to each be a distinct 'ladder'. Each state $|i,n\rangle$ represents a 'rung' of the ladder, and we climb up or down this ladder by adding or removing a photon to the resonator, for example $\hat{a}^\dagger|g,n\rangle = (n+1)|g,n+1\rangle$. We'll see how this procedure can be generalized to the case of Eq. (2.19) in Chapter 5.

There is a vital feature of this Hamiltonian that allows us to read out the state of the qubit – the contribution $2\chi\hat{a}^\dagger\hat{a}\hat{b}^\dagger\hat{b}$ in Eq. (2.22) means that the frequency of the resonator now depends explicitly on the qubit state. This is one of the key requirements for measuring the state of the qubit, which we'll explore further in Sect. 2.2.4.

2.2.2 Adding the environment

To obtain useful information about the system, the resonator must be connected to the environment. We can model this with a master equation, whereby the resonator loses photons

to the environment at a rate κ [54]:

$$\dot{\hat{\rho}} = -i[\hat{H}_{\mathrm{JC}}, \hat{\rho}] + \kappa \mathcal{D}[\hat{a}]\hat{\rho}, \tag{2.25}$$

where the dissipator $\mathcal{D}[\hat{a}]$ acts on the density matrix,

$$\mathcal{D}[\hat{a}]\hat{\rho} = \hat{a}\hat{\rho}\hat{a}^\dagger - \frac{1}{2}\hat{a}^\dagger\hat{a}\hat{\rho} - \frac{1}{2}\hat{\rho}\hat{a}^\dagger\hat{a}. \tag{2.26}$$

To understand the effect of this dissipator, we must write the loss operator $\hat{a}$ in the eigenbasis of the dispersive Hamiltonian given in Eq. (2.22). Thus, we need to apply the same set of unitary transformations which diagonalized Eq. (2.21), yielding [55]:

$$\hat{U}\hat{a}\hat{U}^\dagger = \left[1 - \frac{g^2}{2\Delta_{qr}^2} - \frac{2g^2 E_c}{\Delta_{qr}^2(\Delta + E_c(1 - 2g^2\Delta_{qr}^2))}\hat{b}^\dagger\hat{b}\right]\hat{a} + \frac{g}{\Delta}\hat{b}. \tag{2.27}$$

For simplicity, we shall ignore the second-order contributions g^2/Δ_{qr}^2 in Eq. (2.27). We arrive at a master equation,

$$\dot{\hat{\rho}} = -i[\hat{H}_{\mathrm{Disp}}, \hat{\rho}] + \kappa\mathcal{D}[\hat{a}]\hat{\rho} + \kappa\left(\frac{g}{\Delta_{qr}}\right)^2 \mathcal{D}[\hat{b}]\hat{\rho}, \tag{2.28}$$

where we have performed yet another rotating wave approximation, splitting up the dissipator $\mathcal{D}[\hat{a} + (g/\Delta_{qr})\hat{b}]\rho \approx \mathcal{D}[\hat{a}]\hat{\rho} + \mathcal{D}[(g/\Delta_{qr})\hat{b}]\rho$. This approximation is valid if the difference of the two frequencies corresponding to the two modes $\hat{a}$ and $\hat{b}$ is sufficiently large –we'll see a case where this is not valid in Chapter 6.

The introduction of a decay channel for the qubit, at rate $\Gamma_\kappa = \kappa(g/\Delta_{qr})^2$, is known as Purcell decay [56]. For the large values of κ needed for fast qubit readout, this decay channel typically results in unacceptable lifetimes of the transmon qubit, on the order of hundreds of nanoseconds – strategies to avoid this problem are discussed in more detail in Chapter 6.

2.2.3 Driving the Resonator

Let us now return to Eq. (2.22), with some additional presumptions. We first assume that the qubit is in either the ground or excited state, such that the qubit is an eigenstate of the operator $\hat{b}^\dagger\hat{b}$ – thus, we can replace $\hat{b}^\dagger\hat{b}$ by its expectation value, $\langle\hat{b}^\dagger\hat{b}\rangle$, and set $\langle\hat{b}^{\dagger 2}\hat{b}^2\rangle = 0$, enforcing the presumption that no higher transmon levels are occupied – indeed, $b^{\dagger 2}b^2|n\rangle = n(n-1)|n\rangle$, and thus this operator only has non-zero eigenvalues for $n \geq 2$. Next, we assume that the population of the resonator is small, such that we can neglect

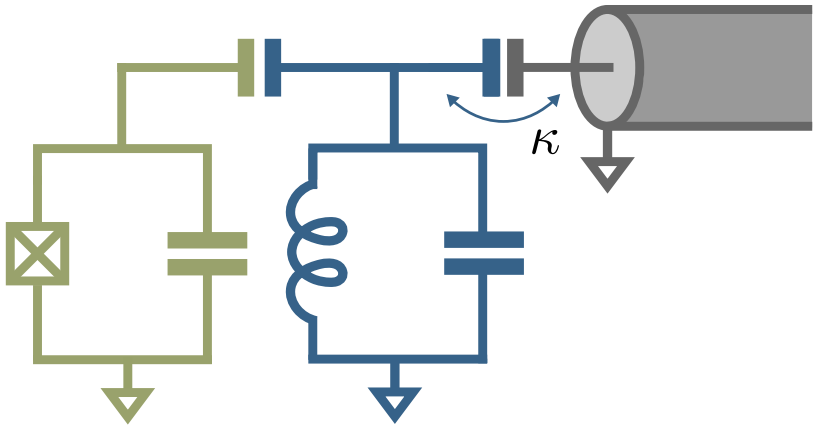

Figure 2.6 Example circuit of a transmon qubit (green) capacitively coupled to a readout resonator (blue). The resonator is coupled to the feedline at rate κ.

the non-linear terms $\hat{a}^{\dagger 2}\hat{a}^2$. We further assume that the decay rate of the qubit derived in Eq. (2.28) is small and can be neglected throughout the readout process. Finally, we include a coherent drive on the resonator at some frequency ω_d and drive amplitude $2\mathcal{E}$. This allows us to write a new, time-dependent master equation:

$$\dot{\hat{\rho}} = -i[(\tilde{\omega}_r + 2\chi\langle\hat{b}^\dagger\hat{b}\rangle)\hat{a}^\dagger\hat{a} + 2i\mathcal{E}\cos(\omega_d t)(\hat{a}^\dagger - \hat{a}), \rho] + \kappa\mathcal{D}[\hat{a}]\hat{\rho}. \tag{2.29}$$

Going into a rotating frame $\hat{U} = \exp(-i\omega_d\hat{a}^\dagger\hat{a}t)$ and, once again, neglecting fast-oscillatory terms, we arrive at the simple equation:

$$\dot{\hat{\rho}}_{\text{RWA}} = -i[(\Delta_{rd} + 2\chi\langle\hat{b}^\dagger\hat{b}\rangle)\hat{a}^\dagger\hat{a} + \mathcal{E}(\hat{a}^\dagger + \hat{a}), \hat{\rho}_{\text{RWA}}] + \kappa\mathcal{D}[\hat{a}]\hat{\rho}_{\text{RWA}}, \tag{2.30}$$

where $\Delta_{rd} = \tilde{\omega}_r - \omega_d$, see [16] for details. We note here that this is simply a driven harmonic oscillator, and as such, we expect a *coherent* response from the resonator. Let us use the notation $|i, \alpha^i(t)\rangle$ to demonstrate a quantum state of this system, where $i = \{g, e\}$ indicates the state of the qubit, and $\alpha^i(t)$ represents the coherent amplitude response of the resonator. It is straightforward to demonstrate that a density matrix initialized in the state $\rho_0 = |j, 0\rangle\langle j, 0|$ will evolve under the action of Eq. (2.30) as

$$|j, 0\rangle\langle j, 0| \to |j, \alpha^j(t)\rangle\langle j, \alpha^j(t)|, \tag{2.31}$$

where the coherent amplitude α^j is governed by the equation

$$\dot{\alpha}^j = -i(\Delta_{rd} + 2\chi\delta_{ej} - i\kappa/2)\alpha^j + \mathcal{E}, \tag{2.32}$$

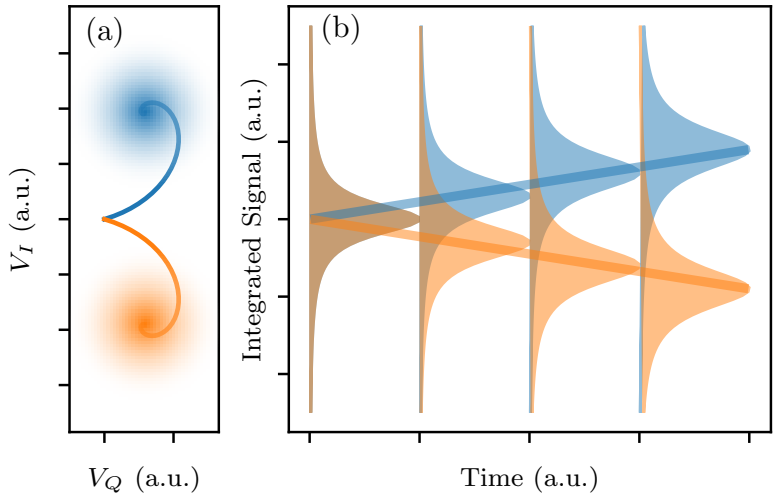

Figure 2.7 (a) Trajectory in phase space of the coherent responses corresponding to the (blue) ground and (orange) excited state. (b) Integrated signal as a function of time. For longer measurements, the responses become increasingly resolved. The overlap between the two distributions yields the readout infidelity.

with the kronecker delta $\delta_{ej} = 1$ if $j = e$ and 0 otherwise. Here, it is clear that the shift 2χ imparted on the resonator by the qubit makes the resonator's coherent response $\alpha^{g,e}$ qubit-state dependent.

2.2.4 Putting it all together

A full understanding of the readout process requires knowledge of all the components in the measurement chain, which include signal generators, amplifiers, attenuators, and IQ mixers, components which are outside the scope of this thesis – Ref. [16] provides an overview of these components, with Refs. [49, 57] going into finer detail. Here, I shall simply jump to the main result – a heterodyne measurement is performed which allows for simultaneous measurement of the two field quadratures – the 'in-phase' $V_I(t)$ and 'in-quadrature' $V_Q(t)$ components. Each of these components contains a contribution from the *internal* field of the resonator $\langle \hat{a}(t) \rangle$ [58]:

$$V_Q(t) = (\sqrt{2\kappa\eta}\,\mathrm{Re}\,\langle \hat{a}(t) \rangle + n_Q(t)), \quad V_I(t) = (\sqrt{2\kappa\eta}\,\mathrm{Im}\,\langle \hat{a}(t) \rangle + n_I(t)), \tag{2.33}$$

where $n_Q(t), n_I(t)$ are Gaussian white noise terms and η is the measurement efficiency, a measure of information loss by the measurement chain [59]. With the assumptions made in Sect. 2.2.3, we see a coherent, qubit-state dependent response from the resonator such

that $\langle \hat{a} \rangle(t) = \alpha^{g,e}(t)$, see the representative example of the coherent responses of the ground and excited states given in Fig. 2.7(a). Consequently, the measurement signal in Eq. (2.33) depends on the qubit state, and by integrating this signal we can discriminate between the two states, see Fig. 2.7(b). Bultink et. al. in Ref. [58] derived that for this measurement the Signal-to-Noise ratio (SNR), used as a measure to discriminate the states, is of the form

$$\mathrm{SNR}(t) = 2\kappa\eta \int_0^t |\alpha^g(t') - \alpha^e(t')|^2 dt', \tag{2.34}$$

where this definition is adapted to be in agreement with the publication in Chapter 6. Assuming that the resonator response is Gaussian, the SNR provides information on the upper bound of the readout fidelity $\mathcal{F}$:

$$\mathcal{F} \leq \frac{1}{2}\left(1 + \mathrm{erf}(\sqrt{\mathrm{SNR}/8})\right). \tag{2.35}$$

We consequently wish to maximize the SNR, achieved by maximizing the distance between the two steady-state resonator responses $\{\alpha^g_{\mathrm{ss}}, \alpha^e_{\mathrm{ss}}\}$, which are of the form

$$\alpha^{g,e}_{ss} = \frac{-i\mathcal{E}}{\Delta_{rd} + 2\chi\delta_{ej} - i\kappa/2}. \tag{2.36}$$

For a fixed drive amplitude $\mathcal{E} = \mathcal{E}_0$, the distance between these responses is maximized by choosing $\omega_d = \tilde{\omega}_r + \chi$, the average frequency of the two qubit-dependent resonator responses.

An obvious way of improving the SNR appears to be to simply increase the drive amplitude $\mathcal{E}$. Unfortunately, the dispersive approximation in Eq. (2.22) quickly breaks down at higher photon numbers, leading to strong resonances which can rapidly degrade the purity of the qubit state and reduce the performance of the readout – we'll see this much greater detail in Chapter 5. A common metric, referred to as the 'critical photon number' $n_{\mathrm{crit}} = \Delta^2_{qr}/(2\tilde{g})^2$ [15] defines a photon population before the dispersive theory is expected to break down. Whilst a useful approximation, experimental results often see a breakdown in the readout process well before this photon number is reached [30, 60, 61, 62]. I contributed to some recent theoretical findings which showed that the readout can fail at very low average photon numbers if the qubit frequency is *below* the resonator frequency [63]. To help understand this effect, a new definition of the critical photon number is provided. This topic is currently under active investigation.

2.2.5 Readout Summary

The journey from Eq. (2.19) to the formula for the coherent state responses in Eq. (2.32) is littered with presumptions, approximations, and transformations. Whilst this derivation is pedagogical and provides insight into how 'ideal' readout occurs, it does not well capture effects seen experimentally at high photon numbers [3, 30, 60]. Indeed, Eq. (2.30) predicts that arbitrarily strong drives can be used without repercussion, indicating that the model is insufficient. Thus, to capture this breakdown of 'quantum non-demolition readout' at high photon numbers, a more sophisticated model is needed. Chapter 5, we'll avoid many of the approximations made above and instead focus on the evolution of Eq. (2.19) under a strong resonator drive, and in Chapter 6, we'll see simulations of a more realistic setup, which includes an additional mode to reduce the qubit decay we derived in Eq. (2.28).

2.3 Numerical Propagation

In Sect. 2.2, many approximations were made to yield a useful form of the Hamiltonian. If we wish to understand these systems to a greater degree of accuracy, we must resort to modelling these systems numerically, as we saw in the diagonalization of the transmon in Sect. 2.1. Given that in Chapters 3 and 4 I introduce a novel numerical integration method, it is appropriate to introduce the basics of how standard numerical integrators propagate states.

The very simplest integrator uses what's known as the 'Euler method' [64], and takes finite step sizes δt:

$$i\frac{d}{dt}|\psi\rangle = \hat{H}(t)|\psi\rangle \longrightarrow |\psi(t+\delta t)\rangle = (1 - i\hat{H}\delta t)|\psi(t)\rangle. \tag{2.37}$$

Such an integrator has an error on the order of $(\delta t)^2$ and does not preserve the norm of the wavefunction so that $|\langle\psi(N\delta t)|\psi(N\delta t)|^2 \neq 1$. More advanced integrators exist, such as Runga-Kutta [65, 66], Adams [67], and bdf [68], which integrate a time step δt in several stages, either explicitly or implicitly. Such integrators generally have errors on the order of $(\delta t)^4$ or smaller, allowing for larger time steps, faster integration, and more accurate simulations.

A limitation of such numerical integrators is the assumption that the time-dependence of the Hamiltonian $\hat{H}(t)$ is *slow* with respect to the time step δt. Given that time-dependent drives generally exhibit *fast-oscillating* behaviour, this imposes significant constraints on

feasible step sizes. In Chapter 3, I will propose an expansion that allows for step sizes that are independent of the drive frequencies, therefore enabling significantly larger steps.

2.3.1 Master Equation Propagation

Evolution of the master equation works in much the same way as the propagation of a ket vector, except we now propagate a *density matrix*:

$$\dot{\hat{\rho}} = \mathcal{L}(t)\hat{\rho}, \tag{2.38}$$

where $\mathcal{L}$ is called the Lindbladian *superoperator*, and describes the density matrix evolution:

$$\mathcal{L}(t)\hat{\rho} = -i[\hat{H}(t), \hat{\rho}] + \sum_i \kappa_i \mathcal{D}[\hat{L}_i]\hat{\rho}. \tag{2.39}$$

A key feature of the Lindbladian is that it preserves the trace of the density matrix, such that $\mathrm{Tr}\,[\mathcal{L}\hat{\rho}] = 0$. Indeed, even the Euler steps described above will preserve the trace:

$$\begin{aligned} \hat{\rho}(t+\delta t) &\approx \hat{\rho}(t) + \delta t \dot{\hat{\rho}}, \\ \mathrm{Tr}[\hat{\rho}(t+\delta t)] &= 1 + \delta t \cdot 0 = 1. \end{aligned} \tag{2.40}$$

Nevertheless, the error associated with this method is still on the order of δt^2 – in other words, the preservation of the trace does not guarantee anything about the accuracy of the solver.

Roots of unity solver

Because it will be useful later [55], we consider here an interesting approximation of the evolution of the density matrix. If we assume that the Lindbladian is time-independent, we can write the exact solution to the evolution:

$$\dot{\hat{\rho}} = \mathcal{L} \rightarrow \hat{\rho}(t) = e^{\mathcal{L}t}\hat{\rho}(0). \tag{2.41}$$

To see how this could be implemented numerically, we can start by considering the Taylor expansion of the exponential

$$e^{\mathcal{L}t}\hat{\rho} = \sum_{n=0}^{\infty} \frac{(\mathcal{L}t)^n}{n!}\hat{\rho}. \tag{2.42}$$

A potential strategy for solving this equation without explicitly calculating powers of $\mathcal{L}$ is to store two matrices – a 'cumulative' density matrix $\hat{\rho}$, and an additional matrix $\hat{\rho}_{i+1} = \mathcal{L}\hat{\rho}_i$

which has successive powers of $\mathcal{L}$ applied on it. More explicitly, starting with an initial matrix $\hat{\rho}_0$, we find

$$\begin{aligned}\hat{\rho} &= e^{\mathcal{L}t}\hat{\rho}_0 \\ &= \hat{\rho}_0 + \underbrace{\mathcal{L}\hat{\rho}_0}_{\hat{\rho}_1} + \frac{1}{2}\underbrace{\mathcal{L}\hat{\rho}_1}_{\hat{\rho}_2} + \frac{1}{6}\underbrace{\mathcal{L}\hat{\rho}_2}_{\hat{\rho}_3} + ...,\end{aligned} \tag{2.43}$$

However, if the norm of $\mathcal{L}$ is large, the norm of the matrices $\hat{\rho}_i$ can diverge and be prone to numerical instability. Instead, let us consider the roots of the $n-$th degree polynomial approximating the Taylor expansion of the exponential function:

$$\begin{aligned}p_n(x) &= \sum_{n=0}^{\infty} \frac{1}{n!}x^n = 1 + x + \frac{x^2}{2} + ... \\ &= \frac{1}{n!}\prod_{i=1}^{n}(x - c_i) = \prod_{i=1}^{n}(1 - x_i/z_i).\end{aligned} \tag{2.44}$$

We can now use this form to approximate the action of the matrix exponential:

$$\hat{\rho}(t, t + \delta t) = \prod_{i=1}^{n}(I - \mathcal{L}_i \delta t/z_i)\hat{\rho}(t). \tag{2.45}$$

This result is particularly useful if the density matrix is extremely large and hard to store in memory, as we only ever have to store one copy – such a method is used in the publication in Chapter 5.

Chapter 3

Dysolve

"I feel the need... the need for speed."

– Tom Cruise, Top Gun

In this chapter, I provide more context to how the Dyson expansion can be used and implemented as a numerical solver. The essential details of the algorithm are found in the publication, in section Sect. 3.6 – consequently, the following sections are to provide more insight into the algorithm, including a comparison against an alternative method, and provide concrete details on how the algorithm can be implemented numerically.

3.1 Motivation

The vast majority of quantum operations – whether single and two-qubit gates, readout schemes or qubit reset – involve one or more time-dependent oscillatory drives. If such an operation does not involve a loss channel, we can describe this evolution with a time-dependent Hamiltonian of the form

$$\hat{H}(t) = \hat{H}_0 + \hat{V}(t), \tag{3.1}$$

where we have a time-dependent perturbation $\hat{V}(t)$,

$$\hat{V}(t) = \sum_i \hat{X}_i \left[\mathcal{E}_{s,i}(t)\sin(\omega_{d,i}t), +\mathcal{E}_{c,i}(t)\cos(\omega_{d,i}t)\right], \tag{3.2}$$

where $\hat{H}_0$ is a diagonal Hamiltonian, $\hat{X}_i$ is a drive operator, and $\mathcal{E}_{s,i}(t), \mathcal{E}_{c,i}(t)$ are drive envelopes which vary slowly with respect to the corresponding drive frequencies $\omega_{d,i}$. We are ultimately interested in the corresponding unitary operator which evolves a system from time t_0 to t – this obeys the differential equation

$$\dot{\hat{U}} = -i\hat{H}(t)\hat{U}, \tag{3.3}$$

which has the corresponding formal solution

$$\hat{U}(t, t_0) = \mathcal{T} \exp\left(-i \int_{t_0}^{t} \hat{H}(t')dt'\right), \tag{3.4}$$

where $\mathcal{T}$ is the time-ordering operator.

As detailed in Sect. 2.3, integrating equation Eq. (3.3) using a standard numerical solver can be challenging since the chosen step size Δt has to be significantly smaller than the oscillation period of the drive, $2\pi/\omega_d$. This requires a significant number of time steps, thus requiring many matrix multiplications and a slow evaluation of the target unitary operator in Eq. (3.4). Consequently, a solver capable of integrating these fast-time dynamics natively would be invaluable in accelerating the simulation of such quantum operations.

3.1.1 Magnus Expansion – A potential solver?

The Magnus expansion [69] is a powerful method for approximating the evolution of quantum systems. Here, the *average* Hamiltonian over the relevant timestep δt is calculated through a sum of commutators:

$$\bar{H} = \sum_{k=1}^{\infty} \bar{H}_k, \quad \bar{H}_1 = \int_{t_0}^{t} \hat{H}(t_1)dt, \quad \bar{H}_2 = \frac{1}{2}\int_{t_0}^{t}\int_{t_0}^{t_1} [\hat{H}(t_1), \hat{H}(t_2)]dt_1 dt_2, \tag{3.5}$$

with higher order contributions $\hat{H}_k$ evaluated from a corresponding set of nested commutators.

This method has two drawbacks: there are strict conditions regarding the norm of $\hat{H}(t)$ for the truncation of the series to be valid, and this method requires the exponentiation of the average Hamiltonian to approximate the unitary $\hat{U}(t, t_0) \approx \exp(-i\bar{H}\delta t)$. Matrix exponentials can be very costly to calculate or approximate numerically, which makes this solver less attractive. Furthermore, in the context of optimal control, which frequently involves gradient ascent [70], it becomes necessary to calculate derivatives of the unitary

operator with respect to the drive amplitudes. However, the Magnus expansion does not automatically provide these derivatives.

Nevertheless, a Magnus solver can be very effective in the correct context, such as understanding and correcting leakage processes [71]. Moreover, in this thesis the Magnus expansion was used to propagate a large density matrix [2] (see Chapter 5), in conjunction with the roots-of-unity solver introduced in Chapter 2. A Magnus and Dyson solver, which was inspired by my work on the Dysolve algorithm detailed in this chapter, was implemented by a team at IBM in the Qiskit Dynamics package [72].

3.2 Dyson Series

The Dyson series is a standard tool to help understand the effect of a perturbation $\hat{V}(t)$ on an unperturbed Hamiltonian $\hat{H}_0$. The original series as posed by Dyson is given in the form [73]

$$\hat{U}(t) = \hat{I} + \sum_{n=1}^{\infty} \hat{U}_n(t, t_0), \tag{3.6}$$

where

$$\hat{U}_n(t, t_0) = \frac{(-i)^n}{n!} \int_{t_0}^{t} dt_1 \int_{t_0}^{t} dt_2 ... \int_{t_0}^{t} dt_n \mathcal{T} \hat{V}_I(t_1) \hat{V}_I(t_2) ... \hat{V}_I(t_n). \tag{3.7}$$

Here, the problem is posed in the *interaction picture*, such that $\hat{V}_I(t) = e^{i\hat{H}_0 t} \hat{V}(t) e^{-i\hat{H}_0 t}$. The error for such a series will be on the order

$$U_{\text{err}} = \hat{U}(t) - \sum_{n=1}^{m} \hat{U}_m(t, t_0) = \frac{(\delta t)^{m+1}}{(m+1)!} \mathcal{O}(\hat{V}^{m+1}), \tag{3.8}$$

which will be negligible for sufficiently large m, and systems with weak perturbations (i.e. drives). However, the form of Eq. (3.7) is not instructive on how to calculate the terms $\hat{U}_n(t, t_0)$ in general, given the complexity of time-ordering and nested integrals. In Ref. [1], I demonstrate how the unitary in Eq. (3.7) corresponding to the Hamiltonian in Eq. (3.2) can be decomposed efficiently. In this chapter, I provide additional context into the decomposition, which can be understood more simply from a time-independent Hamiltonian:

$$\hat{H} = \hat{H}_0 + \Omega \hat{X}, \tag{3.9}$$

where Ω is some small constant, $\hat{H}_0$ is diagonal and $\hat{X}$ is a Hermitian operator:

$$\hat{H}_0 = \sum_k \lambda_k |k\rangle\langle k|, \quad \hat{X} = \sum_{k^{(0)},k^{(1)}} \langle k^{(1)}|\hat{X}|k^{(0)}\rangle |k^{(1)}\rangle\langle k^{(0)}|. \tag{3.10}$$

Here, the superscripts $k^{(0)}$, $k^{(1)}$ are simply used to differentiate different summation indices. Replicating the derivations in Ref. [1] with this simpler Hamiltonian, we arrive at the form of the unitary

$$\hat{U}(t, t+\delta t) = e^{-i\hat{H}_0\delta t} + \sum_{n=1}^{\infty} \Omega^n \hat{S}^n(\delta t), \tag{3.11}$$

where the Dyson series operators $\hat{S}^{(n)}(\delta t)$ are defined as

$$\begin{aligned}&\hat{S}^{(n)}(\delta t) = \\ &(-i\delta t)^n \sum_{\boldsymbol{k}_n} f\left(\boldsymbol{\lambda}_n(\boldsymbol{k}_n)\delta t\right) \langle k^{(n)}|\hat{X}|k^{(n-1)}\rangle\langle k^{(n-1)}|\hat{X}\cdots|k^{(1)}\rangle\langle k^{(1)}|\hat{X}|k^{(0)}\rangle|k^{(n)}\rangle\langle k^{(0)}|.\end{aligned} \tag{3.12}$$

Here, each $\boldsymbol{k}_n = (k^{(0)}, k^{(1)}, \cdots, k^{(n)})$ is a set of indices which specify a set of $(n+1)$ eigenstates $\{|k^{(m)}\rangle\}$ of $\hat{H}_0$ with corresponding eigenvalues $\boldsymbol{\lambda}_n(\boldsymbol{k}_n)$, and we sum over all possible $\boldsymbol{k}_n$. These eigenvalues are written in vector form:

$$\boldsymbol{\lambda}_n(\boldsymbol{k}_n) \equiv \left(\lambda_{k^{(0)}}, \cdots, \lambda_{k^{(n)}}\right). \tag{3.13}$$

The sum over $\boldsymbol{k}_n$ implies a summation over all sets of eigenstates which the operator $\hat{X}$ couples in Eq. (3.12). The functions $f\left(\boldsymbol{\lambda}_n(\boldsymbol{k}_n)\delta t\right)$ are divided difference functions [74, 75], defined by Equations (13) and (17) in the publication, and discussed in more detail below.

3.3 Calculating the divided difference functions

As defined in Ref. [1] and replicated here for clarity, the the n-th order weighting functions $f(\boldsymbol{\lambda}_n)$ can be obtained recursively:

$$f(\boldsymbol{\lambda}_n) = i\frac{f(\boldsymbol{g}(\boldsymbol{\lambda}_n)) - f(\boldsymbol{g}^2(\boldsymbol{\lambda}_n) \cup \boldsymbol{\lambda}_n[n])}{\boldsymbol{\lambda}_n[n-1] - \boldsymbol{\lambda}_n[n]}, \tag{3.14}$$

where $\boldsymbol{g}(\boldsymbol{v}_n)$ returns $\boldsymbol{v}_n$ without its last element, $\boldsymbol{g}^2(\boldsymbol{v}_n) = \boldsymbol{g}(\boldsymbol{g}(\boldsymbol{v}_n))$, and the notation $\cup$ indicates appending an additional element to a vector such that

$$\boldsymbol{\lambda}_n = \boldsymbol{g}(\boldsymbol{\lambda}_n) \cup \boldsymbol{\lambda}_n[n], \tag{3.15}$$

where $\boldsymbol{\lambda}_n[n]$ refers to the n-th element of the vector $\boldsymbol{\lambda}_n$. In the case of degenerate eigenvalues (for example, $\boldsymbol{\lambda}_n[n] = \boldsymbol{\lambda}_n[n-1]$), Eq. (3.14) can be defined in terms of limits, but this is challenging to implement numerically. Take for example the following divided difference function, which from the above definition would suffer from a division-by-zero error:

$$f(\lambda_0, \lambda_1, \lambda_1, \lambda_1) = i\frac{f(\lambda_0, \lambda_1, \lambda_1) - f(\lambda_0, \lambda_1, \lambda_1)}{\lambda_1 - \lambda_1}. \tag{3.16}$$

This can be defined in a limit by replacing the last eigenvalue by λ' and taking $\lambda_1' \to \lambda_1$. However, expanding the equation leads to two additional division by zero occurrences:

$$\begin{aligned} f(\lambda_0, \lambda_1, \lambda_1, \lambda_1) &= \lim_{\lambda_1' \to \lambda} i\frac{f(\lambda_0, \lambda_1, \lambda_1') - f(\lambda_0, \lambda_1, \lambda_1)}{\lambda_1' - \lambda_1} \\ &= \lim_{\lambda_1' \to \lambda} i\frac{\left(\frac{f(\lambda_0,\lambda_1)-f(\lambda_0,\lambda_1)}{\lambda_1-\lambda_1}\right) - \left(\frac{f(\lambda_0,\lambda_1')-f(\lambda_0,\lambda_1')}{\lambda_1'-\lambda_1'}\right)}{\lambda_1' - \lambda_1}. \end{aligned} \tag{3.17}$$

There is a simple way of circumventing these issues. Let us consider the same function but with arbitrary eigenvalues : $f(\lambda_0, \lambda_1, \lambda_2, \lambda_3)$. First, noting from Ref. [1] that we can permute the eigenvalues, we assume $\lambda_0 \leq \lambda_1 \leq \lambda_2 \leq \lambda_3$. Then, we calculate all the *first* order divided difference functions based on sequential pairs of the eigenvalues: $f(\lambda_0, \lambda_1)$, $f(\lambda_1, \lambda_2)$ and $f(\lambda_2, \lambda_3)$. Here, for any pair of degeneracies, the analytical solution $f(\lambda_j, \lambda_j) = \exp(-i\lambda_j)$ can be used.

The next order of divided differences is calculated using the previous functions, but done in the same *sequence* the original divided difference functions were calculated in:

$$f(\lambda_0, \lambda_1, \lambda_2) = i\frac{f(\lambda_0, \lambda_1) - f(\lambda_1, \lambda_2)}{\lambda_0 - \lambda_2}, \quad f(\lambda_1, \lambda_2, \lambda_3) = i\frac{f(\lambda_1, \lambda_2) - f(\lambda_2, \lambda_3)}{\lambda_1 - \lambda_3}. \tag{3.18}$$

We see here that if there is a division by zero in the first function, this requires $\lambda_0 = \lambda_2$, but $\lambda_0 \leq \lambda_1 \leq \lambda_2$ and so $\lambda_1 = \lambda_2$. This means we can once again resort to the third-order analytical expression $f(\lambda_0, \lambda_0, \lambda_0) = \exp(-i\lambda_0)/2$.

We return to the example in Eq. (3.16) and use the above methodology. Assuming $\lambda_0 < \lambda_1$ this would be calculated in a numerically stable fashion as

$$f(\lambda_0, \lambda_1, \lambda_1, \lambda_1) = i\frac{\left(f(\lambda_0, \lambda_1) - e^{-i\lambda_1}\right)/(\lambda_0 - \lambda_1) - e^{-i\lambda_1}/2}{\lambda_0 - \lambda_1}, \tag{3.19}$$

which avoids any more occurrences of division by zero.

3.4 Calculation of the Dyson Series Operators

In this section I provide additional insight into how the Dyson series operators in Eq. (3.12) can be calculated, and how degeneracies can be avoided. For concreteness, we begin by considering the first order operator $\hat{S}^{(1)}(\delta t)$. This term can be simply calculated by a Hadmard product:

$$\begin{aligned}\hat{S}^{(1)}(\delta t) &= (-i\delta t)\sum_{k_1,k_0} f\left(\lambda_{k_1}\delta t, \lambda_{k_0}\delta t\right)\langle k_1|\hat{X}|k_0\rangle|k_1\rangle\langle k_0| \\ &= (-i\delta t)\,\mathcal{F}_1 \odot \hat{X}.\end{aligned} \tag{3.20}$$

where the Hadamard product is defined by:

$$\hat{A} = \sum_{ij} a_{ij}|i\rangle\langle j|,\ \ \hat{B} = \sum_{ij} b_{ij}|i\rangle\langle j|,\ \ \hat{A}\odot\hat{B} = \sum_{ij} a_{ij}b_{ij}|i\rangle\langle j|. \tag{3.21}$$

Here we have defined a matrix $\mathcal{F}_1$, whose elements are the divided difference functions:

$$\mathcal{F}_1 = \sum_{k_0,k_1} |k_1\rangle\langle k_0| f(\lambda_{k_1}\delta t, \lambda_{k_0}\delta t). \tag{3.22}$$

At second order, we must consider the sum over three indices:

$$\begin{aligned}\hat{S}^{(2)}(\delta t) &= (-i\delta t)^2 \sum_{k_2,k_1,k_0} f\left(\lambda_{k_2}\delta t, \lambda_{k_1}\delta t, \lambda_{k_0}\delta t\right)\langle k_2|\hat{X}|k_1\rangle\langle k_1|\hat{X}|k_0\rangle|k_2\rangle\langle k_0| \\ \hat{S}^{(2)}_{k_2,k_0}(\delta t) &= (-i\delta t)^2 \sum_{k_1} \mathcal{F}^{(2)}_{k_2,k_1,k_0}\mathcal{G}^{(2)}_{k_2,k_1,k_0} \\ &= (-i\delta t)^2 \sum_{k_1} \left(\mathcal{F}^{(2)}\odot\mathcal{G}^{(2)}\right)_{k_2,k_1,k_0}.\end{aligned} \tag{3.23}$$

Similarly to Eq. (3.20), this is still represented by a Hadamard product, albeit with *tensors*. These are defined element-wise:

$$\mathcal{G}^{(2)}_{k_2,k_1,k_0} = \langle k_2|\hat{X}|k_1\rangle\langle k_1|\hat{X}|k_0\rangle,\ \ \mathcal{F}^{(2)}_{k_2,k_1,k_0} = f(\lambda_{k_2}\delta t, \lambda_{k_1}\delta t, \lambda_{k_0}\delta t), \tag{3.24}$$

which can be extended to arbitrary order:

$$\hat{S}^{(n)}_{k_n,k_0}(\delta t) = (-i\delta t)^n \sum_{\mathbf{k}_n} \left(\mathcal{F}^{(n)}\odot\mathcal{G}^{(n)}\right)_{\mathbf{k}_n}, \tag{3.25}$$

and where the elements of the tensors defined as

$$\mathcal{G}^{(n)}_{\boldsymbol{k}_n} = \langle k^{(n)}|\hat{X}|k^{(n-1)}\rangle\langle k^{(n-1)}|\hat{X}\cdots|k^{(1)}\rangle\langle k^{(1)}|\hat{X}|k^{(0)}\rangle, \quad F^{(n)}_{\mathbf{k}_n} = f\left(\boldsymbol{\lambda}_n(\boldsymbol{k}_n)\delta t\right). \tag{3.26}$$

These definitions allow for calculating the Dyson operators only with tensor operations, allowing for parallelized computation. The tensor $\mathcal{G}^{(n)}_{\mathbf{k}_n}$ is easy to calculate from outer products of $\hat{X}$, but care must be taken in calculating all of the elements of $\mathcal{F}^{(n)}$ numerically since degeneracies in the eigenvalues can cause divergences. Proceeding, I shall set $\delta t = 1$ for notational simplicity.

3.5 Efficient Calculation of Dyson Operators

If the drive operator $\hat{X}$ is sparse, constructing and contracting the tensors $\mathcal{G}^{(n)}$ and $\mathcal{F}^{(n)}$ in sparse format is generally efficient and straightforward. However, if $\hat{X}$ is *dense*, constructing the tensors $\mathcal{G}^{(n)}$ and $\mathcal{F}^{(n)}$ can become highly memory intensive for a large Hilbert space size N since the number of coefficients of these tensors scales as N^{n+1}.

The Dyson series generally converges best in the *eigenbasis* of the considered problem, since this reduces the number of perturbative terms. The drive operators are generally dense after an arbitrary frame transformation – consequently, we seek a more efficient way of calculating the Dyson operators in this case.

3.5.1 Second order evaluation

Let us start with a pedagogical case, where $\hat{H}_0 = 0$. Here, we trivially find $f(\boldsymbol{\lambda}_n) = \frac{1}{n!}$, and the elements of Dyson operator at $n-$th order simply become

$$\hat{S}^{(n)}_{k_n,k_0}(1) = (-i)^n \frac{1}{n!}\sum_{\mathbf{k}_n}\mathcal{G}^{(n)}_{\mathbf{k}_n}. \tag{3.27}$$

However, this problem is equivalent to exponentiating the Hamiltonian $\hat{H} = \Omega\hat{X}$, and thus the Dyson operator in Eq. (3.27) must simply be proportional to the n-th power of $\hat{X}$:

$$\hat{S}^{(n)}(1) = (-i)^n \frac{1}{n!}\hat{X}^n. \tag{3.28}$$

For this case, calculating the tensor $\mathcal{G}^{(n)}$ was unnecessary – this term can be calculated by matrix multiplication only. Inspired by this, we can ask ourselves whether these higher-order

terms, defined in Eq. (3.25) as a contraction of two tensors, can also be calculated solely through matrix multiplications.

Let us suppose that the second-order divided-difference function can be approximated as a *sum* of the product of two different functions, which each depends on a different pair of eigenvalues:

$$f(\lambda_1, \lambda_2, \lambda_3) \approx \sum_i g_i(\lambda_1, \lambda_2) h_i(\lambda_2, \lambda_3). \tag{3.29}$$

Inserting this term into the definition of $\hat{S}^{(2)}(1)$, we find

$$\begin{aligned}
\hat{S}^{(2)}(1) &\approx (-i)^2 \sum_i \sum_{k_2,k_1,k_0} h_i(\lambda_{k_1}, \lambda_{k_2}) g_i(\lambda_{k_1}, \lambda_{k_0}) \langle k_2|\hat{X}|k_1\rangle\langle k_1|\hat{X}|k_0\rangle|k_2\rangle\langle k_0| \\
&\approx (-i)^2 \sum_i \sum_{k_2,k_1,k_0} \left(h_i(\lambda_{k_1}, \lambda_{k_2}) \langle k_2|\hat{X}|k_1\rangle|k_2\rangle\langle k_1|\right) \left(g_i(\lambda_{k_1}, \lambda_{k_0}) \langle k_1|\hat{X}|k_0\rangle|k_1\rangle\langle k_0|\right) \\
&\approx (-i)^2 \sum_i \left(H_i \odot \hat{X}\right) \left(G_i \odot \hat{X}\right),
\end{aligned} \tag{3.30}$$

where

$$G_i = \sum_{jk} |j\rangle\langle k| g_i(\lambda_j, \lambda_k), \quad H_i = \sum_{jk} |j\rangle\langle k| h_i(\lambda_j, \lambda_k). \tag{3.31}$$

Such an expansion would entirely circumvent the need to construct tensors, relying solely on Hadamard products and matrix multiplications of the size of the original Hilbert space only. This would thus dramatically reduce the memory and potential computational time required to calculate these operators.

Calculating the decomposition described in Eq. (3.29) is no trivial task – to do so rigorously, we would have to define how such a function converges, bases functions and more. Before detailing this method further (in Sect. 3.5.3), we can provide more motivation by an analogy – the Singular Value Decomposition (SVD).

3.5.2 Singular Value Decomposition

To proceed, we must make use of one of the identities used in Ref. [1], namely:

$$f(\lambda_0, \lambda_1, \lambda_2) = e^{i\lambda_0} f(0, \Delta_{10}, \Delta_{20}), \tag{3.32}$$

such that the modified divided difference function $f(0, \Delta_{10}, \Delta_{20})$ depends only on two variables which are eigenvalue *differences*, $\Delta_{ij} = \lambda_i - \lambda_j$. This motivates a simpler form of the decomposition in Eq. (3.29), such that the functions $g_i,\ h_i$ are dependent on one variable

only:

$$f(0, \Delta_{10}, \Delta_{20}) \approx \sum_i g_i(\Delta_{10}) h_i(\Delta_{21}), \tag{3.33}$$

noting that $\Delta_{20} = \Delta_{10} + \Delta_{21}$. To see how these functions $g_i(\Delta_{21})$ could be calculated, we will first consider the function $f(0, \Delta_{10}, \Delta_{21} + \Delta_{10})$ for a large range of Δ_{21}, Δ_{10}, which is displayed in Fig. 3.1 (a). This data is stored in a matrix F, where the two axes are given as a function of Δ_{10} and Δ_{21}. Then, we consider the SVD decomposition of F,

$$F = U\Sigma V^\dagger = \sum_i s_i u_i v_i^\dagger, \tag{3.34}$$

where s_i is the $i-$th singular value of F and Σ is the diagonal matrix of singular values. Given that the sampling interval is small, this SVD decomposition immediately yields discretized approximations to the desired functions, $u_i \to g_i(\Delta_{21})$, $h_i \to g_i(\Delta_{32})$.

In Fig. 3.1(b-e) I plot the functions $g_i(\Delta_{10})h_i(\Delta_{21})$ corresponding to the first four singular values, with the individual functions $g_i(\Delta_{10})$ and $h_i(\Delta_{21})$ plotted on the sides of each subplot. The first function well captures the features in (a) for small values of Δ_{ij}, as well as the behaviour along the lines $\Delta_{10} = 0$ and $\Delta_{21} = 0$. The decomposition, however, struggles to capture the diagonal $\Delta_{21} = -\Delta_{10}$ behaviour corresponding to $\lambda_0 = \lambda_2$. As the order in the SVD decomposition increases, we see an iterative correction along this diagonal – this is more easily demonstrated in Fig. 3.2, where I plot the sum of the functions corresponding to the first (a) 1, (b) 3, (c) 5 and (d) 7 singular values. From this data, it is clear that the usefulness of this method will be primarily dictated by the *order* of this 'SVD-like' decomposition, and the maximum eigenvalue difference Δ_{ij}, recalling that these differences will scale with the time step.

3.5.3 GeneRalized Order Nth degree Decomposition (GROND)

Whilst the SVD decomposition in Sect. 3.5.1 gave a strong intuition on what the functions $h_i(\Delta_{jk})$ would look like, it is not necessarily clear how these functions could be calculated. Moreover, we wish to extend this method beyond the second order, i.e. seek expansions for the third order of the form

$$f(0, \Delta_{10}, \Delta_{20}, \Delta_{30}) \approx \sum_i g_i(\Delta_{10}) h_i(\Delta_{21}) k_i(\Delta_{32}), \tag{3.35}$$

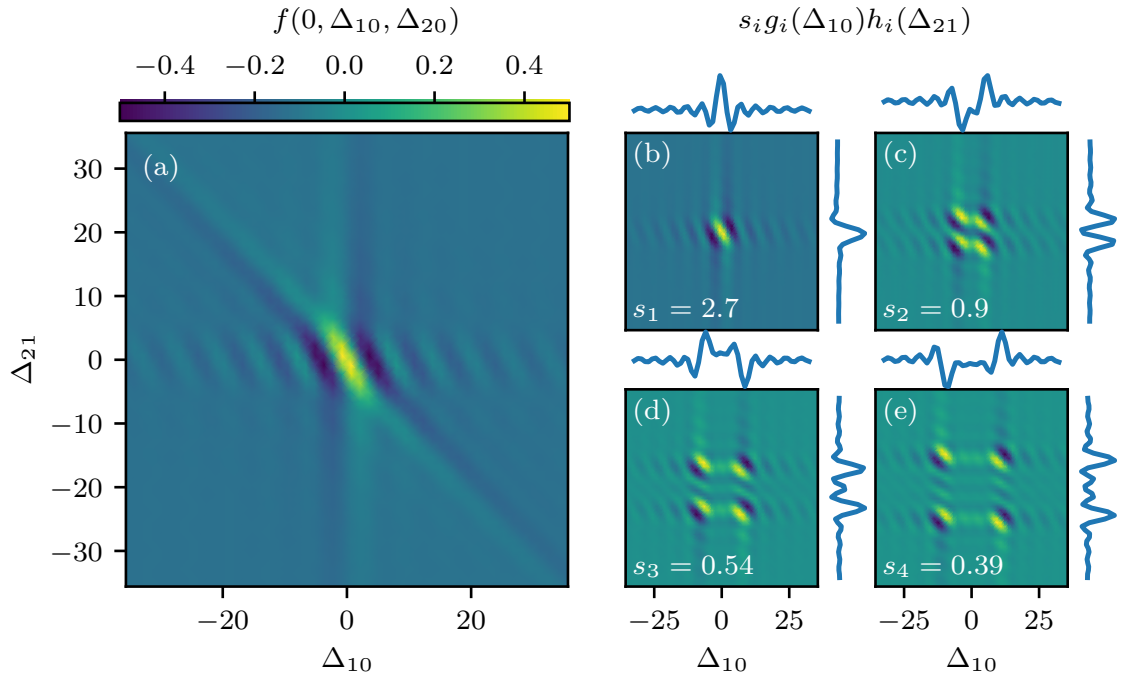

Figure 3.1 (a) Real part of the function $f(0, \Delta_{10}, \Delta_{20})$. (b-d) SVD Decomposition of $f(0, \Delta_{10}, \Delta_{20})$, showing the functions corresponding to the first 4 singular values. The blue curves on the top and right-hand side of the subplots are the functions $g_i(\Delta_{10})$ and $h_i(\Delta_{21})$, respectively. As can be seen, the decomposition fails to capture the diagonal behaviour (where $\lambda_2 = \lambda_0$), meaning that additional terms in this series would be required to capture this resonance for large eigenvalue differences.

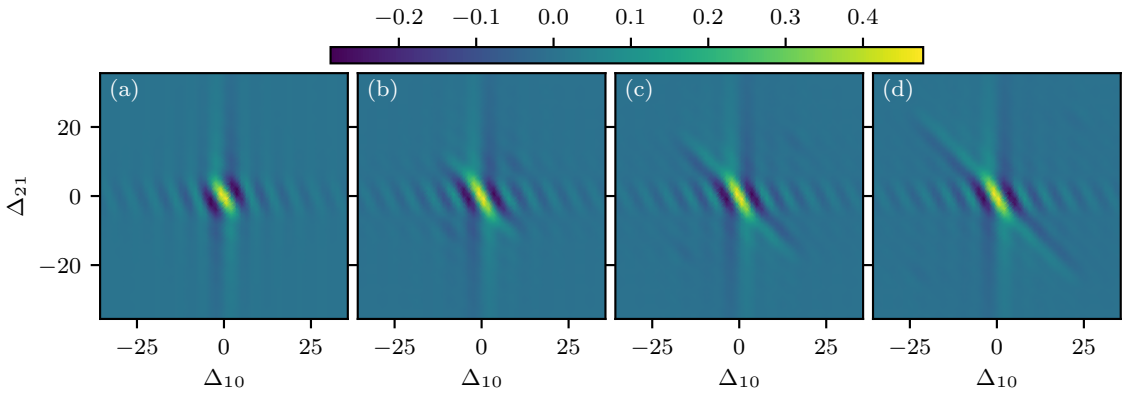

Figure 3.2 Approximating the divided difference function $f(0, \Delta_{10}, \Delta_{20})$ using the SVD (see Fig. 3.1(a)), with (a) one, (b) three, (c) five and (d) seven singular value contributions. The diagonal behaviour along the line $\Delta_{21} = -\Delta_{10}$ is gradually captured as more singular value contributions are used.

such that the third-order Dyson operator would be written as

$$\hat{S}^{(3)}(1) \approx (-i)^3 \sum_i \left(G_i \odot \hat{X}\right)\left(H_i \odot \hat{X}\right)\left(K_i \odot \hat{X}\right), \tag{3.36}$$

with $G_i,\ H_i,\ K_i$ defined as in Eq. (3.31).

Unfortunately, there is no direct analogy to a Singular Value Decomposition of a tensor, only a more general scheme known as a 'Tucker decomposition' [76, 77]. However, there is a closely related alternative that can construct the approximation in Eq. (3.35), referred to as 'Candecomp-parafac' (CP) decomposition [78], which can be extended to arbitrary tensor dimensions. Fig. 3.3 provides a visualization of this decomposition, where an $M \times M \times M$ tensor is decomposed into a sum of outer products of vectors.

Similarly to Sect. 3.5.1, we could theoretically conduct M^3 evaluations of this function, evenly sampled over the three variables Δ_{i1}, store them in an $M \times M \times M$ tensor, and conduct a CP decomposition of this tensor, which would yield a *discretized* approximation to our targeted functions in Eq. (3.35). However, these divided difference functions are highly oscillatory, making the problem challenging without taking a very large M. Additionally, we aim for the solver to be most precise when dealing with small eigenvalues differences, as these correspond to resonant transitions.

As an alternative, one could consider the Taylor expansion of the divided differences functions and construct a CP decomposition of the *coefficient tensor* $\mathcal{C}$, i.e. the tensor of all the coefficients of from the Taylor expansion.:

$$f(0, \Delta_{10}, \Delta_{20}, \Delta_{30}) = \left(\sum_{k_0,k_1,k_2=0}^{d} \mathcal{C}_{k_0,k_1,k_2} \Delta_{10}^{k_0} \Delta_{20}^{k_1} \Delta_{30}^{k_2} \right). \tag{3.37}$$

A CP decomposition of this coefficient tensor $\mathcal{C}$ would yield:

$$\mathcal{C} = \sum_{i=1}^{b} g_i h_i k_i, \quad \mathcal{C}_{k_0,k_1,k_2} = \sum_{i=1}^{b} i(g_i[k_0])(h_i[k_1])(k_i[k_2]), \tag{3.38}$$

where g_i, h_i, k_i are *vectors of coefficients* and b is the order of the CP expansion. We can easily map these vectors to functions with an input Δ,

$$g_i(\Delta) = \sum_j c_j \Delta^j. \tag{3.39}$$

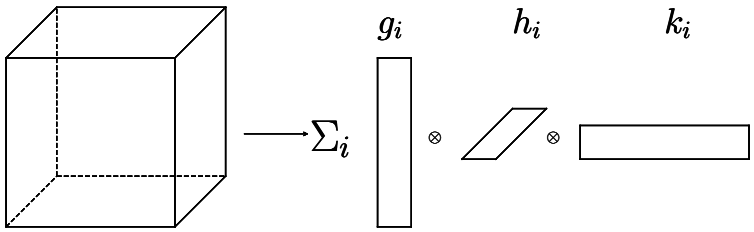

Figure 3.3 Example of the CP decomposition of a $N \otimes N \otimes N$ tensor.

This means we can rewrite Eq. (3.37) in the form

$$f(0, \Delta_{10}, \Delta_{20}, \Delta_{30}) = \sum_i g_i(\Delta_{10}) h_i(\Delta_{21}) k_i(\Delta_{32}) \tag{3.40}$$

which is exactly the form of Eq. (3.35), allowing us to compute the third order Dyson operator as in Eq. (3.36) without the need to calculate tensors.

The above methodology can be extended to arbitrary order – hence the name 'GeneRalized Order Nth Decomposition' (GROND) for this general procedure. This decomposition works excellently for small to moderate eigenvalue differences –from numerical tests, with eigenvalue differences $|\Delta_{ij}| \leq 7$, I found that even the fourth order divided difference expressions converged within machine precision to the analytical functions in Eq. (3.19) with a Taylor expansion of order $d = 20$ and of order $b = 7$ in the CP decomposition. Furthermore, this method only requires operations of matrices of size N^2 at all orders in the Dyson series, resulting in significant memory saving in comparison to the analytic method in Sect. 3.3, which for dense matrices scales as N^{n+1}, with n the Dyson series order. This allows for the simulation of large $N \geq 1000$ system sizes for which the analytical method would be out of reach.

3.5.4 Outlook and Future Work

Whilst the above procedure works excellently for small eigenvalue differences and provides a pathway to building the Dyson operators for large, dense Hilbert space sizes, there are several drawbacks – most notably, this method cannot capture large eigenvalue differences. Calculating the coefficients using an alternative basis – such as spherical Bessel functions or a Fourier series – could result in better stability and convergence over a larger range of Δ. Further, it is unclear how or whether this decomposition should be weighted such that small eigenvalue differences are penalized more, encouraging convergence for small Δ. I leave

addressing these questions to future work.

3.6 Publication – Fast and Differentiable Simulation of Driven Quantum Systems [1]

In this publication, I look at how the Dyson series can be applied more generally to Eq. (3.2) proving the forms of the divided difference expressions and operators defined in this chapter in the presence of an oscillatory drive. I illustrate the solver's numerical prowess showing orders of magnitude speed-up over the implemented solver for unitary dynamics in QuTiP. Finally, I demonstrate that the generated unitary evolution is differentiable with respect to the drive amplitude, allowing for simple pulse optimization, and show how this can be used to optimize a two-qubit gate implemented on superconducting qubit hardware.

For this work, I proposed the idea of the Dysolve algorithm, developed the python scripts and wrote the manuscript. Jonathan Gross helped extensively with the mathematical derivations and inductive proofs, Agustin Di Paolo and Alexandre Blais helped guide the project with the proposed applications of the algorithm, and Élie Genois made the code significantly more efficient, performed many of the simulations and provided the benchmarking results.

The Dysolve algorithm was adapted and implemented in Qiskit Dynamics by a team from IBM and can offer significant performance improvements over standard numerical integrators, see Ref. [72] for more details.

Fast and differentiable simulation of driven quantum systems

Ross Shillito,[1,*] Jonathan A. Gross,[1,†] Agustin Di Paolo,[1,‡] Élie Genois,[1] and Alexandre Blais[1,2]
[1]*Institut quantique and Département de physique, Université de Sherbrooke, Sherbrooke, Quebec, Canada J1K2R1*
[2]*Canadian Institute for Advanced Research, Toronto, Ontario, Canada M5G1M1*

(Received 12 February 2021; accepted 11 August 2021; published 20 September 2021)

The controls enacting logical operations on quantum systems are described by time-dependent Hamiltonians that often include rapid oscillations. In order to accurately capture the resulting time dynamics in numerical simulations, a very small integration time step is required, which can severely impact the simulation run time. Here, we introduce a semianalytic method based on the Dyson expansion that allows us to time-evolve driven quantum systems much faster than standard numerical integrators. This solver, which we name `Dysolve`, efficiently captures the effect of the highly oscillatory terms in the system Hamiltonian, significantly reducing the simulation's run time as well as its sensitivity to the time-step size. Furthermore, this solver provides the exact derivative of the time-evolution operator with respect to the drive amplitudes. This key feature allows for optimal control in the limit of strong drives and goes beyond common pulse-optimization approaches that rely on rotating-wave approximations. As an illustration of our method, we show results of the optimization of a two-qubit gate using transmon qubits in the circuit QED architecture.

DOI: 10.1103/PhysRevResearch.3.033266

I. INTRODUCTION

High-fidelity logical gates are paramount to the realization of useful quantum computation. It is important in the development of these gates that they be simulated with great precision to ensure that they meet the particularly strict requirements for fault-tolerant quantum computation. Several techniques are used for simulating the time dynamics of quantum devices, including dynamical solvers such as Runge-Kutta integrators and direct matrix exponentiation [1,2]. However, capturing the full time dynamics with the necessary accuracy requires integration methods with a sufficiently small time step. As a result, simulations can be computationally very expensive, even for relatively simple cases such as optimizing a two-qubit gate.

To ensure that simulations of quantum systems are feasible, approximations must be made. For example, a common approximation is to neglect counter-rotating terms within the rotating-wave approximation (RWA) to greatly reduce the complexity of the simulation while capturing the dominant dynamics. This approximation renders the Hamiltonian time independent, which can then simply be exponentiated to obtain the propagator. However, effects such as the Bloch-Siegert shift which do not appear under a RWA need to be taken into account to accurately model the system [3]. In addition, many gate optimization methods, such as GRadient Ascent Pulse Engineering (GRAPE), require the calculation of gradients, something which greatly adds to the complexity of the numerical calculations [4]. More specifically, when including the effects of the counter-rotating terms, there exists no simple derivative of time-ordered unitaries with respect to the drive amplitudes, and one must resort to approximating the gradients. Consequently, this approach may not converge to the optimal solution, which is problematic when targeting very high-fidelity gates.

In this work, we develop an algorithm based on a Dyson series expansion of the time-ordered problem that addresses all of the above difficulties. In this approach, which we call `Dysolve`, the time-ordered unitary evolution operator is written as a product of time-independent operators which are weighted by the drive amplitudes and dynamical phases. This algorithm captures the full fast-oscillatory dynamics irrespective of the integration step size, thereby decreasing the complexity of the numerical problem. This also greatly decreases the simulation time in comparison to traditional integration-based solvers without loss of numerical precision. Importantly, this approach trivializes the derivative with respect to the drive strength, which can be calculated to an accuracy equivalent to the order of the Dyson series. Moreover, this approach is compatible with non-Hermitian dynamics, allowing for the simulation of open quantum systems.

We begin by introducing our approach in Sec. II in the case of a single, sinusoidal drive with a constant amplitude, and then extend the formalism to the case of multiple drives, accounting for filtering effects on envelope functions. We then define the `Dysolve` algorithm in Sec. III, and demonstrate its application to driven quantum systems with random drive

*Ross.Shillito@USherbrooke.ca
†Jarthurgross@google.com
‡adipaolo@mit.edu

envelopes. We proceed to apply our algorithm to the GRAPE optimization routine in Sec. IV, and show as an example optimized two-qubit gates in the circuit QED architecture.

II. OSCILLATORY DRIVE PROBLEM

A. Simple time-dependent Hamiltonians

We begin by considering a simple time-dependent Hamiltonian with a cosinusoidally oscillating control drive term,

$$\hat{H}(t) = \hat{H}_0 + \hat{V}(t), \quad \hat{V}(t) = \hat{X}\cos(\omega t). \tag{1}$$

Here, $\hat{H}_0 = \sum_k \lambda_k |k\rangle\langle k|$ is a generic system Hamiltonian expressed in its eigenbasis, while $\hat{X}$ is a dipole operator that connects the eigenstates of $\hat{H}_0$ and which we take to account for the amplitude of the drive. As will become important later, we note that we are not using the RWA.

The propagator under $\hat{H}(t)$ for some time increment δt takes the usual form of a time-ordered integral

$$\hat{U}(t, t+\delta t) = \mathcal{T}\exp\left(-i\int_t^{t+\delta t} dt' \hat{H}(t')\right), \tag{2}$$

with $\mathcal{T}$ the time-ordering operator. Due to the fast oscillatory $\cos(\omega t)$ term, evaluating the propagator $\hat{U}(t, t+\delta t)$ is a numerically challenging problem, and there exist few analytic solutions to even the simplest case of a driven two-level qubit [5]. In special cases, one can invoke the RWA to remove the explicit time dependence from the Hamiltonian, thereby allowing for calculation of the propagator from matrix exponentiation directly [6].

However, when the RWA cannot be used due to a breakdown of the approximation or because a greater degree of accuracy is needed, we propose using a truncated Dyson series. Consider Eq. (2), written using the definition of the time-ordering operator:

$$\hat{U}(t, t+\delta t) = \sum_{n=0}^{\infty}(-i)^n \int_t^{t+\delta t}\int_t^{t_n}\cdots\int_t^{t_2} \hat{H}(t_n)\cdots\hat{H}(t_1)dt_1\cdots dt_n. \tag{3}$$

It is useful to express this expansion in terms of powers of the drive operator $\hat{V}(t)$, forming the Dyson series

$$\hat{U}(t, t+\delta t) = \sum_{n=0}^{\infty}\hat{U}^{(n)}(t, t+\delta t), \tag{4}$$

where we have defined

$$\hat{U}^{(n)}(t, t+\delta t) = \sum_{\boldsymbol{\omega}_n}\exp\left(i\sum_{p=1}^{n}\boldsymbol{\omega}_n[p]t\right)\hat{S}^{(n)}(\boldsymbol{\omega}_n, \delta t). \tag{5}$$

Here, $\boldsymbol{\omega}_n$ is an n-vector whose entries are $\pm\omega$ originating from the decomposition of the $\cos(\omega t)$ of the control into complex exponentials, and $\boldsymbol{\omega}_n[p]$ is the pth element of $\boldsymbol{\omega}_n$. The sum over $\boldsymbol{\omega}_n$ implies a summation over all 2^n possible $\boldsymbol{\omega}_n$ vectors.

Equation (5) also introduces the Dyson series operator $\hat{S}^{(n)}(\boldsymbol{\omega}_n, \delta t)$ which takes the form

$$\hat{S}^{(n)}(\boldsymbol{\omega}_n, \delta t) = \frac{1}{2^n}\sum_{\boldsymbol{m}\in\mathbf{Z}_+^{n+1}}\hat{S}_{\boldsymbol{m}}^{(n)}(\boldsymbol{\omega}_n, \delta t), \tag{6}$$

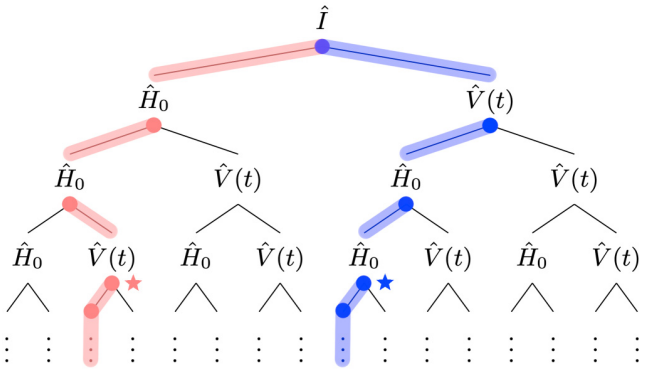

FIG. 1. Tree diagram showing the different branches of the time-ordered integral, with two example branches highlighted.

and which corresponds to a summation over all Dyson path operators of nth order, where $\boldsymbol{m} = [m_n, \ldots, m_0]$ with each index m_i ranging from zero to infinity. These nth-order path operators are given by

$$\hat{S}_{\boldsymbol{m}}^{(n)}(\boldsymbol{\omega}_n, \delta t) = \int_0^{\delta t}\int_0^{t'_M}\cdots\int_0^{t'_2}(-i\hat{H}_0)^{m_n}\hat{X}(-i\hat{H}_0)^{m_{n-1}}\cdots\hat{X}(-i\hat{H}_0)^{m_0} \times (-i)^n\prod_{p=1}^{n}\exp(i\boldsymbol{\omega}_n[p]t_{\iota(p)})dt'_1\cdots dt'_M, \tag{7}$$

where we have introduced

$$M = \sum_{i=0}^{n} m_i + n, \quad \iota(p) = \sum_{j=0}^{p-1} m_j + p. \tag{8}$$

Each $\hat{S}_{\boldsymbol{m}}^{(n)}(\boldsymbol{\omega}_n, \delta t)$ corresponds to different ways to have n contributions from the control $\hat{X}$ (i.e., different terminated branches of the tree diagram), with the m_i's labeling the number of applications of $\hat{H}_0$ before the subsequent application of $\hat{X}$. The total number of operators in the path operator is given by M. To simplify the notation, we introduce an indexing function $\iota(p)$, which can be interpreted as the total number of $\hat{H}_0$ and $\hat{X}$ operators before the pth application of the control $\hat{X}$. These definitions can be visualized as in Fig. 1 where the highlighted branches correspond to sets of path operators $\hat{S}_{[m_1,2]}^{(n)}(\boldsymbol{\omega}_n, \delta t)$ and $\hat{S}_{[m_1,0]}^{(n)}(\boldsymbol{\omega}_n, \delta t)$, respectively, with m_1 determined by where the path terminates. For example, $\hat{S}_{[0,2]}^{(n)}(\boldsymbol{\omega}_n, \delta t)$ and $\hat{S}_{[2,0]}^{(n)}(\boldsymbol{\omega}_n, \delta t)$ terminate at the points in the branches marked by a star.

Below, we give explicit expressions for these operators to zeroth and first order in the control. Building on these results, we then construct expressions that are easily amenable to efficient numerical evaluation to arbitrary orders.

B. Evaluation to zeroth and first order

The zeroth order corresponds to the leftmost branch of the tree diagram of Fig. 1 for which it is straightforward to obtain an explicit expression. Indeed, the path operator simply takes

the form

$$\hat{S}^{(0)}_{[m_0]}(\mathbf{0},\delta t) = \int_0^{\delta t}\int_0^{t_{m_0}}\cdots\int_0^{t_2}(-i\hat{H}_0)^{m_0}dt_1\cdots dt_{m_0-1}dt_{m_0} = \frac{(-i\hat{H}_0\delta t)^{m_0}}{m_0!}. \tag{9}$$

Summing all of the elements in the branch, we obtain

$$\hat{S}^{(0)}(\mathbf{0},\delta t) = \sum_{m_0=0}^{\infty}\hat{S}^{(0)}_{[m_0]}(\mathbf{0},\delta t) = e^{-i\hat{H}_0\delta t}, \tag{10}$$

which corresponds, as expected, to the drift evolution of the Hamiltonian in the absence of drive. In a similar fashion, the path operator to first order takes the form

$$\hat{S}^{(1)}_{[m_1,m_0]}(\boldsymbol{\omega}_1,\delta t) = \int_0^{\delta t}\int_0^{t_M}\cdots\int_0^{t_2}(-i\hat{H}_0)^{m_1}\hat{X}(-i\hat{H}_0)^{m_0} \times \exp(\pm i[\omega]t_{\iota(1)})dt_1\cdots dt_M, \tag{11}$$

and therefore leads to the following Dyson series operator:

$$\hat{S}^{(1)}(\boldsymbol{\omega}_1,\delta t) = -\frac{i\delta t}{2}\sum_{k,k'}f(\lambda_k\delta t,(\lambda_{k'}\mp\omega)\delta t)\langle k|\hat{X}|k'\rangle|k\rangle\langle k'|. \tag{12}$$

This operator weights the matrix elements of $\hat{X}$ by a function f whose inputs are weighted eigenvalues $\{\lambda_k\}$ of the free Hamiltonian $\hat{H}_0$ with corresponding eigenstates $\{|k\rangle\}$, where the second eigenvalue is shifted by the drive frequency $\boldsymbol{\omega}_1[1] = \pm\omega$. This function is defined as

$$f(\lambda_k,\lambda_{k'}) = \frac{i}{\lambda_k-\lambda_{k'}}(e^{-i\lambda_k}-e^{-i\lambda_{k'}}). \tag{13}$$

We refer to this function as the first-order weighting function. The above expressions are derived in the Appendix A, and are unsurprisingly in exact agreement with first-order time-dependent perturbation theory.

Crucially, the Dyson operators depend on the size δt of the time increment, but not the current time of the evolution t. As a result, for a total evolution time $T = P\delta t$, where P is an integer, the set of P incremental evolution operators $\hat{U}(p\delta t,(p+1)\delta t)$ can be evaluated simultaneously. As will become clearer in the next section, this holds true to arbitrary order.

C. Evaluation to *n*th order

To model quantum systems with sufficient accuracy, it is necessary to consider second- and higher-order terms in the Dyson series. This can be done following a similar approach as described above. Indeed, we introduce the nth-order Dyson operator in a similar fashion to Eq. (12):

$$\hat{S}^{(n)}(\boldsymbol{\omega}_n,\delta t) = \left(\frac{-i\delta t}{2}\right)^n\sum_{\boldsymbol{k}_n}f(\boldsymbol{\lambda}_n(\boldsymbol{k}_n)\delta t-\boldsymbol{c}(\boldsymbol{\omega}_n)\delta t)\langle k^{(n)}|\hat{X}|k^{(n-1)}\rangle \times\langle k^{(n-1)}|\hat{X}\cdots|k^{(1)}\rangle\langle k^{(1)}|\hat{X}|k^{(0)}\rangle|k^{(n)}\rangle\langle k^{(0)}|. \tag{14}$$

Here, each $\boldsymbol{k}_n = (k^{(0)},k^{(1)},\ldots,k^{(n)})$ is a set of indices which specify a set of $(n+1)$ eigenstates $\{|k^{(m)}\rangle\}$ of $\hat{H}_0$ with corresponding eigenvalues $\boldsymbol{\lambda}_n(\boldsymbol{k}_n)$, and we sum over all possible $\boldsymbol{k}_n$. These eigenvalues are written in vector form:

$$\boldsymbol{\lambda}_n(\boldsymbol{k}_n)\equiv(\lambda_{k^{(0)}},\ldots,\lambda_{k^{(n)}}),\quad \hat{H}_0|k^{(m)}\rangle = \lambda_{k^{(m)}}|k^{(m)}\rangle. \tag{15}$$

The sum over $\boldsymbol{k}_n$ implies a summation over all sets of eigenstates which the dipole operator $\hat{X}$ couples in Eq. (14). Additionally, we have introduced the cumulative vector

$$\boldsymbol{c}(\boldsymbol{\omega}_n) = \left(\sum_{p=0}^{n-1}\boldsymbol{\omega}_n[n-p],\sum_{p=0}^{n-2}\boldsymbol{\omega}_n[n-p],\ldots,\boldsymbol{\omega}_n[n],0\right). \tag{16}$$

The nth-order weighting function $f(\boldsymbol{\lambda}_n)$ entering Eq. (14) can be obtained recursively using the relation (see Appendix A)

$$f(\boldsymbol{\lambda}_n) = i\frac{f(\boldsymbol{g}(\boldsymbol{\lambda}_n))-f(\boldsymbol{g}^2(\boldsymbol{\lambda}_n)\cup\boldsymbol{\lambda}_n[n])}{\boldsymbol{\lambda}_n[n-1]-\boldsymbol{\lambda}_n[n]}, \tag{17}$$

where $\boldsymbol{g}(\boldsymbol{v}_n)$ returns $\boldsymbol{v}_n$ without its last element, $\boldsymbol{g}^2(\boldsymbol{v}_n) = \boldsymbol{g}(\boldsymbol{g}(\boldsymbol{v}_n))$, and the notation $\cup$ indicates appending an additional element to a vector such that

$$\boldsymbol{\lambda}_n = \boldsymbol{g}(\boldsymbol{\lambda}_n)\cup\boldsymbol{\lambda}_n[n]. \tag{18}$$

In the case of degenerate eigenvalues, we simply define Eq. (17) by taking the limit $\boldsymbol{\lambda}_n[n-1]\to\boldsymbol{\lambda}_n[n]$.

Using Eq. (4) with the above results, we can now form the truncated Dyson series yielding an approximation to the evolution operator to nth order in the perturbation

$$\hat{U}_p\approx\sum_{r=0}^{n}\hat{U}_p^{(r)},\quad \hat{U}_p^{(r)}\equiv\hat{U}^{(r)}(p\delta t,(p+1)\delta t), \tag{19}$$

thus yielding the total evolution operator for time $T = P\delta t$ to the same order,

$$\hat{U}(0,P\delta t)\approx\mathcal{T}\prod_{p=0}^{P}\hat{U}_p. \tag{20}$$

In this formalism, the time-ordering operator $\mathcal{T}$ becomes a trivial operation since we need only arrange the set of matrices $\{\hat{U}_p\}$ in ascending order from right to left. Using this method, we can thus calculate an arbitrary number of terms in the Dyson series, with each subsequent order increasing the accuracy of the approximation to the propagator operator $\hat{U}(0,T)$.

D. Generalization to more complex drives

So far, we have only considered a single, cosinusoidal drive. In practice, for applications such as Derivative Removal Adiabatic Gates [7], it is useful to consider more complicated drives of the form

$$\hat{V}(t) = 2[\Omega_x\cos(\omega t)+\Omega_y\sin(\omega t)]\hat{X}, = (\Omega e^{i\omega t}+\Omega^* e^{-i\omega t})\hat{X}, \tag{21}$$

where $\Omega = \Omega_x + i\Omega_y$ is the complex drive amplitude. This generalization requires only a minor modification to the result of Eq. (5) which now reads

$$\hat{U}^{(n)}(t,t+\delta t) = \sum_{\boldsymbol{\omega}_n}\exp\left(i\sum_{p=1}^{n}\boldsymbol{\omega}_n[p]t\right)\Omega(\boldsymbol{\omega}_n)\hat{S}^{(n)}(\boldsymbol{\omega}_n,\delta t), \tag{22}$$

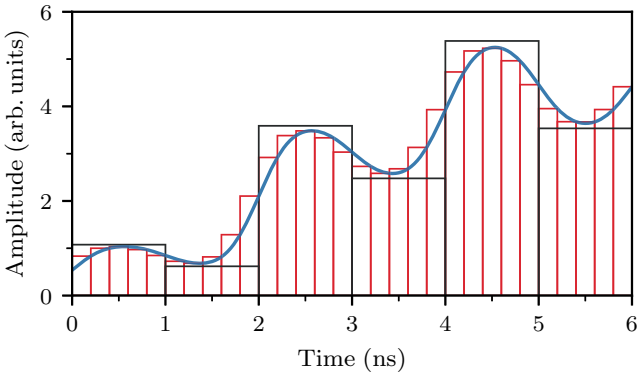

FIG. 2. An example set of pulse amplitudes. The black bars indicate the chosen drive amplitudes u_j, where $\Delta t = 1$ ns. The red bars indicate the subpixels, which provide an approximate interpolation of the true pulse delivered to the system (blue), with a bandwidth $\omega_0/2\pi = 851$ MHz.

and where we have introduced

$$\Omega(\boldsymbol{\omega}_n) = \Omega^{\mu(\boldsymbol{\omega}_n)}\Omega^{*(n-\mu(\boldsymbol{\omega}_n))}. \tag{23}$$

We see from Eq. (21) that Ω and Ω^* will appear according to the number of positive and negative frequency elements in the vector $\boldsymbol{\omega}$, respectively, leading to a simple expression for $\mu(\boldsymbol{\omega}_n)$ in $\Omega(\boldsymbol{\omega}_n)$:

$$\mu(\boldsymbol{\omega}_n) = \left(n + \sum_{p=1}^{n} \boldsymbol{\omega}_n[p]\right)/(2\omega). \tag{24}$$

As shown in the Appendix B, this formalism can be further extended to an arbitrary number of drives of different frequencies, and acting on different system operators.

E. Envelope functions and Gaussian filtering

Having established the necessary notation, we can now take into account control drives with time-dependent amplitudes $\Omega(t)$. To do so, we discretize the drive envelope of total time $T = N_p\Delta t$ into N_p increments, called pixels, each of duration Δt. For a given pixel i, the drive is characterized by its complex amplitude u_i such that the envelope can be expressed as [4]

$$\Omega(t) = \sum_{i=0}^{N_p-1} u_i \sqcap (i\Delta t, (i+1)\Delta t), \tag{25}$$

where $\sqcap(t, t+\Delta t) = \Theta(t) - \Theta(t-\Delta t)$, with $\Theta(t)$ the Heaviside function.

Additionally, following Motzoi *et al.* [8], we take into account the finite bandwidth of the control by applying a Gaussian filter on the discretized enveloped. To do so, each pixel is subdivided into N_s subpixels of width δt with $\Delta t = N_s\delta t$ (see Fig. 2). The subpixel amplitudes s_l are then defined as

$$s_l = \sum_{j=0}^{N_s-1} T_{l,j}u_j, \tag{26}$$

where the Gaussian filter matrix T has elements

$$T_{l,j} = \frac{1}{2}\left\{\mathrm{erf}\left[\omega_0\left(\frac{l\delta t - j\Delta t}{2}\right)\right] - \mathrm{erf}\left[\omega_0\left(\frac{l\delta t - (j+1)\Delta t}{2}\right)\right]\right\}. \tag{27}$$

Following Eq. (22), the nth-order evolution operator over the lth subpixel takes the form

$$\hat{U}_l^{(n)} = \sum_{\boldsymbol{\omega}_n} \exp\left(i\sum_{p=1}^{n} \boldsymbol{\omega}_n[p] l\delta t\right)\Omega_l(\boldsymbol{\omega}_n)\hat{S}^{(n)}(\boldsymbol{\omega}_n, \delta t), \tag{28}$$

where the drive function in Eq. (23) picks up an additional subscript l to denote the lth subpixel:

$$\Omega_l(\boldsymbol{\omega}_n) = s_l^{\mu(\boldsymbol{\omega}_n)}s_l^{*(n-\mu(\boldsymbol{\omega}_n))}. \tag{29}$$

As can be seen in Fig. 2, the subpixels (red) will often overestimate or underestimate the filtered pulse (blue) amplitude depending on its gradient, something which can become a leading contribution to the error in simulations. In Appendix C, we generalize the amplitudes s_l to have a linear time dependence to compensate for the change in amplitude of the continuous pulse across a single subpixel, providing a more accurate approximation to the filtered pulse.

III. THE DYSOLVE ALGORITHM

As already explained, the Dyson series operators $\hat{S}^{(n)}(\boldsymbol{\omega}_n, \delta t)$ for which we have expressions at arbitrary order n are functions of the time increment δt and, crucially, are independent of the total evolution time T. The `Dysolve` algorithm leverages this fact to parallelize the time evolution.

The algorithm operates in two parts: a *preparation* stage and a *contraction* stage. In the preparation stage, the Dyson operators $\hat{S}^{(n)}(\boldsymbol{\omega}_n, \delta t)$ for a Hilbert space size N are computed up to a chosen truncation order n, and arranged in a tensor. This tensor has dimensions $(R \times N \times N)$, where R is the total number of Dyson operators. In the case of a single drive without linear interpolation, $R = 2^{n+1} - 1$, where n is the order of the expansion.

Once the preparation stage is completed, it is in principle possible to consider arbitrary gate times and drive envelopes. Suppose we wish to simulate a time evolution of length $T = P\delta t$, where P is an integer. We first generate a $(P \times R)$ tensor whose elements are the envelopes and oscillatory terms $\exp(i\sum_{p=1}^{n} \boldsymbol{\omega}_n[p]l\delta t)\Omega_l(\boldsymbol{\omega}_n)$ in Eq. (28). We then multiply these tensors to obtain a $(P \times N \times N)$ time-evolution tensor, where the p th $(N \times N)$ matrix corresponds to a time-step operator $\hat{U}_p$. This multiplication constitutes the parallelized portion of the algorithm, with all P time-evolution operators calculated simultaneously. As in Eq. (20), we multiply all of the individual evolution operators in the time-evolution tensor to obtain the final evolution operator $\hat{U}(0, T)$. This whole procedure forms the contraction step of our algorithm.

Below, we will refer to the computation of the evolution operator $\hat{U}(0, T)$ to nth order following the above approach as `Dysolve-n`.

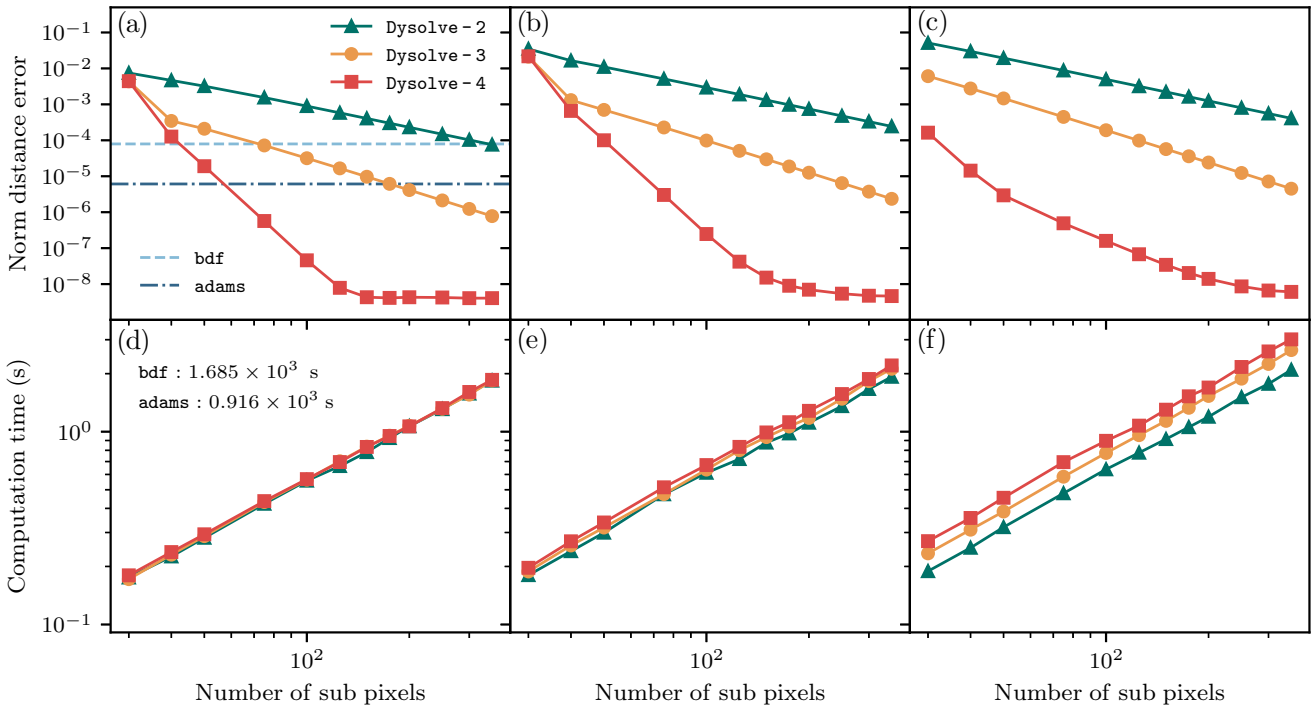

FIG. 3. `Dysolve` benchmarks. Frobenius norm distance from the propagator $\hat{U}_{\text{ref}}(0, T)$ for the `Dysolve` algorithm for a $T = 500$ ns evolution with (a) one input drive, (b) two input drives, and (c) three input drives. Contraction time of the `Dysolve` algorithm as a function of the subpixel number N_s for (d) one, (e) two, and (c) three input drives. The slope of the data is precisely 1, meaning that the computational time scales linearly with the number of subpixels. Preparation stage computation times are (10 ms, 18 ms, 24 ms) for (`Dysolve`-2, 3, 4) and one input drive, (22 ms, 68 ms, 136 ms) for two input drives, and (45 ms, 200 ms, 442 ms) for three input drives. In panels (a) and (d), we compare the error and computation time, respectively, of `Dysolve` with QUTIP `bdf` and `adams` solvers using high-precision settings (see Appendix D for more details).

A. Benchmarking

Before turning to examples of application of `Dysolve`, we first benchmark this algorithm. To do so, there are a number of factors to consider: (i) the size of the Hilbert space, (ii) the drive amplitude, (iii) the number of independent drive channels, and (iv) the shape of the envelope function. To quantify the performance of our algorithm, we use two metrics. Given a number of subpixels N_s, we evaluate (1) the computation time and (2) the Frobenius norm distance between $\hat{U}(0, T)$, the propagator calculated under the `Dysolve` algorithm with a chosen number of subpixels, and a reference $\hat{U}_{\text{ref}}(0, T)$, the same unitary calculated with very high precision. As discussed further in Appendix D, we use `Dysolve`-4 with 10^4 subpixels to compute $\hat{U}_{\text{ref}}(0, T)$, since `Dysolve` is able to reach precisions on the benchmark setup that are up to three orders of magnitude greater than QUTIP's `propagator`, a comparable dynamical solver [9].

For our benchmarks, we use diagonal system Hamiltonians with a Hilbert space size $N = 25$. As a concrete example, we take the eigenvalues to be normally distributed about 7 GHz, the typical operating frequency of superconducting qubits [10]. We consider between one and three input drive operators, at the frequency of the $|0\rangle \leftrightarrow |1\rangle$, $|2\rangle \leftrightarrow |3\rangle$, and $|4\rangle \leftrightarrow |5\rangle$ transitions with $|k\rangle$ the kth lowest-lying system eigenstate. The matrix elements of each operator corresponding to these transitions are set to 1. To emulate an arbitrary drive operator with off-resonant terms, we populate 20% of the remaining matrix elements of the drive operators with complex numbers normally distributed about zero, after which Hermiticity is enforced by the addition of the complex conjugate. The envelope function associated to each drive operator is centered around an amplitude such that the duration of the simulation is equivalent, in the absence of the other drives, to a total of 20 Rabi oscillations. In the context of superconducting qubits with their microwave drives, this corresponds to the simulation a 500-ns evolution with a drive amplitude of 40 MHz. We take the pixel amplitudes u_j of the envelope functions to be normally distributed about 40 MHz with a standard deviation of 1 MHz. To reduce statistical fluctuations, the results presented in Fig. 3 are averaged over 30 simulations, each with different random envelope functions and system eigenvalues.

The numerical simulations reported in Fig. 3 are performed on a 2.8-GHz Intel Xeon Gold 6242 Processor (16 cores/32 threads) using PYTHON. Since the preparation stage only needs to be performed once for a particular system Hamiltonian, the reported computation time accounts only for the contraction stage of the `Dysolve` algorithm. We report the preparation times in the figure caption for reference. Figures 3(a)–3(c) present the norm distance between `Dysolve-n` for $n = 2$, 3, and 4 as a function of the number of subpixels, and for increasing number of input drives. Even with large drive amplitudes and long evolution times, we obtain an excellent approximation to the final unitary operation with relatively few subpixels. For example, using a fourth-order Dyson expansion with 40 subpixels yields a

norm distance error of less than 10^{-5} in all cases. Importantly, as shown in Figures 3(d)–(f), the computation time needed to reach this level of accuracy is only on the order of a few seconds (less than 3 s). In comparison, QUTIP's `propagator` function takes about 16 min to perform the calculation to an equivalent accuracy for a single input drive. Such a significant speedup proves to be particularly useful when, as discussed in section Sec. IV, the contraction stage of the algorithm needs to be repeated many times in an optimization loop. Additional benchmarking results are provided in Appendix D.

IV. APPLICATION TO OPTIMAL CONTROL

Optimal control is an essential tool in the development of high-fidelity gates for quantum computation. Most optimization algorithms, such as GRAPE [4], rely on calculating the gradient of the gate fidelity with respect to the drive amplitude at each pixel. With most approaches, this requires recalculating the evolution operator at each time interval, something which can be as expensive as the original calculation of the unitaries. Moreover, this calculation is generally performed with the presumption that the Hamiltonian is time independent over the duration of a subpixel, thus invoking a form of the RWA. In our expansion, such an assumption is not required.

Here, we show how `Dysolve` can be applied to optimizing control pulse envelopes to maximize gate fidelity. More precisely, we consider optimizing the fidelity of an evolution $\hat{U}(T)$ with respect to a target gate unitary $\hat{U}_{\text{target}}$. In general, $\hat{U}(T)$ acts on the full Hilbert space of the system while $\hat{U}_{\text{target}}$ is defined on its computational subspace. The objective is thus to maximize the performance function [8]

$$\Phi = \frac{1}{d^2}|\text{Tr}\hat{U}^{\dagger}_{\text{target}}\hat{U}(T)\hat{\mathcal{P}})|^2, \tag{30}$$

where d is the dimension of the computational subspace ($d = 2$ for a single qubit), and $\hat{\mathcal{P}}$ is the projector on that subspace. Although our approach can in principle deal with an arbitrary number of drives, for simplicity here we consider control of a single set of complex drive amplitudes u_j.

We use a GRAPE-like approach to maximize the gate fidelity which requires the gradient of the cost function Φ with respect to the drive amplitude at each pixel u_j and subpixel s_l [8]:

$$\begin{aligned}\frac{\partial\Phi}{\partial u_j} &= \sum_{l=0}^{M-1} T_{l,j}\frac{\partial\Phi}{\partial s_l},\\ \frac{\partial\Phi}{\partial s_l} &= \frac{1}{d^2}2\text{Re}\left\{\text{Tr}\left[\hat{U}^{\dagger}_{\text{target}}\frac{\partial\hat{U}}{\partial s_l}\mathcal{P}\right]\text{Tr}[\hat{U}_{\text{target}}\hat{U}^{\dagger}\mathcal{P}]\right\}.\end{aligned} \tag{31}$$

Within the framework of the `Dysolve` algorithm, it is simple to evaluate the operator $(\partial\hat{U}/\partial s_l)$. Indeed, using Eq. (20), we find that for the nth-order Dyson expansion, the derivatives of the unitaries take the form

$$\begin{aligned}\frac{\partial\hat{U}}{\partial s_l} &= \left(\prod_{m=l+1}^{M}\hat{U}_m\right)\frac{\partial\hat{U}_l}{\partial s_l}\left(\prod_{p=0}^{l-1}\hat{U}_p\right),\\ \frac{\partial\hat{U}_l}{\partial s_l} &= \sum_{m=0}^{n}\sum_{\boldsymbol{\omega}_m}\exp\left(i\sum_{p=1}^{m}\boldsymbol{\omega}_m[p]l\delta t\right)\frac{\partial\Omega_l(\boldsymbol{\omega}_m)}{\partial s_l}\hat{S}^{(m)}(\boldsymbol{\omega}_m,\delta t),\\ \frac{\partial\Omega_l(\boldsymbol{\omega}_m)}{\partial s_l} &= \mu(\boldsymbol{\omega}_m)s_l^{(\mu(\boldsymbol{\omega}_m)-1)}s_l^{*(n-\mu(\boldsymbol{\omega}_m))}.\end{aligned} \tag{32}$$

Importantly, note that the Dyson operators $\hat{S}^{(n)}(\boldsymbol{\omega}_n,\delta t)$ remain unchanged by the derivative. As such we only need to perform the preparation stage of the `Dysolve` algorithm once, with the calculation of $\partial\hat{U}/\partial s_l$. Thus, the optimization iterations only require the contraction stage computation to be performed. Further, these derivatives are exact. Consequently, the effects of off-resonant and counter-rotating terms are accounted for in the derivatives. This is the strength of the `Dysolve` algorithm for optimization.

Recall from Eq. (21) that the real and imaginary components of the drive amplitudes are associated with the magnitudes of the cosine and sine quadratures, respectively. Thus, to calculate the relevant amplitude derivatives for the cosine and sine drive envelopes, one simply calculates the appropriate sum or difference of the derivatives in Eq. (31):

$$\frac{\partial\Phi}{\partial u_{j,x}} = \frac{1}{2}\left(\frac{\partial\Phi}{\partial u_j}+\frac{\partial\Phi}{\partial u_j^*}\right),\quad \frac{\partial\Phi}{\partial u_{j,y}} = \frac{-i}{2}\left(\frac{\partial\Phi}{\partial u_j}-\frac{\partial\Phi}{\partial u_j^*}\right), \tag{33}$$

where $u_j = u_{j,x} + iu_{j,y}$. To perform the GRAPE algorithm, the set of drive amplitudes is simply updated by taking steps in the direction of the gradient of the fidelity [4]

$$u_{j,(x,y)} \to u_{j,(x,y)} + \epsilon\frac{\partial\Phi}{\partial u_{j,(x,y)}}, \tag{34}$$

where ϵ is a small, positive number which need not be fixed during the optimization process. For example, one could use a backtracking line search to maximize the gain in fidelity [11,12]. In the examples presented below, we simply use a sufficiently small ϵ for the updating process to be stable.

Example: Cross-resonance gate for coupled transmons

As an example of application of our implementation of the GRAPE algorithm with `Dysolve`, we present results of the optimization of a two-qubit gate. We also demonstrate the importance of off-resonant terms in the optimization process. To this end, we consider two capacitively coupled fixed-frequency transmon qubits, illustrated in Fig. 4. Each qubit is described by a Hamiltonian of the form [10]

$$\hat{H}_j = 4E_{Cj}\hat{n}_j^2 - E_{Jj}\cos\hat{\varphi}_j, \tag{35}$$

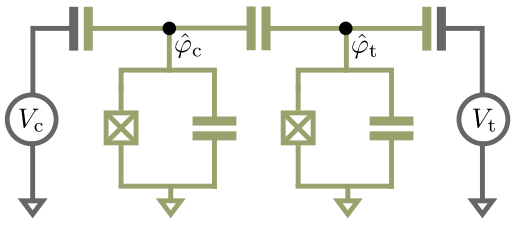

FIG. 4. Superconducting circuit for cross-resonance gates between transmon qubits. The leftmost transmon, of frequency $\omega_c/2\pi$, plays the role of the control qubit and is strongly driven by the voltage source V_c at the frequency $\omega_t/2\pi$ of the target qubit (rightmost transmon). The target qubit is also driven by V_t using a relatively weak tone that serves to give additional control. $\hat{\varphi}_c$ and $\hat{\varphi}_t$ correspond to the phase degrees of freedom associated with the control and target qubits, respectively.

where $\hat{n}_j$ and $\hat{\varphi}_j$ with $j = c, t$ are the conjugate charge and phase operators of the transmon, while E_{Cj} and E_{Jj} are the charging and Josephson energies. More precisely, we consider the cross-resonance gate, which performs an X rotation on a target qubit conditional on the state of a control qubit which is driven at the target qubit's frequency [13,14]. Optimal control has been applied to this gate before in Refs. [15,16]. However, these approaches have made use of simplified models for the device and of the rotating-wave approximation. Taking advantage of the `Dysolve` algorithm, our approach generalizes previous work on the cross-resonance gate by considering the full circuit Hamiltonian and including all counter-rotating terms.

Following the standard circuit-quantization procedure [17], the two-transmon Hamiltonian can be put in the form $\hat{H}(t) = \hat{H}_0 + \hat{V}(t)$ with

$$\hat{H}_0 = \hat{H}_c + \hat{H}_t + \bar{h}g\hat{n}_c\hat{n}_t, \quad \hat{V}(t) = f_c(t)\hat{n}_c + f_t(t)\hat{n}_t, \quad (36)$$

where $\hat{H}_c$ and $\hat{H}_t$ are the Hamiltonians of the control and target qubits, respectively, and $\bar{h}g\hat{n}_c\hat{n}_t$ results from the capacitive interaction between the qubits. The quantities $f_c(t)$ and $f_t(t)$ are the time-dependent drives

$$\begin{aligned} f_c(t) &= 2[\Omega_{cx}(t)\sin(\omega_t t) + \Omega_{cy}(t)\cos(\omega_t t)], \\ f_t(t) &= 2[\Omega_{tx}(t)\sin(\omega_t t) + \Omega_{ty}(t)\cos(\omega_t t)], \end{aligned} \quad (37)$$

with $\{\Omega_{cx}, \Omega_{cy}, \Omega_{tx}, \Omega_{ty}\}$ the drive envelope functions to be optimized by GRAPE. While activating the cross-resonance gate only requires driving the control qubit at the frequency of the target, the additional control over the target qubit $\hat{\Omega}_t(t)$ is useful to eliminate single-qubit rotations and obtain higher fidelities to the target unitary [18,19].

The above Hamiltonians capture the dynamics of the system beyond the rotating-wave approximation. To illustrate the impact of the fast-oscillating terms in this Hamiltonian, we make the RWA by eliminating transitions that are off resonant by more than some threshold frequency ω_{thres}. More specifically, our drive operator now takes the form

$$\begin{aligned} &\hat{V}(t, \omega_{\text{thres}}) \\ &\quad = e^{i\omega_t t}[\Omega_c(t)\hat{n}_c(\omega_t, \omega_{\text{thres}}) + \Omega_t(t)\hat{n}_t(\omega_{\text{thres}}, \omega_t)] \\ &\qquad + e^{-i\omega_t t}[\Omega_c^*(t)\hat{n}_c(-\omega_t, \omega_{\text{thres}}) + \Omega_t^*(t)\hat{n}_t(\omega_{\text{thres}}, -\omega_t)], \end{aligned} \quad (38)$$

where $\Omega_j = \Omega_{jx} + i\Omega_{jy}$ with the modified operators $\hat{n}_{(c,t)}(\omega, \omega_{\text{thres}})$ that are such that matrix elements are set to zero,

$$\hat{n}_{(c,t)}(\omega, \omega_{\text{thres}})[k, k'] \equiv 0, \quad (39)$$

if they connect states $|k\rangle$ and $|k'\rangle$ that are separated in frequency by more than the threshold, $|E_k - (E_{k'} - \omega)| > \omega_{\text{thres}}$. Note that this is equivalent to the usual RWA made on a two-level qubit driven on resonance where $\omega_{\text{thres}} < 2\omega$.

We choose to operate this gate with control and target frequencies of $\omega_c/2\pi = 5.1$ GHz, $\omega_t/2\pi = 4.9$ GHz, and anharmonicities $\alpha_c/2\pi = -355$ MHz and $\alpha_t/2\pi = -352$ MHz, respectively. Moreover, the qubit-qubit coupling is set to $g/2\pi = 4.29$ MHz and the target unitary is taken to be a ZX_{90} gate [20,21]. Additionally, we optimize the pulses for several different choices of threshold frequencies, ω_{thres}, and compute the fidelity of the pulse shapes found in this way using the RWA-free Hamiltonian. Figure 5 shows the infidelity of the cross-resonance gate obtained in this way as a function of the gate time and for different choices of threshold frequencies. We first note that, for a flat pulse (red diamonds), a gate fidelity approaching 99% can be reached. The other data sets in Fig. 5 correspond to different values of ω_{thres} ranging from 200 MHz to 2 GHz (blue symbols, from top to bottom). The darkest blue circles correspond to simulations without RWA and to which the other blue symbols should be compared. The impact of terms oscillating at high frequency can be seen in the results obtained with a ω_{thres} of 200 MHz which coincide with those obtained for an unoptimized flat envelope. Even for $\omega_{\text{thres}}/2\pi = 500$ MHz the results agree with the full simulation only for gate times exceeding 150 ns. To reach agreement with the full simulation for all gate lengths, terms oscillating as fast 1 GHz must be kept. Omitting these fast-oscillating terms can therefore lead to lower-fidelity gates. We stress that, in contrast to the usual solvers, within the framework of `Dysolve`, there is no computational advantage to eliminating any of the off-resonant transitions. In short, `Dysolve` is a numerically efficient approach to optimize pulse envelopes to reach high-fidelity gates.

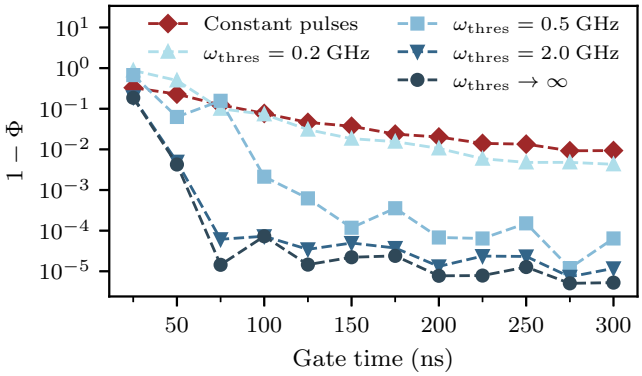

FIG. 5. Infidelity of the cross-resonance gate as a function of the chosen gate time when the gate is operated at the qubit-qubit detuning $\Delta/2\pi = (\omega_c - \omega_t)/2\pi = 210$ MHz. The different curves demonstrate the result of optimization of the pulse shape under RWAs at different threshold frequencies ω_{thres}, followed by a subsequent calculation of the true fidelity in the RWA-free case.

To further demonstrate that GRAPE under `Dysolve` can successfully optimize the cross-resonance gate more generally, Fig. 6 shows the fidelity of an optimized 300-ns gate as a function of the qubit-qubit detuning $\Delta/2\pi = (\omega_c - \omega_t)/2\pi$. In the numerical simulations, this detuning is varied by changing the target qubit frequency while keeping the qubit-qubit coupling and the control-qubit frequency fixed. We note that the performance of the gate at positive detunings is approximately in line with those predicted by Schrieffer-Wolff perturbation theory for an echoed cross-resonance gate performed on a similar device [22]. The results are excluded at $\Delta = 0, \pm\alpha$, as the qubits strongly hybridize due to resonances. The slight variations in the GRAPE fidelity in the wide detuning range are partially attributable to variations in higher-order two-qubit coupling amplitudes across, alongside the performance of gradient ascent for different optimization

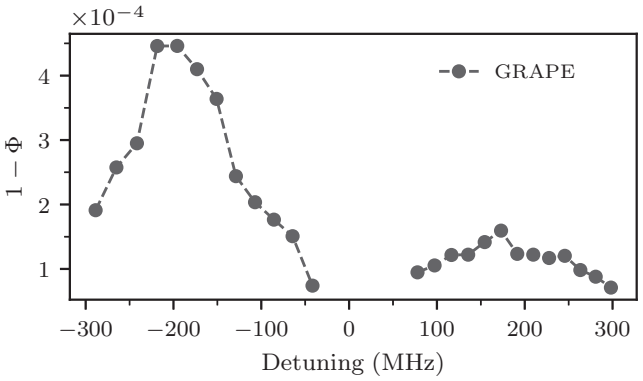

FIG. 6. Infidelity of the cross-resonance gate as a function of the detuning of the qubits when optimized under GRAPE for 5000 iterations. The control qubit frequency is fixed to $\omega_c/2\pi = 5.10$ GHz.

landscapes. This also constitutes evidence for the robustness of our algorithm, as excellent fidelities are obtained in a broad range of parameters.

V. FUTURE WORK AND CONCLUSION

We have demonstrated that the `Dysolve` algorithm provides a means to quickly and accurately simulate driven systems while accounting for all of the effects of the counter-rotating and off-resonant terms. Analysis of ultrafast quantum gates and the quantum speed limit [23], where counter-rotating effects are particularly important, would also be possible with our algorithm. Additionally, this method trivializes the calculation of the gradient, allowing for rapid optimization without the need for additional approximations, and can be modified to include dissipative effects. Indeed, a simple extension of the optimization scheme would allow for optimization of lossy quantum systems specifically, through an open GRAPE-like scheme [24], or the simulation of larger lossy systems through the use of trajectories. We also note that as the expressions depend explicitly on the drive amplitudes, second derivatives and the Hessian matrix can also be calculated exactly, opening the door to second-order optimizers such as Newton's method.

We anticipate that future iterations of the `Dysolve` algorithm improve the efficiency of both the preparation and contraction stages of the algorithm. Given the parallelized nature of this solver, we also envision a direct extension to graphics processing units (GPUs), which could allow for fast simulation of significantly larger systems. We leave this for future work.

Note added. Recently, the authors were made aware of a recent paper [25] which uses a different approach to obtain a result similar to that which we provide in Sec. II. Our derived weighting functions are equivalent to divided differences [26] utilized in their method.

ACKNOWLEDGMENTS

We thank Catherine Leroux for discussions and Éric Giguère for assistance with numerical calculations. This work was undertaken thanks in part to funding from NSERC, the Canada First Research Excellence Fund, and the U.S. Army Research Office Grant No. W911NF-18-1-0411.

APPENDIX A: DERIVATION OF *n*TH-ORDER WEIGHTING FUNCTIONS

In this Appendix, we derive the *n*th-order weighting function and the form of the Dyson series operators. We begin by defining a frequency-dependent set of path operators

$$\hat{R}_{[m_n,\ldots,m_0]}(\boldsymbol{v}_n,\delta t) = (-i)^n \int_0^{\delta t} dt_M \cdots \int_0^{t_2} dt_1 (-i(\hat{H}_0 - \boldsymbol{v}_n[n]))^{m_n} \hat{X} \cdots (-i(\hat{H}_0 - \boldsymbol{v}_n[1]))^{m_1} \hat{X} (-i(\hat{H}_0 - i\boldsymbol{v}_n[0]))^{m_0}, \tag{A1}$$

where, for now, $\boldsymbol{v}_n$ is an arbitrary vector and $M = \sum_{i=0}^{n} m_i + n$. Since there is no explicit time dependence in Eq. (A1), we can evaluate its M integrals simultaneously:

$$\begin{aligned}\hat{R}_{[m_n,\ldots,m_0]}(\boldsymbol{v}_n,\delta t) &= \frac{(-i)^n (-i(\hat{H}_0 - \boldsymbol{v}_n[n]))^{m_n} \hat{X} \cdots (-i(\hat{H}_0 - \boldsymbol{v}_n[1]))^{m_1} \hat{X} (-i(\hat{H}_0 - \boldsymbol{v}_n[0]))^{m_0} (\delta t)^M}{M!} \\ &= (-i\delta t)^n \frac{(-i\delta t(\hat{H}_0 - \boldsymbol{v}_n[n]))^{m_n} \hat{X} \cdots (-i\delta t(\hat{H}_0 - \boldsymbol{v}_n[0]))^{m_0}}{M!}.\end{aligned} \tag{A2}$$

Inserting identity operators $I_N = \sum_k |k\rangle\langle k|$, with $\hat{H}_0|k\rangle = \lambda_k|k\rangle$ and N the size of the system's Hilbert space, between groups of $\hat{H}_0$ and each $\hat{X}$ operator leads to

$$\begin{aligned}\hat{R}_{[m_n,\ldots,m_0]}(\boldsymbol{v}_n,\delta t) = (-i\delta t)^n \sum_{\boldsymbol{k}_n} &\langle k^{(n)}|\hat{X}|k^{(n-1)}\rangle\langle k^{(n-1)}|\hat{X}\ldots|k^{(1)}\rangle\langle k^{(1)}|\hat{X}|k^{(0)}\rangle|k^{(n)}\rangle\langle k^{(0)}| \\ &\times \frac{(-i(\boldsymbol{\lambda}_n(\boldsymbol{k}_n) - \boldsymbol{v}_n)[0]\,\delta t)^{m_0} \cdots (-i(\boldsymbol{\lambda}_n(\boldsymbol{k}_n) - \boldsymbol{v}_n)[n]\,\delta t)^{m_n}}{M!}.\end{aligned} \tag{A3}$$

To obtain the Dyson operator, we sum Eq. (A3) over all $\boldsymbol{m} \in \mathbf{Z}_+^{n+1}$. To this end, we define the nth-order weighting function $f(\boldsymbol{\lambda}_n)$ as

$$\begin{aligned} f(\boldsymbol{\lambda}_n) &= \sum_{m_0,\ldots,m_n} \frac{(-i\boldsymbol{\lambda}_n[0])^{m_0}\cdots(-i\boldsymbol{\lambda}_n[n])^{m_n}}{\left(\sum_{z=0}^{n}(m_z)+n\right)!} \\ &= \sum_{m_0,\ldots,m_{n-2}} (-i\boldsymbol{\lambda}_n[0])^{m_0}\cdots(-i\boldsymbol{\lambda}_n[n-2])^{m_{n-2}} \sum_{P=0}^{\infty}\sum_{p=0}^{P} \frac{(-i\boldsymbol{\lambda}_n[n-1])^{P-p}(-i\boldsymbol{\lambda}_n[n])^{p}}{\left(\sum_{n=0}^{n-2}(m_n)+P-p+p+n\right)!} \\ &= \sum_{m_0,\ldots,m_{n-2},P} \frac{(-i\boldsymbol{\lambda}_n[0])^{m_0}\cdots(-i\boldsymbol{\lambda}_n[n-2])^{m_{n-2}}\left((-i\boldsymbol{\lambda}_n[n-1])^{P+1}-(-i\boldsymbol{\lambda}_n[n])^{P+1}\right)}{\left(\sum_{n=0}^{n-2}(m_n)+P+1+n-1\right)!(-i\boldsymbol{\lambda}_n[n-1]+i\boldsymbol{\lambda}_n[n])} \\ &= i\left[\frac{f(\boldsymbol{g}(\boldsymbol{\lambda}_n))-f(\boldsymbol{g}^2(\boldsymbol{\lambda}_n)\cup\boldsymbol{\lambda}_n[n])}{\boldsymbol{\lambda}_n[n-1]-\boldsymbol{\lambda}_n[n]}\right], \end{aligned} \tag{A4}$$

where $f(\boldsymbol{\lambda}_0)=\exp(-i\boldsymbol{\lambda}_0[0])$. This gives us the form of $\hat{R}^{(n)}(\boldsymbol{v}_n,0,\delta t)$:

$$\hat{R}^{(n)}(\boldsymbol{v}_n,\delta t)=\left(\frac{-i\delta t}{2}\right)^n \sum_{\boldsymbol{k}_n}\langle k^{(n)}|\hat{X}|k^{(n-1)}\rangle\langle k^{(n-1)}|\hat{X}\ldots|k^{(1)}\rangle\langle k^{(1)}|\hat{X}|k^{(0)}\rangle|k^{(n)}\rangle\langle k^{(0)}|\times f((\boldsymbol{\lambda}_n(\boldsymbol{k}_n)-\boldsymbol{v}_n)\,\delta t). \tag{A5}$$

We then note that

$$e^{i\omega t_0}f(\boldsymbol{\lambda}_n)=f(\boldsymbol{\lambda}_n-\omega t_0). \tag{A6}$$

To prove this, first note that this is trivially true for $f(\boldsymbol{\lambda}_0)$. We then make the inductive step, assuming Eq. (A6) to be true for $n-1$. Then,

$$\begin{aligned} e^{i\omega t_0}f(\boldsymbol{\lambda}_n) &= i\left[\frac{e^{i\omega t_0}f(\boldsymbol{g}(\boldsymbol{\lambda}_n))-e^{i\omega t_0}f(\boldsymbol{g}^2(\boldsymbol{\lambda}_n)\cup\boldsymbol{\lambda}_n[n])}{\boldsymbol{\lambda}_n[n-1]-\boldsymbol{\lambda}_n[n]}\right] \\ &= i\left[\frac{f(\boldsymbol{g}(\boldsymbol{\lambda}_n)-\omega_0 t)f(\boldsymbol{g}^2(\boldsymbol{\lambda}_n)\cup\boldsymbol{\lambda}_n[n]-\omega_0 t)}{(\boldsymbol{\lambda}_n[n-1]-\omega_0 t)-(\boldsymbol{\lambda}_n[n]-\omega_0 t)}\right] \\ &= f(\boldsymbol{\lambda}_n-\omega_0 t). \end{aligned} \tag{A7}$$

By linearity, this implies that $e^{ia\delta t}\hat{R}^{(n)}(\boldsymbol{v}_n,\delta t)=\hat{R}^{(n)}(\boldsymbol{v}_n+a,\delta t)$.

We now consider the effect of a set of oscillatory terms. To this end, we first note that the Dyson path operator in Eq. (7), as well as the Dyson operators themselves, can be defined recursively:

$$\hat{S}^{(n)}_{[m_n,\ldots,m_0]}(\boldsymbol{\omega}_n,\delta t)=(-i\hat{H}_0)^{m_n}\hat{X}\int_0^{\delta t}dt_M\cdots\int_0^{t_{M-m_n+1}}dt_{M-m_n}e^{i\boldsymbol{\omega}_n[n]t_{M-m_n}}\hat{S}^{(n-1)}_{[m_{n-1},\ldots,m_0]}(\boldsymbol{g}(\boldsymbol{\omega}_n),t_{M-m_n}), \tag{A8}$$

$$\hat{S}^{(n)}(\boldsymbol{\omega}_n,\delta t)=\frac{1}{2}\sum_{m_n\in\mathbf{Z}_+}(-i\hat{H}_0)^{m_n}\hat{X}\int_0^{\delta t}dt_M\cdots\int_0^{t_{M-m_n+1}}dt_{M-m_n}e^{i\boldsymbol{\omega}_n[n]t_{M-m_n}}\hat{S}^{(n-1)}(\boldsymbol{g}(\boldsymbol{\omega}_n),t_{M-m_n}), \tag{A9}$$

$$\hat{R}^{(n)}(\boldsymbol{v}_n,\delta t)=\frac{1}{2}\sum_{m_n\in\mathbf{Z}_+}(-i\hat{H}_0-i\boldsymbol{v}_n[n])^{m_n}\hat{X}\int_0^{\delta t}dt_M\cdots\int_0^{t_{M-m_n+1}}dt_{M-m_n}\hat{R}^{(n-1)}(\boldsymbol{g}(\boldsymbol{v}_n),t_{M-m_n}). \tag{A10}$$

We now wish to demonstrate that $\hat{R}^{(n)}(\boldsymbol{c}(\boldsymbol{\omega}_n),\delta t)=\hat{S}^{(n)}(\boldsymbol{\omega}_n,\delta t)$, where $\boldsymbol{c}(\boldsymbol{\omega}_n)$ is the cumulative vector introduced in Eq. (16). To first order,

$$\begin{aligned} \hat{S}^{(1)}(\boldsymbol{\omega}_1,\delta t) &= \frac{1}{2}\sum_{m_1\in\mathbf{Z}_+}(-i\hat{H}_0)^{m_1}\hat{X}\int_0^{\delta t}dt_M\cdots\int_0^{t_{M-m_1+1}}dt_{M-m_1}e^{i\boldsymbol{\omega}_1[1]t_{M-m_1}}\hat{S}^{(0)}(\mathbf{0},t_{M-m_1}) \\ &= \frac{1}{2}\sum_{m_1\in\mathbf{Z}_+}(-i\hat{H}_0)^{m_1}\hat{X}\int_0^{\delta t}dt_M\cdots\int_0^{t_{M-m_1+1}}dt_{M-m_1}\left[\sum_i|i\rangle\langle i|e^{i\boldsymbol{\omega}_1[1]t_{M-m_1}}f(\lambda_i t_{M-m_1})\right] \\ &= \frac{1}{2}\sum_{m_1\in\mathbf{Z}_+}(-i\hat{H}_0)^{m_1}\hat{X}\int_0^{\delta t}dt_M\cdots\int_0^{t_{M-m_1+1}}dt_{M-m_1}\left[\sum_i|i\rangle\langle i|f((\lambda_i-\boldsymbol{\omega}_1[1])t_{M-m_1})\right] \\ &= \frac{1}{2}\sum_{m_1\in\mathbf{Z}_+}(-i\hat{H}_0)^{m_1}\hat{X}\int_0^{\delta t}dt_M\cdots\int_0^{t_{M-m_1+1}}dt_{M-m_1}\hat{R}^{(0)}([\boldsymbol{\omega}_1[1]],t_{M-m_1}) \end{aligned}$$

$$
\begin{aligned}
&= \frac{1}{2}\sum_{m_1\in\mathbf{Z}_+}(-i(\hat{H}_0-0))^{m_1}\hat{X}\int_0^{\delta t}dt_M\cdots\int_0^{t_{M-m_1+1}}dt_{M-m_1}\hat{R}^{(0)}(\boldsymbol{g}([\boldsymbol{\omega}_1[1],0]),t_{M-m_1})\\
&= \hat{R}^{(1)}([\boldsymbol{\omega}_1[1],0],\delta t)\\
&= \hat{R}^{(1)}(\boldsymbol{c}(\boldsymbol{\omega}_1),\delta t). \qquad \text{(A11)}
\end{aligned}
$$

We again make the inductive step with the assumption that $\hat{R}^{(n-1)}(\boldsymbol{c}(\boldsymbol{\omega}_n),\delta t)=\hat{S}^{(n-1)}(\boldsymbol{\omega}_n,\delta t)$. To proceed, we first show the following result:

$$
\begin{aligned}
\boldsymbol{c}(\boldsymbol{g}(\boldsymbol{v}_n)) &= \left(\left[\sum_{p=1}^{n-1}\boldsymbol{v}_n[n-p]\right],\left[\sum_{p=1}^{n-2}\boldsymbol{v}_n[n-p]\right],\ldots,\boldsymbol{v}_n[n-1],0\right).\\
&= \left(\left[\sum_{p=0}^{n-1}\boldsymbol{v}_n[n-p]\right],\left[\sum_{p=0}^{n-2}\boldsymbol{v}_n[n-p]\right],\ldots,\boldsymbol{v}_n[n-1]+\boldsymbol{v}_n[n],\boldsymbol{v}_n[n]\right)-\boldsymbol{v}_n[n]\\
&= \boldsymbol{g}\left(\left[\sum_{p=0}^{n-1}\boldsymbol{v}_n[n-p]\right],\left[\sum_{p=0}^{n-2}\boldsymbol{v}_n[n-p]\right],\ldots,\boldsymbol{v}_n[n-1]+\boldsymbol{v}_n[n],\boldsymbol{v}_n[n],0\right)-\boldsymbol{v}_n[n]\\
&= \boldsymbol{g}(\boldsymbol{c}(\boldsymbol{v}_n))-\boldsymbol{v}_n[n]. \qquad \text{(A12)}
\end{aligned}
$$

Then,

$$
\begin{aligned}
\hat{S}^{(n)}(\boldsymbol{\omega}_n,\delta t) &= \frac{1}{2}\sum_{m_n\in\mathbf{Z}_+}(-i\hat{H}_0)^{m_n}\hat{X}\int_0^{\delta t}dt_M\cdots\int_0^{t_{M-m_n+1}}dt_{M-m_n}e^{i\boldsymbol{\omega}_n[n]t_{M-m_n}}\hat{S}^{(n-1)}(\boldsymbol{g}(\boldsymbol{\omega}_n),t_{M-m_n})\\
&= \frac{1}{2}\sum_{m_n\in\mathbf{Z}_+}(-i\hat{H}_0)^{m_n}\hat{X}\int_0^{\delta t}dt_M\cdots\int_0^{t_{M-m_n+1}}dt_{M-m_n}e^{i\boldsymbol{\omega}_n[n]t_{M-m_n}}\hat{R}^{(n-1)}(\boldsymbol{c}(\boldsymbol{g}(\boldsymbol{\omega}_n)),t_{M-m_n})\\
&= \frac{1}{2}\sum_{m_n\in\mathbf{Z}_+}(-i\hat{H}_0)^{m_n}\hat{X}\int_0^{\delta t}dt_M\cdots\int_0^{t_{M-m_n+1}}dt_{M-m_n}\hat{R}^{(n-1)}(\boldsymbol{c}(\boldsymbol{g}(\boldsymbol{\omega}_n))+\boldsymbol{\omega}_n[n],t_{M-m_n})\\
&= \frac{1}{2}\sum_{m_n\in\mathbf{Z}_+}(-i\hat{H}_0-i0)^{m_n}\hat{X}\int_0^{\delta t}dt_M\cdots\int_0^{t_{M-m_n+1}}dt_{M-m_n}\hat{R}^{(n-1)}(\boldsymbol{g}(\boldsymbol{c}(\boldsymbol{\omega}_n)),t_{M-m_n})\\
&= \hat{R}^{(n)}(\boldsymbol{c}(\boldsymbol{\omega}_n),\delta t), \qquad \text{(A13)}
\end{aligned}
$$

finishing the induction and giving the final desired form in Eq. (14).

APPENDIX B: MULTIPLE DRIVE INPUTS

Suppose there are q-independent drive inputs each with its own drive frequency and operator. For an nth-order Dyson expansion, we define an n-dimensional vector $\boldsymbol{\beta}$ with values ranging from 1 to q, each value referring to one of the drive inputs. To simplify the notation, we drop all subscripts n, which are implied. The new Dyson series operators corresponding to a particular vector $\boldsymbol{\beta}$ are

$$
\begin{aligned}
\hat{S}^{(n)}_{\boldsymbol{\beta}}(\boldsymbol{\omega}_{\boldsymbol{\beta}},\delta t) = \sum_{\boldsymbol{k}_n}&\langle k^{(n)}|\hat{X}_{\boldsymbol{\beta}[n]}|k^{(n-1)}\rangle\langle k^{(n-1)}|\cdots\langle k^{(1)}|\hat{X}_{\boldsymbol{\beta}[1]}|k^{(0)}\rangle\\
&\times\ |k^{(n)}\rangle\langle k^{(0)}|f(\boldsymbol{\lambda}-\delta t\,\boldsymbol{v}_{\boldsymbol{\beta}}), \qquad \text{(B1)}
\end{aligned}
$$

with $\boldsymbol{\omega}_{\boldsymbol{\beta}}$ an n-vector where the pth element of this vector is $\pm\omega_{\boldsymbol{\beta}[p]}$, or more simply, plus or minus the drive frequency corresponding to the $\boldsymbol{\beta}[p]$th input. Similarly, $\hat{X}_{\boldsymbol{\beta}[p]}$ refers to the $\boldsymbol{\beta}[p]$th input operator. The definition of $v_{\boldsymbol{\beta}}$ remains unchanged from Eq. (16), albeit with the new drive frequency vector $\boldsymbol{\omega}_{\boldsymbol{\beta}}$.

The drive function is also modified and now reads

$$
\Omega(\boldsymbol{\omega}_{\boldsymbol{\beta}}) = \prod_{p=1}^{n}\boldsymbol{\Omega}^{\mu}_{\boldsymbol{\beta}[p]}\boldsymbol{\Omega}^{*(1-\mu)}_{\boldsymbol{\beta}[p]}, \qquad \text{(B2)}
$$

where

$$
\mu = \tfrac{1}{2}(1+\mathrm{sign}\{\boldsymbol{\omega}_{\boldsymbol{\beta}}[p]\}), \qquad \text{(B3)}
$$

thus allowing us to define our generic nth-order Dyson series,

$$
\hat{U}^{(n)}(t,t+\delta t)=\sum_{\boldsymbol{\beta}}\sum_{\boldsymbol{\omega}_{\boldsymbol{\beta}}}\exp\left(i\sum_{p=1}^{n}\boldsymbol{\omega}_{\boldsymbol{\beta}}[p]t\right)\Omega(\boldsymbol{\omega}_{\boldsymbol{\beta}})\hat{S}^{(n)}_{\boldsymbol{\beta}}(\boldsymbol{\omega}_{\boldsymbol{\beta}},\delta t), \qquad \text{(B4)}
$$

which reads as a version of Eq. (22) with multiple drive operators and frequencies.

APPENDIX C: LINEAR INTERPOLATION OF SUBPIXELS

The accuracy of the Dyson expansion is insensitive to the drive frequency ω—rather, it depends only on the number of subpixels and the drive amplitudes. However, if each subpixel

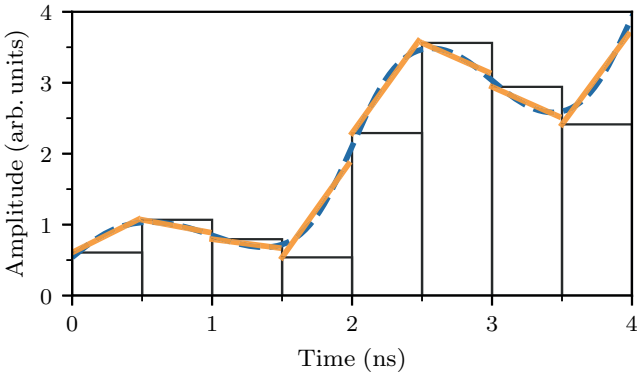

FIG. 7. Linear interpolation of the pulse sequences (solid yellow lines) with two subpixels per pixel. The black bars indicate the original Gaussian filtering, with the blue dashed line the "true" filtered pulse.

assumes a constant amplitude, a leading error can be caused by the change in amplitude of the drive over a single subpixel. In Fig. 7, we illustrate this issue, where the subpixel amplitudes (black bars) will over- or underestimate the true pulse amplitude, depending on the gradient of the envelope function. To circumvent this issue, we consider a new linear interpolation of the drive envelope over a single subpixel. We define $y_l(t)$ as the new time-dependent amplitude for the lth subpixel:

$$y_l(t) = s'_l + \tfrac{s_{l+1}-s_l}{\delta t}(t - l\delta t), \quad l\delta t < t < (l+1)\delta t, \tag{C1}$$

where s'_l is calculated from a modified filter function to ensure that the integral of the subpixel matches that of the continuous pulse. The $y_l(t)$ time-dependent subpixels are shown in yellow, and well approximate the true pulse even for relatively few subpixels. To determine the impact of the linear time term t_0 on the time-ordered integral, we consider the $(n-j)$th iteration of a Dyson series operator,

$$\begin{aligned}
&\int_0^{\delta t} dt_M \cdots \int_0^{t_{M-m_{(j+1)}}} dt_{M-m_j} t_{M-m_j} \exp(i\boldsymbol{\omega}_n[j] t_{M-m_j}) \\
&\quad \times \hat{S}^{(j-1)}(\boldsymbol{g}^{(j-1)}(\boldsymbol{\omega}_n), t_{M-m_j}), \\
&= \int_0^{\delta t} dt_M \cdots \int_0^{t_{M-m_{(j+1)}}} dt_{M-m_j} (-i\partial_{\boldsymbol{\omega}_n,j}) \exp(i\boldsymbol{\omega}_n[j] t_{M-m_j}) \\
&\quad \times \hat{S}^{(j-1)}(\boldsymbol{g}^{(j-1)}(\boldsymbol{\omega}_n), t_{M-m_j}), \\
&= (-i\partial_{\boldsymbol{\omega}_n,j}) \int_0^{\delta t} dt_M \cdots \int_0^{t_{M-m_{(j+1)}}} dt_{M-m_j} \exp(i\boldsymbol{\omega}_n[j] t_{M-m_j}) \\
&\quad \times \hat{S}^{(j-1)}(\boldsymbol{g}^{(j-1)}(\boldsymbol{\omega}_n), t_{M-m_j}), \\
&= (-i\partial_{\boldsymbol{\omega}_n,j}) \hat{S}^{(j)}(\boldsymbol{g}^{(j)}(\boldsymbol{\omega}_n), t_{M-m_{j+1}}),
\end{aligned} \tag{C2}$$

where we define $\partial_{\boldsymbol{\omega}_n,j}$ as the derivative with respect to the jth component of the vector $\boldsymbol{\omega}_n$. For example,

$$\partial_{\boldsymbol{\omega}_n,j}(\boldsymbol{\omega}_n) = [0, 0, \ldots, \underbrace{1}_{j\text{th}}, \ldots, 0]. \tag{C3}$$

This result allows us to modify Eq. (29) to the following result, replacing the drive amplitude function by a drive amplitude operator:

$$\begin{aligned}
&\hat{\Omega}_l(\boldsymbol{\omega}_n) = \prod_{p=1}^{n} \left(s_l - i s'_l \partial_{\boldsymbol{\omega}_n,p}\right)^{\mu} \left(s_l^* + i s_l'^* \partial_{\boldsymbol{\omega}_n,p}\right)^{1-\mu}, \\
&\text{where } \mu = \frac{1}{2}(1 + \text{sign}\{\boldsymbol{\omega}[p]\}),
\end{aligned} \tag{C4}$$

where the derivatives are to act upon the Dyson series operators. It is important to note that $s'_l \ll s_l$ in the majority of cases; as such, it is possible to consider a partial truncation of the series, where only up to a certain power of derivative terms s'_l are included in the set of Dyson series operators.

To make use of the drive amplitude function, we must determine the derivatives of the weighting function $f(\boldsymbol{\lambda}_n)$. We

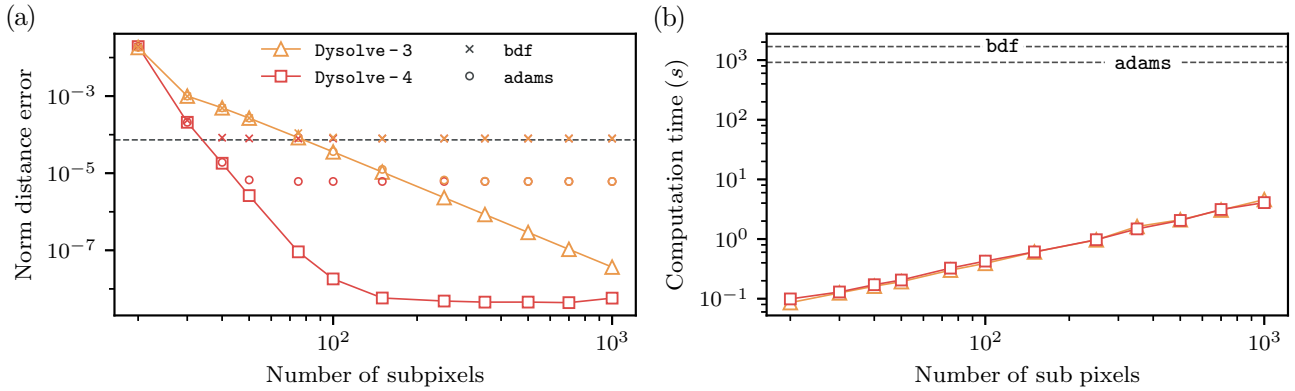

FIG. 8. Comparison of the `Dysolve` algorithm and QUTIP's `propagator` over an evolution of 20 Rabi oscillations, with the same parameters as in Fig. 3 of the main text. (a) Frobenius norm distance between the `Dysolve` algorithm with different subpixel numbers and QUTIP's `propagator` with `adams` (circles) and `bdf` (x's) methods. The squares are the norm distance between `Dysolve-n` and `Dysolve-4` with 10^4 subpixels. The dashed line indicates the error between `propagator` with the `adams` and `bdf` methods. (b) Computational time of `Dysolve` and `propagator` for results of the left panel. The dashed lines refer to the computational time for `propagator` with `adams` and `bdf` with the high-precision settings described in Appendix D.

begin with the base case with $\boldsymbol{\lambda}_0 = [\lambda_0]$,

$$\frac{d}{d\lambda_0} f(\boldsymbol{\lambda}_0) = \frac{d}{d\lambda_0} e^{-i\lambda_0} = -ie^{-i\lambda_0} = \lim_{\epsilon \to 0} \frac{e^{-i(\lambda_0+\epsilon)} - e^{-i\lambda_0}}{\epsilon} = if(\lambda_0 \cup \boldsymbol{\lambda}_0). \quad \text{(C5)}$$

Now, we assume that $\partial_{\boldsymbol{\lambda}_n,j} f(\boldsymbol{\lambda}_n) = if(\boldsymbol{\lambda}_n[j] \cup \boldsymbol{\lambda}_n)$. Then,

$$\begin{aligned}
\partial_{\boldsymbol{\lambda}_n,j} f(\boldsymbol{\lambda}_n) &= i\partial_{\boldsymbol{\lambda}_n,j} \frac{f(g(\boldsymbol{\lambda}_n)) - f(g^2(\boldsymbol{\lambda}_n) \cup \boldsymbol{\lambda}_n[n])}{\boldsymbol{\lambda}_n[n-1] - \boldsymbol{\lambda}_n[n]} \\
&= \frac{i}{\boldsymbol{\lambda}_n[n-1] - \boldsymbol{\lambda}_n[n]} \\
&\quad \times \left[\partial_{\boldsymbol{\lambda}_n,j} f(g(\boldsymbol{\lambda}_n)) - \partial_{\boldsymbol{\lambda}_n,j} f(g^2(\boldsymbol{\lambda}_n) \cup \boldsymbol{\lambda}_n[n])\right] \\
&= \frac{i}{\boldsymbol{\lambda}_n[n-1] - \boldsymbol{\lambda}_n[n]} [if(g(\boldsymbol{\lambda}_n[j] \cup \boldsymbol{\lambda}_n)) \\
&\quad - if(g^2(\boldsymbol{\lambda}_n[j] \cup \boldsymbol{\lambda}_n) \cup \boldsymbol{\lambda}_n[n])] \\
&= if(\boldsymbol{\lambda}_n[j] \cup \boldsymbol{\lambda}_n). \quad \text{(C6)}
\end{aligned}$$

APPENDIX D: BENCHMARK CALCULATIONS

In order to calculate the Frobenius norm error used as a metric in our benchmarks, we required an excellent approximation to the reference propagator operator $\hat{U}_{\text{ref}}(0, T)$. We first considered an alternative numerical solver, `propagator` from the PYTHON package QUTIP. We selected highly accurate settings for this computation, namely, absolute and relative tolerances of 10^{-16}, two different numerical methods with their maximal order (12 for `adams` and 5 for `bdf`), discretizing each pixel into 10^4 subpixels and allowing for 10^{17} internal substeps. After averaging over many simulations of $T = 500$ ns with a single drive, it became apparent that the `propagator` function was converging and unable to return solutions with a Frobenius norm error less than 5×10^{-6}, as seen in Fig. 8. The shortcoming of this numerical method was confirmed by the fact that the norm distance between `adams` and `bdf` was found to be significantly greater than the norm distance between `adams` and the `Dysolve` algorithm for a sufficient subpixel number. This is in addition to a computational time at least two orders of magnitude larger with the standard QUTIP approach as seen in Fig. 8(b).

[1] J. C. Tremblay and T. Carrington, Using preconditioned adaptive step size Runge-Kutta methods for solving the time-dependent Schrödinger equation, J. Chem. Phys. **121**, 11535 (2004).

[2] C. Moler and C. Loan, Nineteen dubious ways to compute the exponential of a matrix, SIAM Rev. **20**, 801 (1978).

[3] F. Bloch and A. Siegert, Magnetic resonance for nonrotating fields, Phys. Rev. **57**, 522 (1940).

[4] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, Optimal control of coupled spin dynamics: Design of NMR pulse sequences by gradient ascent algorithms, J. Magn. Reson. **172**, 296 (2005).

[5] E. Barnes and S. Das Sarma, Analytically Solvable Driven Time-Dependent Two-Level Quantum Systems, Phys. Rev. Lett. **109**, 060401 (2012).

[6] Y. Wu and X. Yang, Strong-Coupling Theory of Periodically Driven Two-Level Systems, Phys. Rev. Lett. **98**, 013601 (2007).

[7] F. Motzoi, J. M. Gambetta, P. Rebentrost, and F. K. Wilhelm, Simple Pulses for Elimination of Leakage in Weakly Nonlinear Qubits, Phys. Rev. Lett. **103**, 110501 (2009).

[8] F. Motzoi, J. M. Gambetta, S. T. Merkel, and F. K. Wilhelm, Optimal control methods for rapidly time-varying Hamiltonians, Phys. Rev. A **84**, 022307 (2011).

[9] J. Johansson, P. Nation, and F. Nori, QUTIP 2: A PYTHON framework for the dynamics of open quantum systems, Comput. Phys. Commun. **184**, 1234 (2013).

[10] 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, Phys. Rev. A **76**, 042319 (2007).

[11] L. Armijo, Minimization of functions having Lipschitz continuous first partial derivatives, Pac. J. Math. **16**, 1 (1966).

[12] P. A. Absil, R. Mahony, and B. Andrews, Convergence of the iterates of descent methods for analytic cost functions, SIAM J. Optim. **16**, 531 (2005).

[13] C. Rigetti and M. Devoret, Fully microwave-tunable universal gates in superconducting qubits with linear couplings and fixed transition frequencies, Phys. Rev. B **81**, 134507 (2010).

[14] J. M. Chow, A. D. Córcoles, J. M. Gambetta, C. Rigetti, B. R. Johnson, J. A. Smolin, J. R. Rozen, G. A. Keefe, M. B. Rothwell, M. B. Ketchen, and M. Steffen, Simple All-Microwave Entangling Gate for Fixed-Frequency Superconducting Qubits, Phys. Rev. Lett. **107**, 080502 (2011).

[15] J. L. Allen, R. Kosut, J. Joo, P. Leek, and E. Ginossar, Optimal control of two qubits via a single cavity drive in circuit quantum electrodynamics, Phys. Rev. A **95**, 042325 (2017).

[16] S. Kirchhoff, T. Keßler, P. J. Liebermann, E. Assémat, S. Machnes, F. Motzoi, and F. K. Wilhelm, Optimized cross-resonance gate for coupled transmon systems, Phys. Rev. A **97**, 042348 (2018).

[17] M. H. Devoret *et al.*, in Quantum Fluctuations in Electrical Circuits, *Les Houches Summer School of Theoretical Physics, LXIII* (North-Holland, Amsterdam, 1995), Vol. 7.

[18] S. Sheldon, E. Magesan, J. M. Chow, and J. M. Gambetta, Procedure for systematically tuning up cross-talk in the cross-resonance gate, Phys. Rev. A **93**, 060302(R) (2016).

[19] A. D. Patterson, J. Rahamim, T. Tsunoda, P. A. Spring, S. Jebari, K. Ratter, M. Mergenthaler, G. Tancredi, B. Vlastakis, M. Esposito, and P. J. Leek, Calibration of a Cross-Resonance Two-Qubit Gate Between Directly Coupled Transmons, Phys. Rev. Appl. **12**, 064013 (2019).

[20] J. M. Chow, J. M. Gambetta, A. D. Corcoles, S. T. Merkel, J. A. Smolin, C. Rigetti, S. Poletto, G. A. Keefe, M. B. Rothwell, J. R. Rozen *et al.*, Universal Quantum Gate Set Approaching Fault-Tolerant Thresholds with Superconducting Qubits, Phys. Rev. Lett. **109**, 060501 (2012).

[21] A. D. Córcoles, J. M. Gambetta, J. M. Chow, J. A. Smolin, M. Ware, J. Strand, B. L. T. Plourde, and M. Steffen, Process

verification of two-qubit quantum gates by randomized benchmarking, Phys. Rev. A **87**, 030301(R) (2013).

[22] M. Malekakhlagh, E. Magesan, and D. C. McKay, First-principles analysis of cross-resonance gate operation, Phys. Rev. A **102**, 042605 (2020).

[23] Y. Shao, B. Liu, M. Zhang, H. Yuan, and J. Liu, Operational definition of a quantum speed limit, Phys. Rev. Research **2**, 023299 (2020).

[24] M. Abdelhafez, D. I. Schuster, and J. Koch, Gradient-based optimal control of open quantum systems using quantum trajectories and automatic differentiation, Phys. Rev. A **99**, 052327 (2019).

[25] A. Kalev and I. Hen, An integral-free representation of the Dyson series using divided differences, arXiv:2010.09888.

[26] C. de Boor, Divided differences, Surv. Approximation Theory **1**, 46 (2005).

Chapter 4

Master Equation Integrator for Driven Time-Dependent Systems

"I've decided to get into shape, and the shape I've selected is a triangle."

– Howie Mandel

In Chapter 3, I demonstrated that the Dysolve algorithm offers a significant improvement over standard numerical integrators for time-dependent unitary evolution. However, some of the most complex quantum simulations require modelling a master equation, such as the readout master equation introduced in Chapter 2 whereby excitations can exit the system via a 'bath'. As we will see in Sect. 4.1, it is possible to apply Dysolve directly to a master equation, but at great expense. In this chapter, we propose a method of approximating the master equation by evaluating sets of Volterra integrals, which results in a significant speed-up over Qutip's Mesolve integrator whilst constraining memory requirements.

4.1 Master Equation

In Chapter 2, I introduced some of the concepts required to simulate a master equation. The goal here is to see how Dysolve could be extended to tackle this more complex class of equations. We begin with a simple master equation

$$\dot{\hat{\rho}} = -i[\hat{H}_0 + \hat{X}\sin(\omega_d t), \hat{\rho}] + \kappa\mathcal{D}[\hat{L}]\hat{\rho}, \tag{4.1}$$

where $\hat{L}$ is a collapse operator, $\hat{H}_0$ is a diagonal Hamiltonian and $\hat{X}$ is some drive operator with corresponding drive frequency ω_d acting in a Hilbert space of size N. Like in Chapter 3, we choose the simplest case to demonstrate the algorithm – it can naturally be generalized to multiple loss operators, which we'll consider in Sect. 4.1.3.

To facilitate numerical simulations, we rewrite this master equation using a non-Hermitian Hamiltonian:

$$\dot{\hat{\rho}} = -i[\hat{\mathcal{H}}(t)\hat{\rho} - \hat{\rho}\hat{\mathcal{H}}(t)^\dagger] + \kappa\hat{L}\hat{\rho}\hat{L}^\dagger, \tag{4.2}$$

where the non-Hermitian Hamiltonian is defined as

$$\hat{\mathcal{H}}(t) = \hat{H}_0 + \hat{X}\sin(\omega_d t) - i\frac{\kappa}{2}\hat{L}^\dagger\hat{L}. \tag{4.3}$$

Let us briefly imagine that the 'jump' term $\kappa\hat{L}\hat{\rho}\hat{L}^\dagger$ wasn't present in Eq. (4.2). In this case, we can simply evolve the density matrix using a propagator: $\hat{\rho}(t) = \hat{\mathcal{V}}(t)\hat{\rho}_0\hat{\mathcal{V}}^\dagger(t)$, where

$$\hat{\mathcal{V}}(t) = \mathcal{T}\exp\left(-i\int_0^t \hat{\mathcal{H}}(t')dt'\right) \tag{4.4}$$

is a time-ordered non-Hermitian matrix exponential which can be easily evaluated by the methods detailed in Chapter 3. Unfortunately, the jump term $\kappa\hat{L}\hat{\rho}\hat{L}^\dagger$ means that we cannot rely on the non-Hermitian propagator in Eq. (4.4) alone to solve the master equation evolution.

4.1.1 Dysolve – Superoperator notation

There is a rather straightforward method to make Eq. (4.1) compatible with the Dysolve algorithm – to calculate the matrix corresponding to the Lindblad superoperator $\mathcal{L}$ [79]. To do so, we begin by employing superoperator notation and rewrite Eq. (4.2) as

$$\mathcal{L}\hat{\rho} = \left(-i(\hat{\mathcal{H}}(t)\bullet\hat{I} - \hat{I}\bullet\hat{\mathcal{H}}(t)^\dagger) + \kappa\hat{L}\bullet\hat{L}^\dagger\right)\hat{\rho}, \tag{4.5}$$

where we use the notation $(\hat{A} \bullet \hat{B})\hat{\rho} = \hat{A}\hat{\rho}\hat{B}$, and $\hat{I}$ is simply the identity matrix. Next, we note that Eq. (4.5) can be mapped to a matrix-vector form, whereby the $N \times N$ matrix $\hat{\rho}$ is rearranged to form an $N^2 \times 1$ vector, and $\mathcal{L}$ is constructed as an $N^2 \times N^2$ matrix by employing the Kronecker product:

$$\hat{A} \bullet \hat{B} \equiv \hat{A} \otimes \hat{B}^T. \tag{4.6}$$

See Ref. [79] for more details on this procedure. As a result of this mapping, we arrive at a matrix-vector equation of the form

$$\begin{aligned} \dot{\hat{\rho}} = \Big[& -i(\hat{H}_0 \otimes \hat{I} - \hat{I} \otimes \hat{H}_0) - i(\hat{X} \otimes \hat{I} - \hat{I} \otimes \hat{X}) \sin(\omega_d t) \\ & + \kappa \hat{L} \otimes \hat{L}^{\dagger T} - \frac{\kappa}{2} \hat{L}^\dagger \hat{L} \otimes \hat{I} - \frac{\kappa}{2} \hat{I} \otimes (\hat{L}^\dagger \hat{L})^T \Big] \hat{\rho}. \end{aligned} \tag{4.7}$$

Here we can directly apply the Dysolve algorithm, defining the required eigenvalues from the matrix $(\hat{H}_0 \otimes \hat{I} - \hat{I} \otimes \hat{H}_0)$ and treating every other operator as a perturbative drive term. However, the size of these superoperators scales as N^4, severely limiting the usefulness of this approach. To put this into context, a density matrix with Hilbert space size of $N = 1000$, assuming complex double precision, requires about 15 MB of memory – the corresponding Lindbladian would require more than 14 TB of memory to store, putting such a simulation out of reach of the vast majority of computers. As a result, a general-purpose solver cannot rely on the above superoperator forms.

4.1.2 A Dyson Series of Dyson Series?

As we saw in Chapter 3, using the Dyson series as a solver has excellent convergence properties if the drive amplitudes are sufficiently small and the order in the series is adequately large. Superconducting qubits generally have frequencies on the order of GHz and control field amplitudes on the order of tens of MHz, making them only a perturbation to otherwise unitary dynamics. The loss rates in these circuits are also generally small, with a typical readout resonator having a loss rate $\kappa/2\pi$ on the order of tens of MHz, with qubit intrinsic loss as low as KHz or less. Inspired by this, we can construct a Dyson series in the 'jump' superoperator term, $\kappa \hat{L} \bullet \hat{L}^\dagger$. To do so, we first choose to enter the *non-unitary* frame defined by the matrix exponential in Eq. (4.4). Given that the Hamiltonian is non-Hermitian, care must be taken using the identities: in particular, $\hat{\mathcal{V}}(t,0)\hat{\mathcal{V}}^\dagger(t,0) \neq I$, but $\hat{\mathcal{V}}(t,0)\hat{\mathcal{V}}(0,t) = I$.

It is further useful to note the following identities:

$$\begin{aligned}\dot{\hat{\mathcal{V}}}(t,0) &= -i\hat{\mathcal{H}}\hat{\mathcal{V}}(t,0), \quad \dot{\hat{\mathcal{V}}}(0,t) = i\hat{\mathcal{V}}(0,t)\hat{\mathcal{H}}, \\ \dot{\hat{\mathcal{V}}}^\dagger(t,0) &= i\hat{\mathcal{V}}^\dagger(t,0)\hat{\mathcal{H}}^\dagger, \quad \dot{\hat{\mathcal{V}}}^\dagger(0,t) = -i\hat{\mathcal{H}}^\dagger\hat{\mathcal{V}}^\dagger(0,t).\end{aligned} \tag{4.8}$$

Defining the density matrix in the interaction picture $\hat{\rho}_I(t) = \hat{\mathcal{V}}(0,t)\hat{\rho}(t)\hat{\mathcal{V}}^\dagger(0,t)$ and taking the time derivative, we obtain

$$\begin{aligned}\dot{\hat{\rho}}_I &= i\hat{\mathcal{V}}(0,t)\hat{\mathcal{H}}\hat{\rho}(t)\hat{\mathcal{V}}^\dagger(0,t) - i\hat{\mathcal{V}}(0,t)\hat{\rho}(t)\hat{\mathcal{H}}^\dagger\hat{\mathcal{V}}^\dagger(0,t) \\ &\quad + \hat{\mathcal{V}}(0,t)\left(-i[\hat{\mathcal{H}}(t)\hat{\rho}(t) - \hat{\rho}(t)\hat{\mathcal{H}}(t)^\dagger] + \kappa\hat{L}\hat{\rho}(t)\hat{L}^\dagger\right)\hat{\mathcal{V}}^\dagger(0,t) \\ &= \kappa\hat{\mathcal{V}}(0,t)\hat{L}\hat{\rho}(t)\hat{L}^\dagger\hat{\mathcal{V}}^\dagger(0,t) \\ &= \kappa\hat{L}_I(t,0)\hat{\rho}_I\hat{L}_I^\dagger(t,0),\end{aligned} \tag{4.9}$$

where we define $\hat{L}_I(t,0) = \hat{\mathcal{V}}(0,t)\hat{L}\hat{\mathcal{V}}(t,0)$. As in the original definition of the Dyson series [73], we can define the density matrix self-consistently:

$$\begin{aligned}\hat{\rho}_I(T) &= \hat{\rho}_0 + \kappa\int_0^T \hat{L}_I(t',0)\hat{\rho}_I(t')\hat{L}_I^\dagger(t',0)^\dagger dt' \\ &= \hat{\rho}_0 + \kappa\int_0^T \hat{L}_I(t',0)\left(\hat{\rho}_0 + \kappa\int_0^{t'} \hat{L}_I(t'',0)\hat{\rho}_I(t'')\hat{L}_I^\dagger(t'',0)^\dagger dt''\right)\hat{L}_I^\dagger(t',0)^\dagger dt' \\ &\quad + ...\end{aligned} \tag{4.10}$$

Finally, we can revert to the original frame, $\hat{\rho}(T) = \hat{\mathcal{V}}(T,0)\hat{\rho}_I(T)\hat{\mathcal{V}}^\dagger(T,0)$, in which the trace of the density matrix is preserved. This yields a sum of operators:

$$\hat{\rho}(T) = \sum_{i=0}^{\infty}\hat{\rho}_i(T), \tag{4.11}$$

where

$$\begin{aligned}\hat{\rho}_0(T) &= \hat{\mathcal{V}}(T,0)\hat{\rho}(0)\hat{\mathcal{V}}^\dagger(T,0), \\ \hat{\rho}_1(T) &= \kappa\int_0^T \hat{\mathcal{V}}(T,t)\hat{L}\mathcal{V}(t,0)\hat{\rho}(0)\hat{\mathcal{V}}^\dagger(0,t)\hat{L}^\dagger\hat{\mathcal{V}}^\dagger(T,t)dt, \\ \hat{\rho}_2(T) &= \kappa^2\int_0^T\int_0^{t_1} \hat{\mathcal{V}}(T,t_1)\hat{L}\hat{\mathcal{V}}(t_1,t)\hat{L}\mathcal{V}(t)\hat{\rho}(0)\hat{\mathcal{V}}^\dagger(t,0)\hat{L}^\dagger\hat{\mathcal{V}}^\dagger(t_1,t)L^\dagger\hat{\mathcal{V}}^\dagger(T,t_1)dtdt_1,\end{aligned} \tag{4.12}$$

and so on. Here, each contribution $\hat{\rho}_n$ corresponds to an evolution of the original density matrix $\hat{\rho}$ which undergoes n 'jumps', or applications of the jump superoperator $\hat{L}\bullet\hat{L}^\dagger$ during the time interval $[0,T]$. This is closely related to the technique of using jump trajectories to approximate the master equation [54, 80], which evolves a quantum state with the non-

Hermitian propagator $\hat{\mathcal{V}}(t)$, and uses Monte-Carlo sampling to determine when jumps occur. I'd like to stress that the formulation in Eq. (4.12) differentiates itself from Monte-Carlo sampling since each $\hat{\rho}_n(T)$ corresponds to *all* possible trajectories with n jumps during this interval. To make this more concrete, consider the evolution of a Monte-Carlo trajectory from time 0 to T which contains two jumps, at times t and t'. This trajectory, along with all other possible two-jump trajectories occurring at times $0 \leq t \leq t' \leq T$, is contained in $\hat{\rho}_2(T)$. As the error of Monte-Carlo integration scales as $1/\sqrt{M}$, where M is the number of trajectories, evaluating each integral directly in Eq. (4.12) shows great promise in reducing the computational complexity of integrating a master equation.

Given that the decay rate κ is small, we anticipate the contribution to the series of the $(n+1)$-th order to be significantly smaller than the n-th order, just like the Dyson series used in Chapter 3. Moreover, this form naturally tells us about the relative probabilities of having several jumps in a given evolution:

$$p_n = \text{Tr}\left[\hat{\rho}_n(T)\right] : \text{The probability of } n \text{ jumps in the interval } [0, T]. \tag{4.13}$$

Consequently, the trace can be used as a useful metric for the convergence of this series, as we can choose to truncate the series at nth order when the trace loss of the approximated density matrix $(1 - \sum_i^n p_i)$ is negligible. This is the goal of this method – to evaluate up to the nth order in Eq. (4.12) such that we have an excellent approximation to the true density matrix.

Given that the equations in Eq. (4.12) are Volterra integral equations [81], and to differentiate this method from the Dysolve algorithm, I will refer to this method as 'Voltsolve'.

4.1.3 Superoperator Notation

Before proceeding to demonstrate how each of the integrals in Eq. (4.12) can be calculated, it is useful to simplify the notation. We switch to superoperator notation, keeping in mind that these superoperators are never explicitly calculated for reasons listed in Sect. 4.1.1, and also generalize to allow for an arbitrary number of loss operators $\hat{L}_i$. In this new notation,

$$\begin{aligned} \hat{\rho}(T) &= \sum_{i=0}^{\infty} \hat{\rho}_i(T), \\ \hat{\rho}_i(T) &= \mathcal{M}_i \hat{\rho}_0, \end{aligned} \tag{4.14}$$

where

$$\begin{aligned}\mathcal{M}_0 &= \mathcal{E}(T,0),\\ \mathcal{E}(T,0) &= \hat{\mathcal{V}}(T,0)\bullet\hat{\mathcal{V}}^\dagger(T,0),\\ \hat{\mathcal{V}}(T,0) &= \mathcal{T}\exp\left(-i\int_0^T \hat{H}(t) - \sum_j \frac{\kappa_j}{2}\hat{L}_j^\dagger\hat{L}_j dt\right).\end{aligned} \tag{4.15}$$

We also define $\mathcal{S} = \sum_j \kappa_j \hat{L}_j \bullet \hat{L}_j^\dagger$. Then,

$$\mathcal{M}_1 = \int_0^T \mathcal{E}(T,t)\mathcal{S}\mathcal{E}(t,0)dt, \tag{4.16}$$

$$\mathcal{M}_2 = \int_0^T\int_0^{t_1} \mathcal{E}(T,t_1)\mathcal{S}\mathcal{E}(t_1,t)\mathcal{S}\mathcal{E}(t,0)dtdt_1, \tag{4.17}$$

and so on, making the evolution captured in the integrals much easier to parse.

Here, I wish to emphasize that each set of collapse superoperators $\mathcal{S}$ captures all possible combinations of jumps at different times – for example, let us consider two collapse operators, $\hat{L}_1$ and $\hat{L}_2$, with corresponding decay rates κ_1 and κ_2. Then, we can expand $\mathcal{M}_2$ to yield

$$\begin{aligned}\mathcal{M}_2 &= \int_0^T\int_0^{t_1} \mathcal{E}(T,t_1)(\kappa_1\hat{L}_1\bullet\hat{L}_1^\dagger + \kappa_2\hat{L}_2\bullet\hat{L}_2^\dagger)\mathcal{E}(t_1,t)(\kappa_1\hat{L}_1\bullet\hat{L}_1^\dagger + \kappa_2\hat{L}_2\bullet\hat{L}_2^\dagger)\mathcal{E}(t,0)dtdt_1\\ &= \kappa_1^2\int_0^T\int_0^{t_1} \mathcal{E}(T,t_1)(\hat{L}_1\bullet\hat{L}_1^\dagger)\mathcal{E}(t_1,t)(\hat{L}_1\bullet\hat{L}_1^\dagger)\mathcal{E}(t,0)dtdt_1\\ &+ \kappa_1\kappa_2\int_0^T\int_0^{t_1} \mathcal{E}(T,t_1)(\hat{L}_1\bullet\hat{L}_1^\dagger)\mathcal{E}(t_1,t)(\hat{L}_2\bullet\hat{L}_2^\dagger)\mathcal{E}(t,0)dtdt_1\\ &+ \kappa_1\kappa_2\int_0^T\int_0^{t_1} \mathcal{E}(T,t_1)(\hat{L}_2\bullet\hat{L}_2^\dagger)\mathcal{E}(t_1,t)(\hat{L}_1\bullet\hat{L}_1^\dagger)\mathcal{E}(t,0)dtdt_1\\ &+ \kappa_2^2\int_0^T\int_0^{t_1} \mathcal{E}(T,t_1)(\hat{L}_2\bullet\hat{L}_2^\dagger)\mathcal{E}(t_1,t)(\hat{L}_2\bullet\hat{L}_2^\dagger)\mathcal{E}(t,0)dtdt_1.\end{aligned} \tag{4.18}$$

As such, the form of Eq. (4.17) allows for efficient and accurate encoding of the dynamics.

4.2 Volterra Integral Evaluation

Whilst the method proposed in Sect. 4.1.2 seems promising in theory, evaluating the action of the maps on the density matrix $\mathcal{M}_i\hat{\rho}$ poses a serious challenge – it is not immediately clear how to analytically or numerically approach these integrals. In this section, I provide intuition on how these maps can be approximated using Gaussian quadrature [82], a class of numerical integration methods used extensively in applied mathematics and engineering.

4.2.1 Basics – Linear Interpolation

The goal of this and the succeeding sections is to demonstrate that the integrals encoded in each superoperator $\mathcal{M}_i$ can be approximated by evaluating their integrands at different points and taking a linear combination of the results. For pedagogical reasons, let us assume that the integrand of $\mathcal{M}_1$ is well approximated as a linear function of time – that is,

$$\mathcal{E}(T,t)\mathcal{SE}(t,0) \approx B_0 + tB_1, \tag{4.19}$$

where B_0 and B_1 are some unknown superoperators. Then, we can evaluate the corresponding integral trivially,

$$\mathcal{M}_1 \approx \int_0^T (B_0 + tB_1)dt = T(B_0 + TB_1/2). \tag{4.20}$$

Consequently, by simply setting $t = T/2$ in Eq. (4.19), we can approximate the action of the map $\mathcal{M}_1$ on $\hat{\rho}$ as

$$\hat{\rho}_1 \approx T\mathcal{E}(T,T/2)\mathcal{SE}(T/2,0)\hat{\rho}. \tag{4.21}$$

We can extend this linear method to arbitrary order – for example, the n-th order map $\mathcal{M}_n$ would evaluate the approximate integrand

$$\mathcal{M}_n \approx \int_0^{t_1} ... \int_0^{t_{n-1}} \int_0^T \left(\hat{B}_0 + \sum_{i=1}^n B_i t_i\right) dt dt_1 ... dt_{n-1}. \tag{4.22}$$

Here, $0 < t < t_1 < ... < t_{n-1} < T$ refers to the times of the subsequent collapses. The region of integration in this n-dimensional 'time-space' thus defines an n-dimensional simplex, or the generalization of a triangle/trapezoid, see Ref. [83] for more details. Since the integrand in Eq. (4.22) varies linearly, the map can be calculated by simply evaluating the function at the 'centre-of-mass' of the simplex, which corresponds to equal times between the trajectory jumps: $t_i - t_{i-1} = T/n$. This yields the approximation for the n-th order map,

$$\mathcal{M}_n\hat{\rho} \approx \frac{T^n}{n!}\mathcal{E}(T,(n-1)T/n)\mathcal{S}...\mathcal{E}(2T/n,T/n)\mathcal{SE}(T/n,0)\hat{\rho}, \tag{4.23}$$

where $T^n/n!$ corresponds to the volume of the integrated simplex.

4.2.2 Gauss-Legendre Quadrature – First Order

Unless we consider a very small time step T, we anticipate that the above integrands will change much more rapidly than just linearly, thus rendering the approximation in Sect. 4.2.1 inaccurate. Thankfully, there exists a simple method to accurately approximate the integrals

of arbitrary one-dimensional *polynomials* over the interval $[-1, 1]$. This method, known as Gauss-Legendre Quadrature [84], requires multiple function evaluations at locations $\{x_i\}$, which are chosen to be the roots of the n-th order Legendre polynomials $P_n(x)$, with corresponding weights

$$w_i = \frac{2}{(1-x_i)^2[P_n'(x_i)]^2}. \tag{4.24}$$

This form is *exact* if $f(x)$ is polynomial of degree $2n-1$ or lower, where n can be extended almost arbitrarily [85]. More explicitly, we see that the approximation to the integral of some function $f(x)$,

$$\int_{-1}^{1} f(x)dx \approx \frac{5}{9} f\left(-\sqrt{\frac{3}{5}}\right) + \frac{8}{9} f(0) + \frac{5}{9} f\left(\sqrt{\frac{3}{5}}\right), \tag{4.25}$$

will be exact if $f(x)$ is an arbitrary polynomial of 5th order or lower. The idea for evaluating the first-order Volterra integral is almost the same – we simply map the region of integration $[-1, 1] \to [0, T]$, and obtain, for arbitrary order,

$$\mathcal{M}_1\hat{\rho} \approx T\sum_{i=0}^{n} w_i \mathcal{E}(T, t_i)\mathcal{S}\mathcal{E}(t_i, 0)\hat{\rho}, \quad \sum w_i = 1. \tag{4.26}$$

For standard numerical integrators, computing Eq. (4.26) would be prohibitively expensive, since computing the action of the $\mathcal{E}$ is computationally difficult with fast oscillating drives, see Chapter 2. However, this evolution is much simpler for Dysolve (see Chapter 3), thus opening a pathway for computing these functions.

4.2.3 Gauss-Legendre Quadrature – Second and Higher Orders

Whilst evaluating Gauss-Legendre quadrature on one-dimensional integrals is a standard and well-understood problem, extrapolating this method to higher orders is non-trivial and is the subject of ongoing research [86, 87, 88]. Further, approximating these integrals requires many more function evaluations than for the one-dimensional case to achieve the same $(2n-1)$ scaling of the Gauss-Legendre quadrature in Sect. 4.2.2. Nevertheless, since we are assuming the decay rates of the system to be small, the relative contribution of the second order jump term $\hat{\rho}_2$ should be small with respect to $\hat{\rho}_1$, and consequently tolerate a greater degree of error, and thus a lower order polynomial approximation.

Reddy et. al. [89] showed that by calculating the roots of integrated two-dimensional polynomials over the 2D simplex of order 4 and 7, a solution can be expressed with 4 and 7 points, respectively - see Fig. 4.1 for an illustration. Thus, the action of the map $\mathcal{M}_2$ would

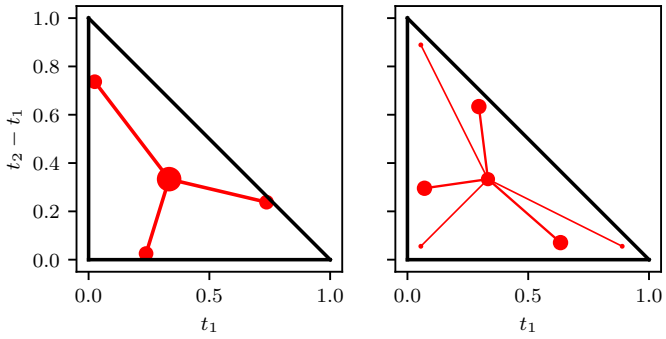

Figure 4.1 Locations of the $2D$ jumps for $n = 4$ and $n = 7$, able to accurately integrate up to 4-th and 7th order polynomials, respectively. The sizes of the red dots indicate their relative weight, and the lines are drawn for visualization. In each case, the central trajectory is located at (1/3,1/3), the centre-of-mass of the triangle.

be approximated as:

$$\mathcal{M}_2\hat{\rho} \approx T^2 \sum_{i=1}^{7} w_i \mathcal{E}(T, t_{i2})\mathcal{SE}(t_{i2}, t_{i1})\mathcal{SE}(t_{i1}, 0)\hat{\rho}, \quad \sum_{i=1}^{7} w_i = \frac{1}{2}, \tag{4.27}$$

where t_{ix}, t_{iy} correspond to the set of function evaluations shown in Fig. 4.1. This method also works with symmetric coordinates [90, 91], but at the cost of more function evaluations. Keeping in mind that each function evaluation requires building time-ordered propagators and several matrix multiplications, we must balance speed and accuracy.

Simpler, albeit less efficient forms, exist for the higher order simplexes [92, 93]. We can, in theory, continue approximating each order in the series of Volterra integrals until convergence, although again we have to compromise between the number of function evaluations and the accuracy of the solver. For Voltsolve, we limit ourselves to the 4 and 7-point evaluation of the second-order Volterra integral, and a 4 point evaluation of the third-order, equivalent to integrating a cubic polynomial exactly [92]. For fourth and higher orders, there is generally negligible advantage from integrating the function more accurately than the linear approximation, and so I choose to simply evaluate a single trajectory at the centre-of-mass, as in Eq. (4.23).

4.3 Alternative Method – Correction

From Sect. 4.2.3, it is clear that calculation of the maps $\mathcal{M}_i$ using Gauss-Legendre quadrature becomes increasingly challenging for higher orders. An alternative way of performing this calculation is to attempt to calculate all orders simultaneously by iterating over the previous order. Instead of calculating $\mathcal{M}_1\hat{\rho}$ explicitly, let us consider the following expansion:

$$\mathcal{M}\hat{\rho} = \left[\prod_{i=1}^{n} \mathcal{E}(t_{i+1}, t_i)(I + w_i T\mathcal{S})\mathcal{E}(t_i, 0)\right]\hat{\rho}, \tag{4.28}$$

such that w_i satisfy the roots of the n-th order Gauss-Legendre polynomial. Then, we note that

$$\begin{aligned}\mathcal{M}\hat{\rho} &= \mathcal{M}_0\hat{\rho} + \mathcal{M}_1\hat{\rho} + \prod_{j>i}^{n} T^2 w_i w_j \mathcal{E}(T, t_j)\mathcal{S}\mathcal{E}(t_j, t_i)\mathcal{S}\mathcal{E}(t_i, 0)\hat{\rho} + ... \\ &= \mathcal{M}_0\hat{\rho} + \mathcal{M}_1\hat{\rho} + \sum_{i=2}^{n} \mathcal{M}'_i\hat{\rho}.\end{aligned} \tag{4.29}$$

We now see that there are now additional contributions above just $\mathcal{M}_1$ – indeed, there are a number of maps $\mathcal{M}'_i$, each consisting of terms with i jumps. Specifically, we see that each two-jump contribution has a weight $w_i w_j$ and times t_i and t_j associated with the first collapse and second collapses, respectively. To illustrate this, I plot the relative distributions of the jump times in Fig. 4.3, which are unevenly distributed around the centre-of-mass of the triangle, $(1/3, 1/3)$.

The primary problem here is that to preserve trace, we require the weight of the n-th order contribution to be $1/n!$ (corresponding to the volume of the n-th order simplex), which is not satisfied by this expansion. To see this more explicitly, I plot in Fig. 4.2 the convergence of the relative weight of each order as a function of the Gauss-Legendre order. For example, at second order, the weight W_2 is

$$W_2 = \sum_{j>i} w_i w_j < \frac{1}{2}. \tag{4.30}$$

Consequently, we seek a correction term, $\mathcal{M}^c_n$, such that we approximate the map $\mathcal{M}_n$:

$$\mathcal{M}_n \approx \mathcal{M}'_n + \mathcal{M}^c_n. \tag{4.31}$$

A simple remedy is to calculate the centre-of-mass of the contributions from $\mathcal{M}'_2$ and add a single correction of weight w_c and jump times t_{1c}, t_{2c} such that the total weight and

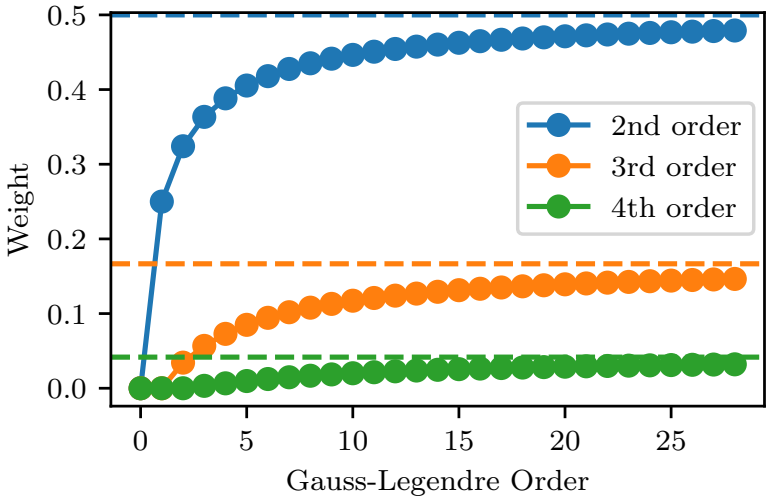

Figure 4.2 Convergence of the weights of the corresponding map $\mathcal{M}_n$. The dashed lines show the asymptote at $1/n!$, the weight required to preserve the trace of the density matrix.

centre-of-mass of the jumps realigns with the simplex, as seen in Sect. 4.2.1. This yields:

$$\sum_{j>i}^{n} w_i w_j + w_c = \frac{1}{2}, \quad \sum_{j>i}^{n} w_i w_j t_i + w_c t_{1c} = \sum_{j>i}^{n} w_i w_j (t_j - t_i) + w_c (t_{2c} - t_{1c}) = \frac{1}{6}. \quad (4.32)$$

An identical methodology follows for the higher orders, allowing for an excellent approximation to each order whilst maintaining the correct contribution from the corresponding jump term. The times and corresponding weights of these correction trajectories are plotted in Fig. 4.3 as blue dots. Capturing how this series converges, and the impact of including additional corrections, is part of my ongoing research.

4.3.1 Performance Benchmark – Transmon Readout

The expansion utilized in Voltsolve relies on the decreasing likelihood of a larger number of jumps in a time interval Δt. As such, we expect the solver to behave optimally for driven systems with small decay rates –for example, simulating a Surface-17 Quantum Error Correction cycle with qubit T_1 lifetimes in the tens of microseconds [20]. This notwithstanding, qubit readout is known to be challenging to simulate, thanks to the presence of loss and periodic driving of the readout resonator. Further, as we will see in Chapter 5, it is important to keep many levels of the transmon and resonator subspaces and avoid rotating-wave approximations in order to accurately capture the dynamics, which

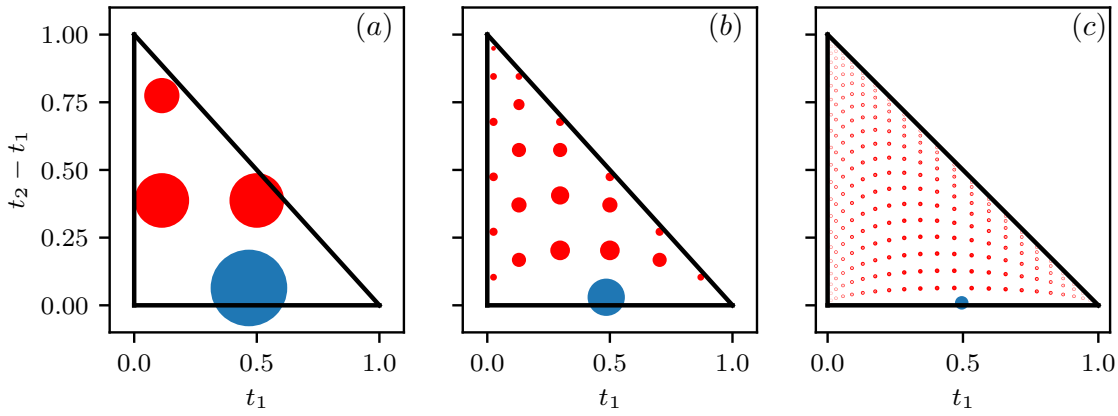

Figure 4.3 Locations of the jump evaluations $(t_1, t_2 - t_1)$ corresponding to the second order map $\mathcal{M}_2$ at different orders in the Gauss-Legendre expansion, with (a) $n = 3$, (b) $n = 7$ and (c) $n = 21$. The sizes of the dots indicate their relative weight, and the blue dots correspond to the evaluation of the correction term $\mathcal{M}_2^c$.

adds to the complexity. In keeping with the themes of the following chapters, we consider the simulation of a driven transmon-resonator system with a lossy resonator to test this solver. The simulated Hamiltonian and master equation are of the form

$$\begin{aligned}\dot{\hat{\rho}} &= -i[\hat{H}(t), \hat{\rho}] + \kappa \mathcal{D}[\hat{a}]\hat{\rho}, \\ \hat{H}(t) &= 4E_c(\hat{n}_t - n_g)^2 - E_J \cos\hat{\varphi}_t + \omega_r \hat{a}^\dagger \hat{a} - ig(\hat{n}_t - n_g)(\hat{a} - \hat{a}^\dagger) + i\Omega_0(\hat{a}^\dagger - \hat{a})\sin(\omega_d t).\end{aligned} \tag{4.33}$$

Here, we use the parameters $\omega_r/2\pi = 5.156$ GHz, $g/2\pi = 200$ MHz, $E_J/2\pi = 10.512$ GHz, $E_C/2\pi = 281.2$ MHz, $\kappa/2\pi = 35.3$ MHz, $\Omega_0/2\pi = 80$ MHz and $\omega_d/2\pi = 5.19$ GHz for the simulation. Further, we use $t_{\text{dim}} = 13$ states to represent the transmon subspace and $r_{\text{dim}} = 12$ states for the resonator, corresponding to evolving a density matrix of size $N = 156$, and consider an evolution time of $t = 20$ ns.

We consider two versions of Voltsolve – one where each order is calculated individually (see Sect. 4.2.3) and the other where the orders are calculated simultaneously and corrections added (see Sect. 4.3). For the former solver, we use 7, 4 and 1 trajectories to evaluate the second, third and fourth-order Volterra integrals respectively. For the latter solver, which I will refer to as Voltsolve(alt), we consider an expansion up to the 7th order in the Gauss-Legendre quadrature and add corrections up to 4th order in the jump operator, see Sect. 4.3. The Dysolve algorithm is used in both cases to calculate the propagators in Eq. (4.4), where I consider up to 4th order in the drive term $i\Omega_0(\hat{a} - \hat{a}^\dagger)$.

For comparison, we shall use the standard numerical integrator built into qutip, Mesolve

[94], which uses an adaptive step Runga-Kutta solver [66]. In doing so, we must realize that Mesolve is optimal for *sparse* inputs. For the optimal performance of both solvers, we solve the master equation in the bare frame for Mesolve, and in the eigenbasis of the undriven Hamiltonian for Voltsolve. Further, we consider Mesolve for three different tolerances, ranging from 10^{-9} to 10^{-11}. Some simulation results are summarized in the table below:

$t_{\text{dim}} = 13,\ r_{\text{dim}} = 12$				
Method	Tol.	Δt (2π ns)	Time	Accuracy
mesolve	$1e^{-9}$	adaptive	137s	$1.83e^{-4}$
mesolve	$1e^{-10}$	adaptive	145s	$1.75e^{-5}$
mesolve	$1e^{-11}$	adaptive	161s	$\times$
voltsolve	$\times$	0.5	31s	$2.22e^{-6}$
voltsolve(alt)	$\times$	0.5	10.4s	$3.08e^{-6}$

Table 4.1 Comparative performance of Mesolve and Voltsolve for integrating the master equation in Eq. (4.33), corresponding to a transmon-readout simulation.

The accuracy of the solver is simply defined as the L_2 norm $||\hat{\rho}(t_f) - \hat{\rho}_c(t_f)||$, where $\hat{\rho}(t_f)$ is the final density matrix from one of the solvers and $\hat{\rho}_c(t_f)$ is calculated from Mesolve using a tolerance of 10^{-11}. We find that not only do both versions of Voltsolve outperform Mesolve concerning computational time, boasting at least a 4.4× and 13× speed-up respectively for the two solvers, but additionally find that the accuracy of Voltsolve is *greater* than qutip for error tolerances of $1e^{-9}$ and $1e^{-10}$, with a similar computational advantage found at larger system sizes. This demonstrates that, whilst Voltsolve is a perturbative solver, we are still able to achieve excellent precision and speed.

4.4 Future Work and Outlook

The Voltsolve algorithm is very much a work in progress – the code is far from optimized, and there likely exist many different approaches, both theoretically and programmatically, to evaluate the Volterra integrals more efficiently. As demonstrated by IBM in Ref. [72], the Dysolve algorithm offers a powerful way of solving time-dependent systems, and adapting the algorithm to master equations could unlock currently out-of-reach simulations of highly complex systems. Preliminary analysis indicates that the dominant errors for Voltsolve(alt) in the above simulations are from the corrections at second order. To compensate, we could consider multiple correction trajectories, rather than just a single correction trajectory as detailed in Sect. 4.3. This would likely yield a greater accuracy of the solver than the original method detailed in Sect. 4.2.3. I leave this to future work.

My thanks to Élie Genois and Ronan Gautier for many discussions and assistance in developing the method.

Chapter 5

Transmon Ionization

"There are three ways to do things. The right way, the wrong way, and the Max Power way!

"Isn't that the wrong way?"

"Yeah, but faster!"

– Homer and Bart Simpson, *The Simpsons*

In this chapter, I provide some insights into non-QND effects observed when readout is performed with a *strong* drive, and provide the necessary background for my publication in collaboration with Sandbox [2]. I additionally demonstrate the importance of the external gate charge and how it can impact readout.

5.1 Transmon Readout – Background

Over the past few years the field of circuit QED has evolved greatly, with recent demonstrations of a surface code implemented with transmon qubits [20, 21]. However, readout remains a limiting factor of these processes – for example, the readout time and average fidelity in Ref. [21] was 500 ns and 98%, which lags significantly behind the two-qubit gate performance (35 ns and 99.5%, respectively). The readout makes up for more than half of the error correction cycle time (921 ns) and accounts for 12.1% of the 'error budget' for each error correction cycle – however, when including the error associated with data qubit idling (predominantly due to the long readout process), the contribution to the total error budget

increases to more than 30%. Similarly, the readout fidelities and times in Ref. [20] were 99% and 400 ns respectively, with a cycle time of 1.1 μs. Increasing the readout fidelity and reducing the integration time is thus critical to advancing the field.

The primary reason that readout is comparatively slow is that the measurement drives need to be relatively weak to ensure that the measurement remains QND. Using stronger drives could theoretically accelerate this process, but this is known to result in 'leakage', or population of the higher energy levels of the transmon, which greatly reduces the readout fidelity and can disrupt the subsequent operations in the surface code cycles. A commonly used metric is the 'critical photon number', n_{crit}, which is used to indicate the maximum number of photons in the readout resonator before the theory of dispersive measurement is predicted to fail – however, experimental results have demonstrated that the readout can break down well before this critical photon number is reached [30]. Consequently, there is a need to better understand what photon numbers cause measurements to breakdown, so as to optimize readout parameters, fidelities, and integration times.

The breakdown of the readout can be challenging to capture from both a numerical and theoretical point of view. Since many of the damaging processes are caused by resonances of the low-lying energy states with states that lie outside of the transmon confining potential, the Kerr nonlinear oscillator model does not capture these effects, as we saw in Chapter 2. Moreover, to simulate the readout process whilst capturing these resonances, it can be necessary to preserve up to 23 transmon eigenstates, which greatly increases the numerical complexity. In this chapter, I first describe how analysing the static Hamiltonian (without the readout drive) can help us identify these resonances, and how we can identify sets of states which become populated during the readout process. This puts into greater context my publication (Ref. [2]). Finally, I demonstrate how the offset gate charge can impact readout, which could explain some experimentally observed effects [95].

5.2 Transmon Readout – Diagonalization

In this section, we return to the Hamiltonian introduced in Sect. 2.2. The transmon-resonator model, without making an RWA or using the Fock basis for the transmon, is given by

$$\hat{H}_{tr} = 4E_c(\hat{n}_t - n_g)^2 - E_J \cos\hat{\varphi}_t + \omega_r \hat{a}^\dagger \hat{a} - ig(\hat{n}_t - n_g)(\hat{a} - \hat{a}^\dagger). \tag{5.1}$$

The goal here is to diagonalize the above Hamiltonian and associate its eigenstates with qubit and resonator states. In this section, all calculations are done with the following set of parameters:

$$E_J/h = 22.66 \text{ GHz},\ E_J/E_c = 134,\ g/2\pi = 240 \text{ MHz},\ \omega_r/2\pi = 6.852 \text{ GHz}. \tag{5.2}$$

These parameters correspond to a qubit frequency $\tilde{\omega}_q/2\pi = 5.279$ GHz, an anharmonicity $\alpha/2\pi = 163$ MHz, a resonator pull $\chi/2\pi = -8.8$ MHz and critical photon number $n_{\text{crit}} \approx 11$.

5.2.1 Identifying the computational states

As described in Sect. 2.1, the charge and Fock bases are the appropriate choices to represent the transmon and resonator subspaces, respectively. After constructing the Hamiltonian in Eq. (5.1) and using a standard numerical diagonalization procedure, we obtain the spectral decomposition of the Hamiltonian $\hat{H}_{tr} = \sum_\lambda \lambda|\lambda\rangle\langle\lambda|$. Without further information on the structure of the eigenstates, such a decomposition gives little insight into the underlying physics. We would like to label each eigenstate according to its transmon-like or resonator-like character, $\hat{H}_{tr} = \sum_{i,n} E_{\overline{i,n}}|\overline{i,n}\rangle\langle\overline{i,n}|$. Here, i and n represent the 'transmon' and 'resonator' indices, respectively, and $E_{\overline{i,n}}$ the corresponding energy.

The simplest state to identify is the ground state, $|\overline{g,0}\rangle$, which is the eigenstate associated with the smallest eigenvalue: $E_{\overline{g,0}} = \min\{\lambda\}$, noting that the minimum energy may be non-zero due to the presence of counter-rotating terms in the charge-charge coupling of Eq. (5.1). Moreover, in the dispersive regime, we expect the excited state of the qubit to only be dressed *weakly* by the resonator. Consequently, we can identify the excited state to be the eigenstate which has maximum overlap with the *bare* state $|e,0\rangle = |e\rangle \otimes |0\rangle$:

$$|\overline{e,0}\rangle \equiv \text{argmax } |\langle\lambda|e,0\rangle|^2. \tag{5.3}$$

From the dispersive theory in Sect. 2.2 we anticipate this overlap to be of the order $1 - g^2/\Delta^2$, and thus close to unity and simple to identify.

5.2.2 Generalizing the Jaynes-Cummings Ladder

Recall from Sect. 2.2.1 in Chapter 2 that we could approximate the transmon-resonator model with a Hamiltonian similar to the Jaynes-Cummings model [50]. After diagonalization, this yielded a distinct ladder for each qubit state, which we can climb or down by adding or

removing photons. Here we seek to generalize the notion of the ladder to work for Eq. (5.1), including the effects of the counter-rotating terms and the cosine potential. In other words, we seek divide the Hilbert space into subspaces $\{|\overline{g,n}\rangle\}$, $\{|\overline{e,n}\rangle\}$, ..., with each subspace consisting of n states, such that the drive on the resonator induces dynamics predominately within each subspace. Let us reconsider the *driven* system:

$$\hat{H}(t) = \sum_{i,n} E_{\overline{i,n}} |\overline{i,n}\rangle\langle\overline{i,n}| + i\mathcal{E}(\hat{a}^\dagger - \hat{a})\sin(\omega_d t). \tag{5.4}$$

Inspired by the Jaynes-Cummings ladder, we can therefore ask ourselves a simple question – what eigenstate of the Hamiltonian is principally populated after adding a resonator photon to the state $|\overline{i,n}\rangle$? This motivates a simple definition:

$$|\overline{i,n+1}\rangle \equiv \underset{|\lambda\rangle\in S}{\text{argmax}} \, |\langle\lambda|\hat{a}^\dagger|\overline{i,n}\rangle|^2, \tag{5.5}$$

where S is the set of eigenstates $|\lambda\rangle$ which have not been labelled. Extending this concept further, we can define the normalized overlap metric [2]:

$$|\overline{i,n+1}\rangle \equiv \underset{|\lambda\rangle\in S}{\text{argmax}} \, C_i(n), \quad C_i(n) = \left| \frac{\langle\lambda|\hat{a}^\dagger|\overline{i,n}\rangle}{\langle\overline{i,n}|\hat{a}\hat{a}^\dagger|\overline{i,n}\rangle} \right|^2. \tag{5.6}$$

To differentiate this method from the Jaynes Cummings ladder, we refer to each set of n states $\{i\} = \{|\overline{i,n}\rangle\}$ as a *branch* of the resonator.

Using this metric, we can systematically label every eigenstate of the system. Suppose that the initial diagonalization of Eq. (5.1) was performed with t_{dim} transmon states and r_{dim} resonator states. Then, the branches are generated by the following procedure:

- Identify the ground state $|\overline{0,0}\rangle$;
- Identify the remaining $(r_{\text{dim}} - 1)$ states in the branch by iteratively maximizing the overlap metric as in Eq. (5.6), and subsequently removing each identified state from the set S;
- Increase the transmon index i by 1 and reset the resonator index $n = 0$. Label the lowest energy state in the set S as $|\overline{i,0}\rangle$, the first element in the new branch $\{i\}$;
- Repeat the previous two steps.

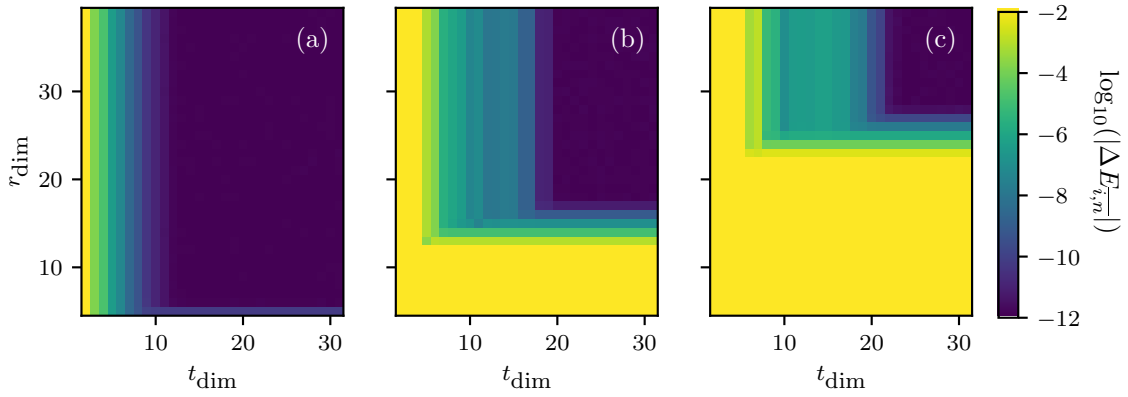

Figure 5.1 Convergence of the eigenvalue difference $\log_{10}(|\Delta E_{\overline{i,n}}|)$ corresponding to states (a) $|\overline{0,0}\rangle$, (b) $|\overline{0,10}\rangle$ and (c) $|\overline{0,20}\rangle$ when diagonalized with t_{dim} transmon eigenstates and r_{dim} resonator states as compared to diagonalization with 40 transmon and 70 resonator states.

5.2.3 Eigenvalue convergence

A common pitfall in simulating and diagonalizing the transmon-resonator system is to use an insufficient number of transmon states, t_{dim}, in the diagonalization procedure. In Fig. 5.1, I plot $\log_{10}(|\Delta E_{\overline{i,n}}|)$, the logarithm of the difference between the converged eigenvalues (in GHz) corresponding to states (a) $|\overline{0,0}\rangle$, (b) $|\overline{0,10}\rangle$, and (c) $|\overline{0,20}\rangle$ and the eigenvalues computed from diagonalization of the system with the dimensions quoted in the x and y axes. The converged eigenvalues were computed by diagonalizing the system with $t_{\mathrm{dim}} = 40$, $r_{\mathrm{dim}} = 70$ – increasing the dimension further only leads to machine precision fluctuations in the eigenvalues. We see a surprising need to conserve as many as 22 transmon eigenstates for convergence of the $|\overline{0,20}\rangle$ state due to the hybridization of states in the $\{0\}$ branch with highly energetic states.

5.2.4 Master Equation Example – Branch Population

To demonstrate the usefulness of this branch analysis, I plot in Fig. 5.2 (a) the overlap metric for the ground branch $C_g(n)$ as a function of the resonator index n. This metric remains close to unity over the entire range of n. This indicates that neighbouring states within a branch, $|i,n\rangle$ and $|i,n+1\rangle$, are strongly coupled by the creation operator $\hat{a}^\dagger$, even though some weak resonances with states in other branches will occur – for example, the feature at $n \approx 30$, where the overlap metric drops to 0.9. In comparison, in Fig. 5.2 (b) I plot the overlaps of each of the identified eigenstates in the ground state branch with

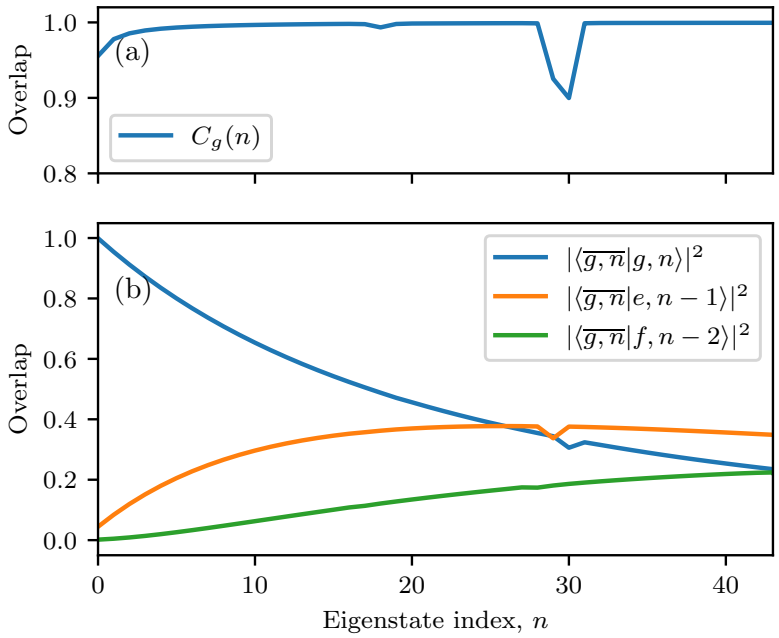

Figure 5.2 (a) Overlap metric $C_g(n)$ used to define the ground state branch. (b) Overlap of the identified states in the ground state branch with corresponding bare states $|g,n\rangle$ and $|e,n-1\rangle$.

the bare states $|g,n\rangle, |e,n-1\rangle$ and $|f,n-2\rangle$. Somewhat surprisingly, we see the overlap between the identified eigenstate and the equivalent bare state $|\langle\overline{g,n}|g,n\rangle|^2$ drop *below* the overlap with the excited state, $|\langle\overline{g,n}|e,n-1\rangle|^2$ around index $n = 26$. We further see a significant contribution of the bare states $|f,n-2\rangle$ at high photon numbers, which would naïvely suggest that populating the resonator with a relatively low photon number results in an unmanageable level of leakage.

Whilst this may appear alarming, it is of remarkably little consequence. The reason is straightforward: the dispersive approximation is not valid at this resonator population, and consequently, the eigenstates cannot be expected to be perturbatively close to the bare basis states as I have defined them. To illustrate this point more concretely, I consider a simple readout experiment. I initialize the qubit in the ground state, $\rho_0 = |\overline{g,0}\rangle\langle\overline{g,0}|$ and evolve the system using the master equation:

$$\dot{\rho} = -i[\hat{H}(t), \hat{\rho}] + \kappa\mathcal{D}[\hat{a}]\hat{\rho}, \tag{5.7}$$

where I choose $\kappa/2\pi = 30$ MHz. For the first 48 nanoseconds, the drive amplitude is chosen to be constant, $\mathcal{E}/2\pi = 180$ MHz, corresponding to a steady state resonator population $\langle\hat{a}^\dagger\hat{a}\rangle \approx 35$ for the chosen parameters. After time $t = 48$ ns, the drive is turned off and the

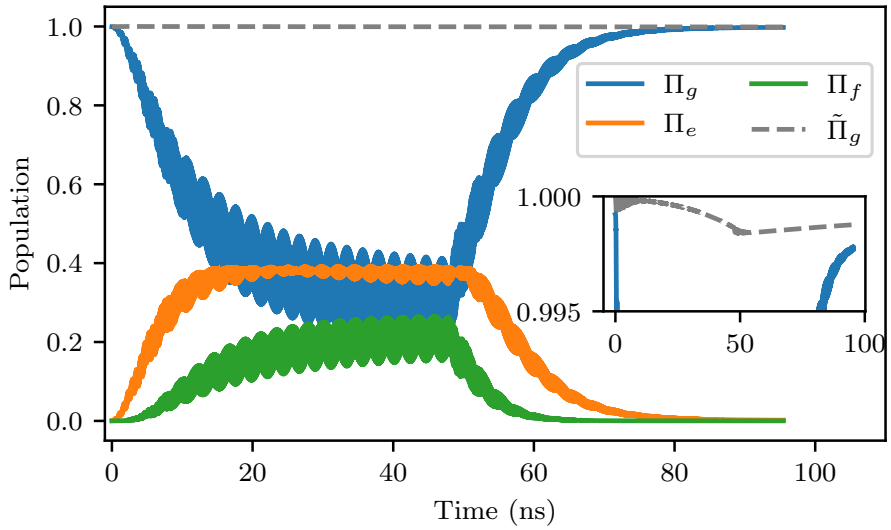

Figure 5.3 Evolution of the ground state $\rho_0 = |\overline{g,0}\rangle\langle\overline{g,0}|$ governed by the master equation in Eq. (5.7). The drive is turned off at 48 nanoseconds. The blue, orange and green curves correspond to the population of the ground, excited and leakage *bare* subspaces – the dashed line indicates the population of the *dressed* subspace. The inset details the deviation of the dressed subspace population away from unity.

system is allowed to relax. In Fig. 5.3, I plot the population of the density matrix projected onto the different *bare* subspaces $\Pi_i = \sum_n |i,n\rangle\langle i,n|$, and compare this to the occupation of the *dressed* ground-state branch $\tilde{\Pi}_g = \sum_n |\overline{i,n}\rangle\langle\overline{i,n}|$. Whilst we see a significant population of the *bare* excited and leakage subspaces Π_e, Π_f, which first exceeds the population of the bare ground subspace Π_g at approximately 15 nanoseconds, there is a negligible deviation from the population of the dressed ground-state subspace $\tilde{\Pi}_g$ throughout the entire evolution. Moreover, we see vanishing population remaining in the bare excited and leakage subspaces at the end of the evolution, which confirms that the population of the bare leakage state $|f\rangle$ is irrelevant for readout. However, as we'll see in Sect. 5.3, the population of transmon eigenstates *close to and outside of the cosine potential well* is a strong indicator of ionization.

5.3 Publication – Dynamics of Transmon Ionization [2]

In this publication, we simulate the readout of a superconducting transmon qubit using a strong drive. These simulations included up to 32 transmon states and 1024 resonator states to ensure we captured all of the relevant dynamics. The resulting computational complexity required the use of Tensor Processing Units (TPUs) for numerical integration. We observe and make sense of ionization dynamics, whereby eigenstates of the transmon close to and above its confining potential well are populated.

In this joint effort, Alexandru Petrescu, Joachim Cohen and Alexandre Blais gave invaluable theory support and helped direct the project, with the Sandbox team providing extensive help in using and managing the Tensor Processing Units for the simulations. I developed and launched the simulation scripts, proposed the branch and semiclassical analyses, and wrote the manuscript.

Editors' Suggestion

Dynamics of Transmon Ionization

Ross Shillito,[1,2,*] Alexandru Petrescu,[1,3] Joachim Cohen,[1] Jackson Beall,[2] Markus Hauru,[2] Martin Ganahl,[2] Adam G.M. Lewis,[2] Guifre Vidal,[2,4,5] and Alexandre Blais[1,5]

[1] *Institut Quantique and Département de Physique, Université de Sherbrooke, Sherbrooke, Quebec, J1K 2R1 , Canada*

[2] *Sandbox@Alphabet, Mountain View, California 94043, USA*

[3] *Laboratoire de Physique de l'Ecole Normale Supérieure, Mines-Paristech, CNRS, ENS-PSL, Inria, Sorbonne Université, PSL Research University, Paris 75005, France*

[4] *Google Quantum AI, Mountain View, California 94043, USA*

[5] *Canadian Institute for Advanced Research, Toronto, Ontario M5G 1M1, Canada*

(Received 21 April 2022; revised 20 July 2022; accepted 11 August 2022; published 13 September 2022)

Qubit measurement and control in circuit quantum electrodynamics (QED) rely on microwave drives, with higher drive amplitudes ideally leading to faster processes. However, degradation in qubit coherence time and readout fidelity has been observed even under moderate drive amplitudes corresponding to a few photons populating the measurement resonator. Here, we numerically explore the dynamics of a driven transmon-resonator system under strong and nearly resonant measurement drives and find clear signatures of transmon ionization where the qubit escapes out of its cosine potential. Using a semiclassical model, we interpret this ionization as resulting from resonances occurring at specific resonator-photon populations. We find that the photon populations at which these spurious transitions occur are strongly parameter dependent and that they can occur at low resonator-photon population, something that may explain the experimentally observed degradation in measurement fidelity.

DOI: 10.1103/PhysRevApplied.18.034031

I. INTRODUCTION

Dispersive readout in circuit quantum electrodynamics (QED) is realized by driving a measurement resonator coupled to the qubit [1]. In principle, increasing the drive amplitude, and thereby the resonator-photon population, increases the measurement rate, something that is expected to lead to fast, high-fidelity, and quantum nondemolition (QND) readout [2]. However, experimentally the fidelity and QND character of the measurement of transmon qubits [3] is often observed to decrease beyond a photon-number threshold, which can be as small as a few photons [4,5]. Perturbative models have been made in attempts to explain these observations but have limited applicability beyond low photon numbers [6–8].

In this paper, we go beyond perturbative treatments by numerically investigating the full dynamics of a strongly driven transmon-resonator system. At distinct resonator populations, we find clear signatures of transmon ionization where transmon states above the Josephson-junction potential are occupied [9,10]. Because these states are not strongly influenced by the Josephson potential, they are well described by charge states. Consequently, for states above the transmon well, the transmon-resonator coupling appears longitudinal and the system dynamics are consequently modified. Importantly, the term *ionization* is used loosely here and refers only to the escape of the transmon population to states lying above the Josephson cosine-potential well.

Accurate simulation of the dynamics of transmon ionization requires description of the density matrix of the system on a truncated Hilbert space of very large dimension, something that is made possible here by the use of large-scale computational accelerators known as tensor processing units (TPUs). While other studies have been limited to steady-state calculations with strongly detuned drives [9], the computational power of TPUs allows us to simulate the full time dependence with drives that are resonant with the resonator, as is relevant for qubit measurement. Accounting for tens of transmon levels and hundreds of resonator states, we moreover see signatures of the high-power readout [11–13]. We interpret these numerical results using a semiclassical model capturing the nonlinear

*Ross.Shillito@USherbrooke.ca

response of the driven system. Using this model, we identify parameter regimes where ionization is expected to occur at sufficiently small photon number to affect dispersive readout, observations that are in qualitative agreement with experiments [4].

This paper is organized as follows. In Sec. III, we introduce the model, provide details on its TPU implementation, and present the numerical results on the dynamics of transmon ionization. Next, in Sec. III, we formulate a semiclassical theory that allows us to interpret the dynamics of the coupled transmon-resonator system as governed by transitions between qubit-state-dependent effective resonators (Sec. III A) obeying nonlinear equations of motion (Sec. III B). In Sec. III C, we choose system parameters that illustrate how ionization can reduce readout fidelity even at the low drive powers that are typical of dispersive readout in circuit QED. We summarize our findings in Sec. IV.

II. MASTER-EQUATION SIMULATIONS WITH TENSOR PROCESSING UNITS (TPUs)

A. Model

We consider a transmon capacitively coupled to a resonator [see Fig. 1(a)]. In the presence of a drive of amplitude $\mathcal{E}$ and frequency ω_d on the resonator, the Hamiltonian of the system takes the usual form ($\hbar = 1$) [2]

$$\hat{H}(t) = 4E_c\hat{n}_t^2 - E_J\cos\hat{\varphi}_t + \omega_r\hat{a}^\dagger\hat{a} - ig\hat{n}_t(\hat{a} - \hat{a}^\dagger) - i\mathcal{E}(\hat{a} - \hat{a}^\dagger)\sin(\omega_d t). \quad (1)$$

The first two terms correspond to the free transmon Hamiltonian with charging energy E_c, Josephson energy E_J, charge operator $\hat{n}_t$, and phase operator $\hat{\varphi}_t$. We denote the eigenenergies and eigenstates of the free transmon Hamiltonian E_i and $|i\rangle$, respectively. The first two of those eigenstates, labeled $\{|0\rangle, |1\rangle\}$, span the computational basis of the qubit. Of the higher excited states, an approximate number $2E_J/\omega_p$ are bound states lying within the cosine-potential well illustrated in Fig. 1(b). Here, $\omega_p = \sqrt{8E_CE_J}$ is the plasma frequency, which is approximately the 0–1 transition frequency of the transmon [3]. Moreover, we label as $|n\rangle$ the eigenstates of the free-resonator Hamiltonian of frequency ω_r, corresponding to the bosonic annihilation operator $\hat{a}$. The transmon-resonator coupling of amplitude g in the second line of Eq. (1) includes fast-rotating terms that are beyond the rotating-wave approximation (RWA) and are important to capture the contribution of high-energy states [14]. In the absence of the drive, the dressed energies and states of the coupled system are denoted $E_{\overline{i,n}}$ and $|\overline{i,n}\rangle$.

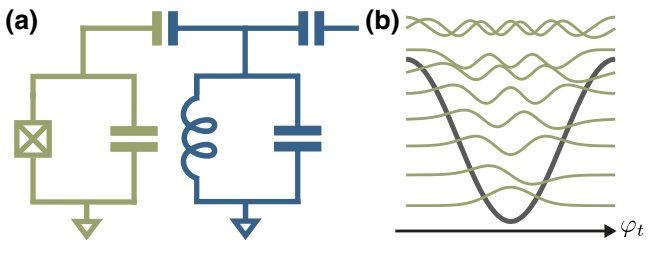

FIG. 1. (a) A transmon (green) capacitively coupled to a driven resonator (blue). (b) The cosine potential of the transmon with its first nine eigenstates. For our choice of parameters, the last three states are ionized states lying above the top of the potential.

Including cavity loss at a rate κ, the driven transmon-resonator system is described by the usual Lindblad master equation [2],

$$\dot{\hat{\rho}} = -i[\hat{H}(t), \hat{\rho}] + \kappa\mathcal{D}[\hat{a}]\hat{\rho}, \quad (2)$$

with the dissipator

$$\mathcal{D}[\hat{O}]\hat{\rho} = \hat{O}\hat{\rho}\hat{O}^\dagger - \frac{1}{2}\left\{\hat{O}^\dagger\hat{O}, \hat{\rho}\right\}. \quad (3)$$

Because we are interested in capturing the dynamics of the system in the presence of a large-amplitude nearly resonant drive on the resonator, leading to several hundred of photons and highly excited states of the transmon, we keep up to 32 states in the transmon and 1024 states in the resonator for the most demanding simulations, corresponding to a total Hilbert-space dimension $2^{15} = 32\,768$. That is, the joint density matrix $\hat{\rho}$ for the transmon-resonator system is a Hermitian matrix of size $2^{15} \times 2^{15}$, which thus contains $2^{30} = 1\,073\,741\,824$ time-dependent complex coefficients. Furthermore, given that the unbound transmon states are approximately eigenstates of the charge operator, their eigenvalues increase quadratically rather than linearly, as is the case for the bound states. This increases the complexity of numerical simulations, as larger eigenvalues require smaller integration step sizes for convergence. Additional details, including a prescription for adapting the resonator Hilbert-space size, are given in Appendix A.

In order to perform these challenging numerical simulations, we resort to TPUs. Google's TPUs are application-specific integrated circuits designed exclusively to accelerate large-scale machine-learning workloads [15]. Recently, they have been repurposed for other high-performance computational tasks [16–28], including simulations of quantum systems in large Hilbert spaces [24–28]. Due to their (i) matrix-multiply units capable of accelerating matrix multiplication, (ii) large amounts of high-bandwidth memory, and (iii) fast intercore interconnects directly connecting up to thousands of cores, TPUs are particularly fast at performing large-scale dense linear-algebra operations, which are required, e.g., to numerically integrate the above Lindblad master equation. More details on TPUs can be found in Appendix B.

For a given drive amplitude $\mathcal{E}$, simulations begin with the system initialized in either the dressed ground $|\overline{0,0}\rangle$ or excited $|\overline{1,0}\rangle$ states. To reduce the simulation time while capturing the transient dynamics including the ionization of the transmon, we simulate the evolution over a time of at least κ^{-1}. After each period of the drive, the reduced transmon, $\hat{\rho}_t$, and resonator, $\hat{\rho}_r$, density matrices are recorded.

Unless otherwise stated, in all of our simulations as well as in the semiclassical model discussed in Sec. III, we use the parameters $E_J/E_C = 50$ and $E_C/h = 280$ MHz, corresponding to a transmon of frequency $\omega_t/2\pi = 5.304$ GHz and anharmonicity $\alpha/2\pi = 322$ MHz. These parameters result in approximately six bound transmon states within the cosine potential (see Fig. 1). The resonator frequency is $\omega_r/2\pi = 7.5$ GHz and $g/2\pi = 250$ MHz. A resonator loss rate of $\kappa/2\pi = 20$ MHz is chosen to ensure fast resonator dynamics and the drive amplitude takes values in the range $\mathcal{E}/2\pi \in [0, 440]$ MHz. We fix the drive frequency to the bare-resonator frequency, $\omega_d = \omega_r$. The above parameters place the system in the dispersive regime, with $\chi/2\pi = -5.6$ MHz and $\chi/\kappa = 0.28$, where the dispersive shift is defined here as $\chi = (E_{11} - E_{10} - E_{01} + E_{00})/2$. Note that the resonator population is much smaller than the critical photon number $n_{\text{crit}} = (\Delta/2g')^2 \approx 15$, with $g' = (E_J/32E_C)^{1/4}g$ [2]. Finally, to minimize the number of simulations that are required, we note that we ignore a possible offset gate charge on the transmon. The effect of the gate charge on the ionization will be studied in a future work [29].

B. Numerical results

Figures 2(a) and 2(b) illustrate the average photon number $N_r = \langle \hat{a}^\dagger \hat{a} \rangle$ (red lines) and the average transmon population $N_t = \sum_i i \langle i|\hat{\rho}_t|i\rangle$ (black lines) as a function of time for the drive amplitude $\mathcal{E}/2\pi = 280$ MHz. The system is initialized in the dressed ground state $|\overline{0,0}\rangle$ in Fig. 2(a) and in the dressed excited state $|\overline{1,0}\rangle$ in Fig. 2(b). Also shown is the instantaneous distribution $N_{i,t} = \langle i|\hat{\rho}_t|i\rangle$ of the transmon states (color scale). The difference between the two initial states is striking. For this drive amplitude, when starting in the ground state, the transmon leaks out of its initial state but the distribution largely remains confined within the cosine potential of the Josephson junction. In contrast, when the transmon is initialized in its excited state, we observe a sudden jump of the average transmon population with a distribution extending well above the top of the cosine potential. This is a clear illustration of the ionization of the transmon. For the drive amplitudes

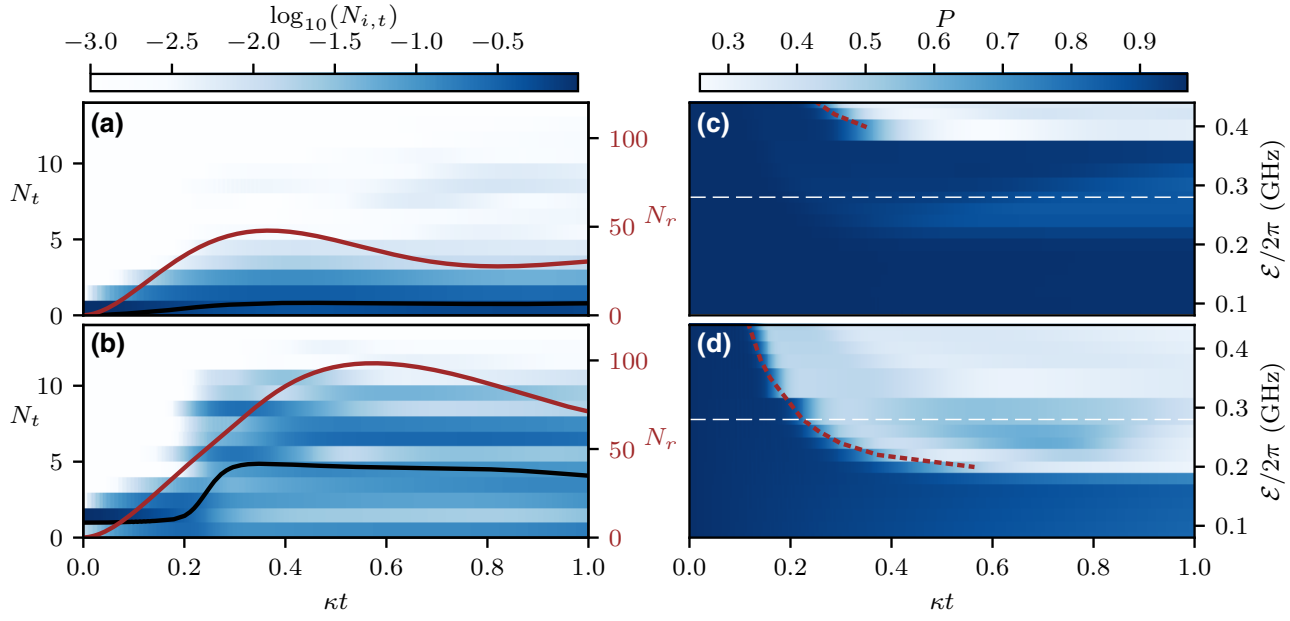

FIG. 2. (a),(b) The transmon and photon population as a function of time with $\mathcal{E}/2\pi = 0.28$ GHz for (a) the ground state and (b) the excited state. The black line indicates the average transmon population N_t (left axis) and the solid red line depicts the resonator average population N_r (right axis). The full distribution $\log_{10}(N_{i,t})$ of the transmon population for each transmon level $|i\rangle$ is also plotted, with the color bar above providing the scale. (c),(d) The purity of the reduced transmon density matrix as a function of the drive amplitude and time for the transmon initialized (c) in the ground state and (d) in the excited state. The dashed red lines indicate when the resonator reaches a population of 105 and 42 photons for the ground and excited states, respectively. The white dashed lines indicate the drive amplitude $\mathcal{E}/2\pi = 0.28$ GHz of (a) and (b). Because the drive amplitude is not sufficiently large to reach a population of $N_r = 105$ photons when the qubit is the ground state, ionization is not observed in (a).

considered here, the simulated measurement is far from QND and the dynamics are therefore not well described by a dispersive Hamiltonian [2].

As illustrated in Figs. 2(c) and 2(d), transmon ionization is accompanied by a sudden drop in the purity $P = \mathrm{Tr}[\hat{\rho}_t^2]$ of the reduced transmon density matrix. The rapid decline of the purity is observed at specific resonator-photon populations indicated by the dashed red line: $N_r = 105$ when initialized in the ground state in Fig. 2(c) and $N_r = 42$ for the excited state in Fig. 2(d). The dashed red lines terminate whenever the drive is too weak for the resonator to reach those populations. That transmon ionization occurs at specific photon populations suggests that the phenomenon is due to resonances, which we discuss in more detail below. These observations are compatible with Ref. [9], where steady-state calculations with an off-resonant drive have also shown drops in transmon purity in steady-state numerical calculations. Resonances at large photon numbers have also been observed in Ref. [14].

As a further illustration of transmon ionization at distinct photon numbers, we plot in Fig. 3(a) the transmon population N_t as a function of the resonator population N_r, both values taken from time traces such as shown in Figs. 2(a) and 2(b). The different curves correspond to different drive amplitudes and the initial state of the transmon is identifiable from the starting transmon population. Remarkably, the responses essentially collapse to single curves for each of the initial states. Below the ionizing photon population, the transmon population remains close to the computational manifold and the resonator population exhibits transient oscillations due to the drive being off resonant with the Lamb-shifted resonator frequency. Together with Purcell decay, these transient oscillations are responsible for the features observed at small N_r. Above the ionizing photon populations, the transmon population rapidly increases. This coincides with the sudden drop in the purity observed in Figs. 2(c) and 2(d).

Focusing now on the dynamics of the resonator, Fig. 4 shows the Wigner function of the resonator at different times with the transmon starting in its excited state and with the same drive amplitude $\mathcal{E}/2\pi = 0.28$ GHz as used in Figs. 2(a) and 2(b). For small features to appear more clearly, we plot the logarithm of the absolute value of the Wigner function. As expected, the cavity is initially well described by the single coherent states $|\alpha_1\rangle$ associated with the first excited state of the transmon [see Fig. 2(a)]. As the transmon population in other levels increases, additional features emerge corresponding to the coherent states associated with the now occupied transmon levels. Additionally, "bananization" caused by transmon-induced nonlinearities becomes apparent [30,31]. The dashed lines correspond to fixed photon numbers and are used as a guide to the eye. The full colored lines overlying the Wigner functions are obtained from a semiclassical approximation that we now introduce.

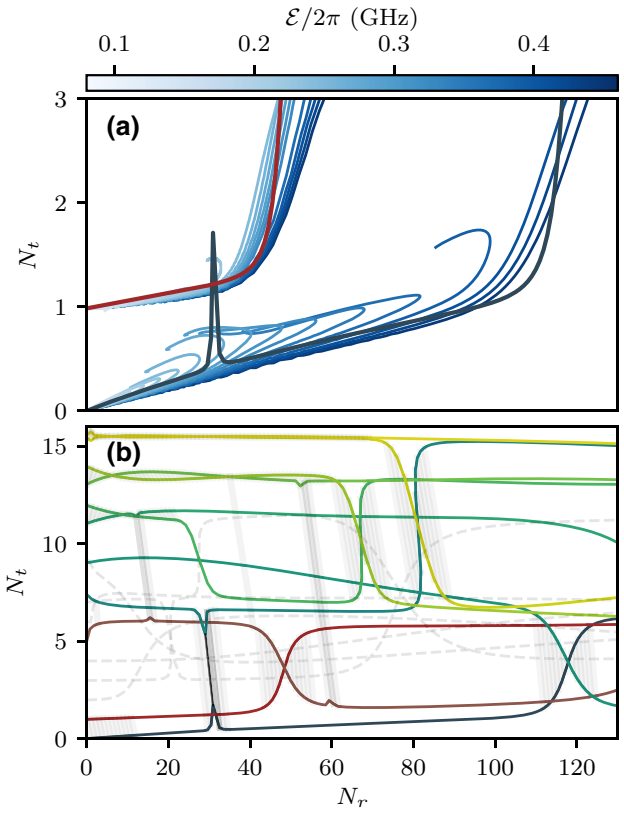

FIG. 3. (a) A parametric plot of the average transmon population versus the average resonator population for an evolution time of κ^{-1}. Each line represents a unique simulation with a different drive amplitude $\mathcal{E}$ and starting either in the ground or excited state. These metrics are calculated in the bare basis and are defined as in Sec. II B. (b) The average transmon population as a function of the bare-resonator population for each resonator branch $\{i\}$. Each branch is represented by a different color, with dashed gray lines indicating branches that do not feature significantly in the dynamics. The nearly vertical dark lines represent the overlap $C_i(n)$ between the different branches, darker lines indicating stronger overlap.

III. STATE IDENTIFICATION AND SEMICLASSICAL INTERPRETATION

The TPU-based large-scale numerical simulations presented in Sec. clearly illustrate the breakdown of the dispersive approximation. Notably, at specific resonator-photon populations, we observe a sudden jump in transmon population above the cosine-potential well, which is associated with a sharp drop in transmon purity. This results in complex dynamics of the resonator field, as illustrated by its Wigner function. In this section, we develop a semiclassical model to understand the main features of these observations.

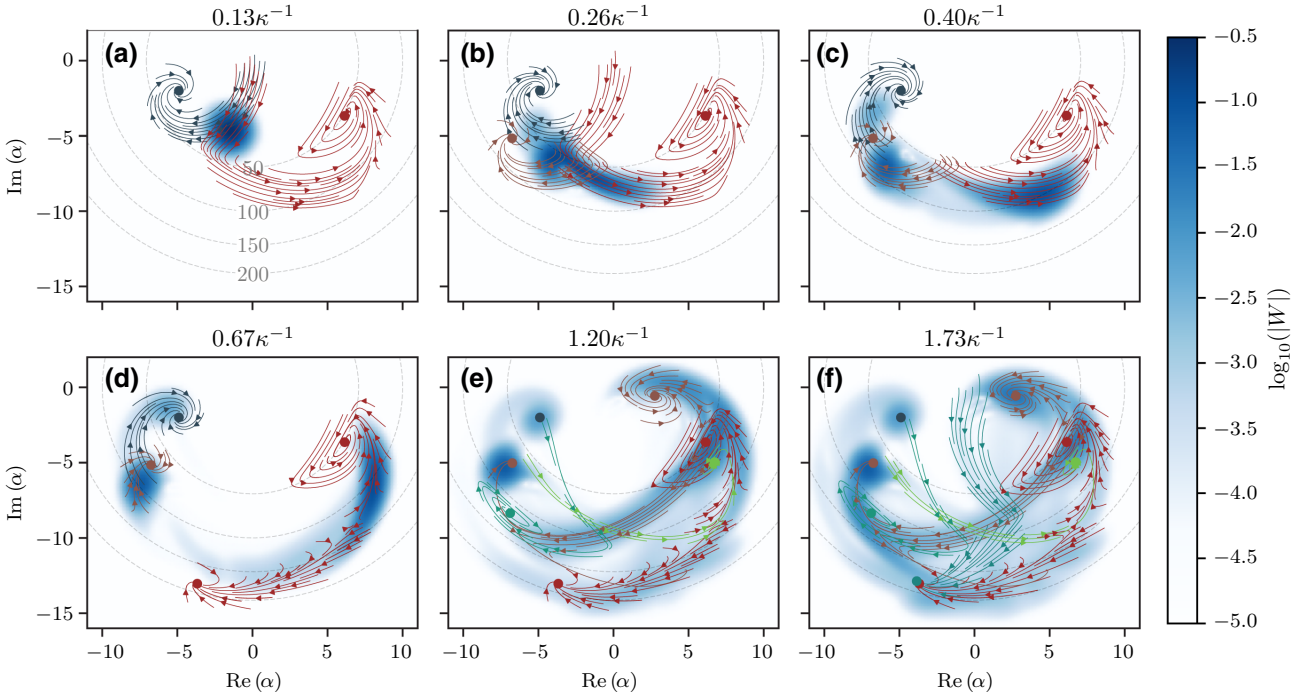

FIG. 4. The logarithm of the absolute value of the Wigner function $\log_{10}(|W|)$ of the resonator at different times for a drive amplitude $\mathcal{E}/2\pi = 0.28$ GHz and the transmon initialized in its dressed excited state. The logarithm is chosen to reveal the smaller features. Overlaid are the semiclassical flows of the relevant branches following the color convention of Fig. 3(b). The dashed circles are a guide to the eye at constant photon numbers (50, 100, 150, and 200).

A. Identification of dressed states and resonator branches

Our semiclassical approach is based on the dressed states $|\overline{i,n}\rangle$ and energies $E_{\overline{i,n}}$ of the transmon-resonator Hamiltonian in Eq. (1) in the absence of the drive. More precisely, at arbitrary photon number, we identify the dressed states that are closest to transmon eigenstates. At low photon number, this identification is simple in the dispersive regime. However, as the photon population approaches and even exceeds n_{crit}, the dressed states are highly entangled and identification becomes difficult.

Building on Ref. [13], our approach relies on first identifying the eigenstates of the Hamiltonian of Eq. (1) for $\mathcal{E} = 0$ obtained from numerical diagonalization with the largest overlap with $|i, 0\rangle$, the bare transmon states at zero photon population. Then, given an identified state $|\overline{i,n}\rangle$ for $n \geq 0$, the next state $|\overline{i,n+1}\rangle$ is chosen from the subset of remaining eigenstates $\{|\lambda\rangle\}$ such that the overlap

$$C_i(n) = \left| \frac{\langle \lambda | \hat{a}^\dagger | \overline{i,n} \rangle}{\langle \overline{i,n} | \hat{a}\hat{a}^\dagger | \overline{i,n} \rangle} \right|^2 \tag{4}$$

is maximized. Following this procedure recursively, we obtain a set of states $\{|\overline{i,n}\rangle\}$, where the bare transmon label i is fixed and n spans a desired range of resonator population. We refer to each such set of states as a branch $\{i\}$ of the resonator. In Fig. 3(b), we plot the average transmon population of the first 16 of these resonator branches as a function of the photon number. The full colored lines are branches that play an important role in the understanding of the numerical results. Branches playing a more minor role for our particular choice of parameters are illustrated as dashed gray lines. At finite transmon-resonator coupling and because we include non-RWA terms in the system Hamiltonian, the different branches have a nonzero overlap $C_i(n)$.

To illustrate this, in Fig. 3(b) we plot, with dark and nearly vertical lines, the overlap $C_i(n)$ between branches whenever it rises above 0.01. Darker lines indicate a stronger overlap between the states. At small photon number, the overlap between {0} and {1} corresponds to Purcell decay, which is seen to decrease with increasing photon number [32]. More importantly, we note a strong overlap between the branches {0} and {9} for $N_r \sim 110$ and between {1} and {5} for $N_r \sim 50$, in agreement with the photon numbers at which sudden drops of transmon purity are observed in Figs. 2(c) and 2(d) (red dashed lines). To further emphasize the link between the large overlaps and transmon ionization, we reproduce as solid blue and red lines the first two branches {0} and {1} of Fig. 3(b) together with the numerical data in Fig. 3(a). The sharp elbows, which are indicative of the fact that transmon

ionization results from a resonance, align well between the semiclassical results and the numerical data.

A strong overlap is also observed between {0} and {8} close to $N_r = 30$ [see Figs. 3(a)and 3(b)]. In contrast to what is observed at $N_r \sim 50$ and 110, this sharper resonance results in little population transfer, seen as a small drop of purity for $\mathcal{E} \in (0.2 - 0.3)/2\pi$ GHz in Fig. 2(c). This can be understood from the framework of Landau-Zener transitions. In our semiclassical picture, as the resonator rings up toward its transmon-state-dependent steady state, the photon number is swept at a rate related to κ and to the drive amplitude. A large drive amplitude leads to a rapid sweep through the narrow feature at $N_r \sim 30$, resulting in a diabatic passage with little leakage [see the darker-blue numerical lines in Fig. 3(b)]. We note that similar effects have been observed in the study of a resonator-induced phase gate, in which qubit-resonator leakage is maximized near avoided crossings [33]. On the other hand, for our choice of parameters, the features at $N_r \sim 50$ and 110 are broader and the sweep rate therefore comparatively slower. This slower sweep rate results in a nondiabatic process, leading to leakage out of the transmon computational subspace and to the observed drop in purity. This interpretation is further confirmed in Appendix C for numerical experiments with a different drive frequency and loss rates. In the steady state, the above concepts are not applicable and ionization is determined only by the chosen parameters and drive amplitude [9].

B. Semiclassical dynamics

Following Ref. [13], we assign to each branch $\{i\}$ an effective oscillator of photon-number dependent frequency

$$\omega_i(n) = E_{\overline{i,n+1}} - E_{\overline{i,n}}. \tag{5}$$

As seen in Fig. 5, where we plot $\omega_i(n)$ versus n, Eq. (5) accounts for the photon-number dependence of the resonator frequency, including that the resonator responds at its bare frequency at large photon numbers. Assuming that the transmon remains in a given branch $\{i\}$, the dynamics of this effective oscillator approximately obey the classical equation of motion of a driven damped oscillator,

$$\dot{\alpha_i} = -i[\omega_i(|\alpha_i|^2) - \omega_d]\alpha_i - i\mathcal{E}/2 - \kappa\alpha_i/2, \tag{6}$$

where nonlinear effects are encapsulated in the dependence of the branch frequency on the photon number $\omega_i(|\alpha_i|^2)$ according to Eq. (5). Because the quantity $|\alpha_i|^2$ takes arbitrary real values, we generalize Eq. (5) from discrete values n to a continuous function. To do this, we smooth $\omega_i(n)$ with a first-order Savitzky-Golay filter and linearly interpolate between each n. Doing this additionally removes large discontinuities in the photon-number dependence of the effective frequencies $\omega_i(n)$ caused by strong interactions with other branches. This semiclassical approximation is expected to accurately describe the system for large photon number $|\alpha_i|^2 \gg 1$.

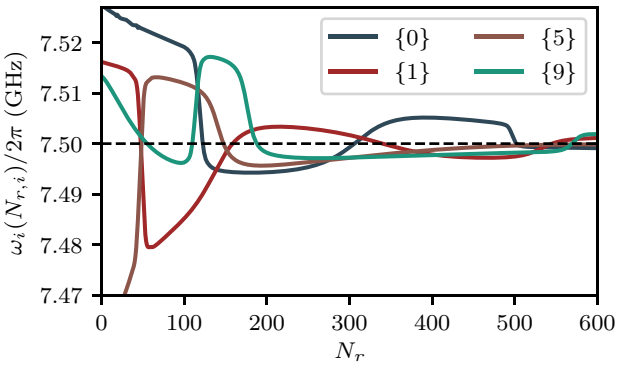

FIG. 5. The effective frequency $\omega_i(N_r)$ of the resonator branches as a function of the resonator population. The frequency is smoothed using a Savitzky-Golay filter to remove discontinuities caused by strong interactions with other branches. The dashed line indicates the bare-cavity frequency $\omega_r/2\pi = 7.5$ GHz.

Numerically integrating Eq. (6), we plot in Fig. 4(a) the time dependence of $\alpha_1(t)$ (red lines) together with a snapshot of the Wigner function at time $t = 0.13\kappa^{-1}$ assuming that the effective oscillator starts in vacuum. The arrows indicate the flow of time and the different lines are obtained from Eq. (6) for a set of initial conditions in an area close to vacuum representing zero point fluctuations. The red dot corresponds to the steady-state value α_1^s.

To account for transitions between transmon states, the evolution of $\alpha_i(t)$ associated with other branches is illustrated whenever the semiclassical evolution reaches photon numbers at which we anticipate non-negligible rates for transitions into those branches according to the overlaps observed in Fig. 3(b). To distinguish between the different branches, we use the color scheme of Fig. 3(b). For example, the blue lines in Figs. 4(a)–4(d) correspond to branch {0}, which appears as a result of Purcell decay. In the same way, starting in Fig. 4(b), we include the flow of $\alpha_5(t)$ associated with branch {5} (brown lines), which has a strong overlap with {1} when $N_r \sim 50$. Following the evolution of the Wigner function from panel to panel, it is possible to see a feature following the flow of $\alpha_5(t)$ and settling at the expected steady-state value α_5^s in Fig. 4(d) (brown dot). Applying this procedure whenever the resonator number is such that an occupied branch has a strong overlap $C_i(n)$ with another branch, the vast majority of the Wigner-function transient behavior can be understood. In particular, starting in Fig. 4(d), we see the bistable behavior of $\{\alpha_1\}$ in both the semiclassical results (red lines) and the numerical data.

Rather than focusing on the steady-state response, in Fig. 6(a) we plot, for different drive amplitudes and as a

function of time, the difference in the resonator populations ΔN_r obtained by solving the semiclassical expression in Eq. (6), given that the transmon is initialized in its ground or excited state. As expected from Fig. 5, at some threshold power, the photon number rapidly increases if the qubit is initially in its excited state, while the increase is not as pronounced for ground state. This results in the observed large ΔN_r. Figure 6(b) shows the same quantity obtained from numerical simulations. Given that the population of the resonator $N_r \geq 450$ for the strongest drive amplitude $\mathcal{E}/2\pi = 440$ MHz, these simulations are particularly demanding and require Hilbert-space sizes up to 2^{15}. The agreement between the full numerical simulations and the simple semiclassical model is remarkable. In particular, in both approaches, we observe that the photon-number difference ΔN_r goes to zero at the strongest drive amplitudes. This is expected from Fig. 5, where at very large photon numbers, the frequencies $\omega_i(|\alpha_i|^2)$ eventually collapse to the bare value ω_r. The different drive amplitude at which this collapse occurs depend on the initial transmon state and the resulting large ΔN_r is exploited in the high-power qubit readout [11–13]. The choice of drive amplitude range $\mathcal{E}/2\pi \in [0, 440]$ MHz for the TPU simulations is made to demonstrate the full range of the high-power readout behavior and demonstrate the decrease in ΔN_r at the largest amplitudes.

A possible interpretation for the observed large response at the bare-resonator frequency is that once the transmon is ionized, mostly chargelike states are occupied. Because these states couple longitudinally to the resonator, they do not lead to a resonator frequency pull [10]. However, the numerical results show that even for the highest drive powers, a significant distribution of states inside the well remain populated (not shown). As a result, the collapse of the resonator to its bare frequency cannot be explained as resulting alone from the longitudinal-type coupling of the unbound states. We leave a more detailed analysis of this effect to future work [29].

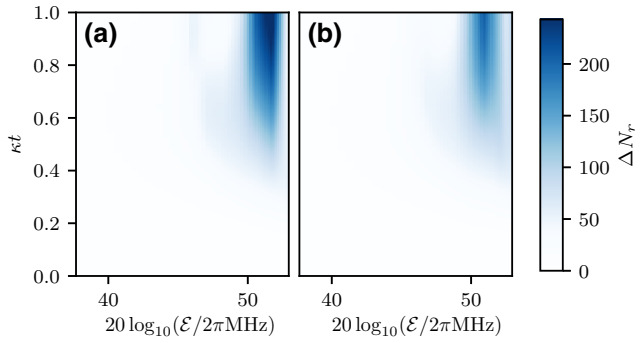

FIG. 6. The difference in the resonator population ΔN_r given that the transmon is initialized in its ground or excited state as a function of the drive amplitudes and time as obtained from (a) the semiclassical model and (b) the numerical TPU data.

C. Resonances at low photon number

For the parameters used above, the first resonance leading to non-QNDness occurs at a relatively large photon number of approximately 30, followed by resonances at even larger photon numbers. Such resonances can, however, occur at much lower photon numbers. To demonstrate this, we now use parameters based on the experiment of Ref. [4]: $E_C/h = 314$ MHz, $E_J/E_C = 55.47$, $g/2\pi = 211$ MHz, and $w_r/2\pi = 4.804$ GHz, corresponding to $n_{\text{crit}} \approx 10$. The cavity damping rate is set to $\kappa/2\pi = 40$ MHz. The drive amplitude is varied in the range $\mathcal{E}/2\pi \in [0.02, 0.14]$ MHz and the evolution time is set to 48 ns, corresponding to the smallest measurement time used in Ref. [4]. We note that these are bare system parameters chosen such as to approach the dressed parameters reported in Ref. [4].

In a similar fashion to Fig. 3(a), in Fig. 7(a) we parametrically plot the transmon population N_t against the resonator population N_r for the transmon initialized in the excited state. For this choice of parameters and integration time, Purcell decay is apparent for the three lowest drive amplitudes (see the sharp decrease of N_t). For larger

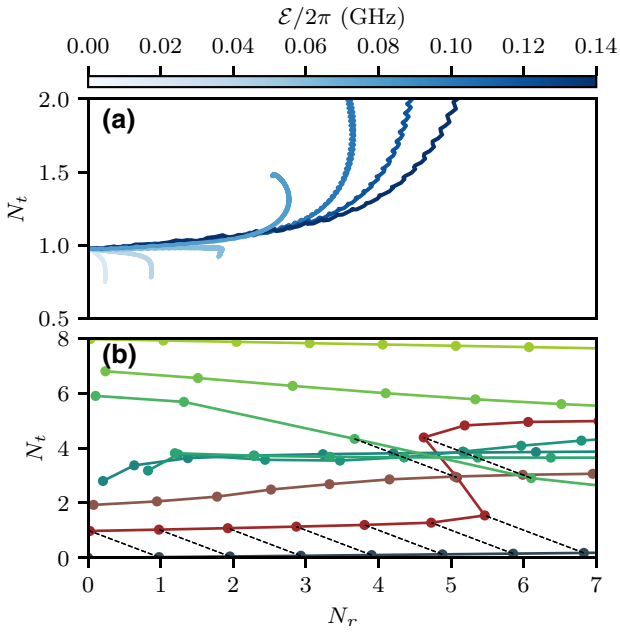

FIG. 7. (a) A parametric plot of the average transmon population versus the average resonator population for an evolution time of 48 ns. Each line represents a unique simulation with a different drive amplitude $\mathcal{E}$ and starting in the excited state. Because of the small Hilbert-space size, these results are obtained using CPUs. (b) The average transmon population as a function of the bare-resonator population for each resonator branch $\{i\}$. The dashed lines indicate the overlap between different branches when $C_i(n) > 0.01$.

measurement amplitudes corresponding to an average resonator population of approximately 2.5 photons, upward transition of the transmon population is clearly observed. Interestingly, this corresponds to the average photon number at which a decrease in measurement fidelity is observed in Ref. [4] [see Fig. 3(b) of that reference].

To understand the origin of this population leakage under measurement, we show in Fig. 7(b) the average photon population for the resonator branches $\{i\}$ as obtained from our semiclassical model. At $N_r \approx 5$, the transmon population rapidly rises for branch $\{1\}$ associated with the first excited state of the transmon (red line). In the numerical simulation shown in Fig. 7(a) as well as in experiments, the resonator field is in a coherent state. As a result, because of the $\sqrt{n}$ fluctuations of coherent states, the resonance at $N_r \approx 5$ already plays a role at $N_r \approx 2.5$, leading to the observed non-QNDness. While care must be taken when using a semiclassical model at such a small photon population, these observations are suggestive of the fact that the resonances that are observed here could play a role in the drop of measurement fidelity and QNDness that is experimentally observed at low photon number.

IV. CONCLUSIONS

We leverage the computational power of TPUs to perform large Hilbert-space size time-dependent master-equation simulations of transmon readout in circuit QED. From these simulations, we identify resonances occurring at a distinct resonator-photon number where the transmon population escapes to states above the Josephson cosine-potential well. To interpret these results, we develop a semiclassical model capturing the nonlinear transmon-state-dependent change of the resonator frequency with photon number. In particular, this model correctly captures which states play a significant role in transmon ionization. Using a different set of parameters, we show that these resonances can occur at small photon population. These results suggest that the non-QND nature of the dispersive readout experimentally observed at small drive amplitudes could be due to these resonances. The location of ionization events can be predicted from the semiclassical branch analysis, thus providing an upper bound on the maximum population of the measurement photons to maintain the QND character of the readout. Moreover, because the location of these ionizing events is strongly parameter dependent, this work suggests a pathway to avoid these spurious effects when optimizing device parameters for readout.

ACKNOWLEDGMENTS

We thank Cristóbal Lledó and Adrien Moulinas for insightful discussions. This work was undertaken in part thanks to funding from the Natural Sciences and Engineering Research Council (NSERC), the Canada First Research Excellence Fund, the Ministère de l'Économie et de l'Innovation du Québec, the U.S. Army Research Office under Grant No. W911NF-18-1-0411, and the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers and Quantum Systems Accelerator. This research was also supported by Cloud TPUs from Google's TPU Research Cloud (TRC). G.V. is a Distinguished Invited Professor at the Institute of Photonic Sciences (ICFO) and a Distinguished Visiting Research Chair at Perimeter Institute. Research at Perimeter Institute is supported by the Government of Canada through the Department of Innovation, Science and Economic Development and by the Province of Ontario through the Ministry of Research, Innovation and Science.

APPENDIX A: NUMERICAL IMPLEMENTATION

To simulate the master equation in Eq. (2), we move to a rotating frame defined by

$$\hat{U}_{\rm rf}(t) = \exp(i\omega_d \hat{a}^\dagger \hat{a} t). \tag{A1}$$

In that frame, the system Hamiltonian takes the form

$$\begin{aligned}\hat{H}_I(t) &= \hat{U}_{\rm rf}(t)\hat{H}(t)\hat{U}^\dagger_{\rm rf}(t) - i\hat{U}_{\rm rf}(t)\dot{\hat{U}}^\dagger_{\rm rf}(t),\\ &= \hat{H}_{I0} + \hat{X}_I(t),\end{aligned} \tag{A2}$$

with

$$\hat{H}_{I0} = 4E_c\hat{n}_t^2 - E_J\cos(\hat{\varphi}_t) + (\omega_r - \omega_d)\hat{a}^\dagger\hat{a} - \frac{\mathcal{E}}{2}(\hat{a}^\dagger + \hat{a}) \tag{A3}$$

and

$$\hat{X}_I(t) = -ig\hat{n}_t(\hat{a}^\dagger e^{i\omega_d t} - \hat{a}e^{-i\omega_d t}) + \frac{\mathcal{E}}{2}(\hat{a}^\dagger e^{2i\omega_d t} + \hat{a}e^{-2i\omega_d t}). \tag{A4}$$

Moreover, in this rotating frame the master equation reads

$$\dot{\hat{\rho}}_I(t) = -i[\hat{H}_I(t), \hat{\rho}_I] + \kappa\mathcal{D}[\hat{a}]\hat{\rho}_I \equiv \hat{\mathcal{L}}\hat{\rho}_I, \tag{A5}$$

with $\hat{\rho}_I(t) = \hat{U}_{\rm rf}(t)\hat{\rho}(t)\hat{U}^\dagger_{\rm rf}(t)$ and where the dissipator $\mathcal{D}[\hat{a}]$ is unaffected by the transformation.

To solve this master equation on TPUs, we approximate the action of the time-ordered Lindbladian exponential $\mathcal{T}\exp\left(\int_0^t \hat{\mathcal{L}}(t')dt'\right)$. To do this, we first solve for the roots $\{z_i\}$ of the Taylor approximation of the exponential to nth order and rearrange the terms:

$$e^x \approx \sum_{i=0}^{n}\frac{x^i}{i!} = \frac{1}{n!}\prod_{i=1}^{n}(x - z_i) = \prod_{i=1}^{n}(1 - x/z_i), \tag{A6}$$

where $\prod_i z_i = n!$. We then evolve the system by time δt by iterating through the n roots:

$$\hat{\rho}(t+\delta t) = \hat{\rho}_n(t), \quad \hat{\rho}_{i=0}(t) = \hat{\rho}(t)$$
$$\hat{\rho}_i(t) = \left[1 - (\delta t/z_i)\overline{\hat{\mathcal{L}}(t, t+\delta t)}\right]\hat{\rho}_{i-1}(t), \tag{A7}$$

where $\overline{\hat{\mathcal{L}}(t, t+\delta t)}$ is calculated using a second-order Magnus expansion. The order n and step size δt required for convergence depend on the number of transmon and resonator states, with values of $n = 8 - 15$ and $1/(\omega_d \delta t) \approx 50$–$100$ found to be optimal for our choice of parameters.

The Hilbert space of the cavity is adapted throughout the evolution. At each step, the quantity $E = |1 - \mathrm{Tr}\{\rho(N\delta t)[\hat{a}, \hat{a}^\dagger]\}|$ is calculated to ensure minimal occupation of the highest cavity excited states. The cavity size is doubled if $E > 10^{-6}$, with the previous step recalculated with the larger Hilbert-space size once this condition is met.

In the case of a transmon-resonator system driven *off resonantly*, a simple frame change can be used to follow the linear response of the oscillator, subsequently reducing the required Hilbert space [9]. In the case of a resonant drive, such a transformation does not exist.

APPENDIX B: TENSOR PROCESSING UNITS

Google's TPUs are application-specific integrated circuits (ASICs) originally designed to accelerate and scale up machine-learning workloads [15]. By leveraging the JAX library [34], it is possible to repurpose TPUs to also accelerate other large-scale computational tasks [16–28]. For instance, in Refs. [24–27], TPUs are used to simulate the wave function of up to 36–40 (two-level) qubits. In this work, we use the power of TPUs to simulate, instead, the time evolution of the joint density matrix of a transmon-resonator system, as given by a Lindblad master equation.

We employ TPUs of the third generation, denoted v3. Each single TPU v3 core is equipped with two matrix-multiply units (MXUs) to formidably accelerate matrix-matrix multiplication (matmul), resulting in about ten teraflops (TFLOPs) of measured single-core matmul performance in single precision.

The smallest available TPU configuration consists of eight TPU v3 cores with a total of 128 GB of dedicated high-bandwidth memory (HBM), controlled by a single host with 48 CPU cores. The largest configuration is a pod with 2048 TPU v3 cores and 32 TB of HBM, controlled by 256 hosts. Given a choice of configuration, the available TPU cores are directly connected to nearest neighbors in a two-dimensional torus network through fast inter-core interconnects (ICIs). The ICIs play an essential role in the ability of the TPUs to maintain high performance when distributing matmuls and other dense linear-algebra operations over all available TPU cores. In this work, we use the JAX library [34] to write single-program multiple data (SPMD) code and execute it on configurations made of multiple TPU cores. Specifically, for the largest density matrix under consideration, of size $2^{15} \times 2^{15}$, we use 128 TPU v3 cores. The density matrix is distributed over all available cores and updated according to the different terms in the Lindblad operator, with a typical update $\hat{\rho}(t) \to \hat{\rho}(t+\delta t)$ for $n = 15$ in Eq. (A7) taking on the order of seconds.

APPENDIX C: LANDAU-ZENER TRANSITIONS FOR DIFFERENT PARAMETER CHOICES

Figure 8(a) shows the average transmon population versus the average resonator population for the same

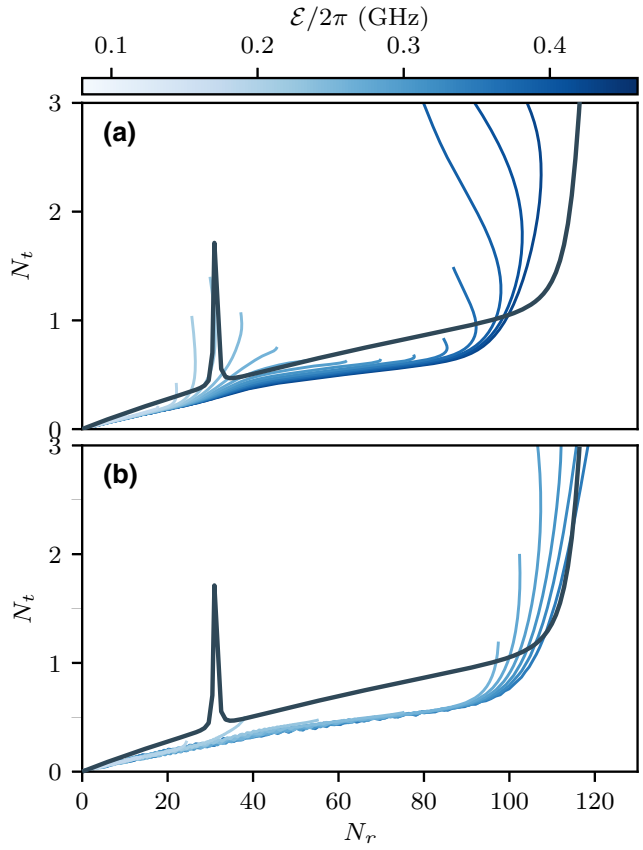

FIG. 8. Parametric plots of the average transmon population versus the average resonator population for an evolution time of κ^{-1} for (a) $\kappa/2\pi = 20$ MHz as in the main text and (b) $\kappa/2\pi = 80$ MHz. Each line represents a unique simulation with a different drive amplitude $\mathcal{E}$ starting in the ground state. The frequency of the drive is set to the mean of the pulled resonator frequencies associated with the ground and excited states of the transmon, as is often the case in dispersive readout. For the larger κ, the system sweeps through the resonance $N_r \sim 30$ at a faster rate. The transition becomes more diabatic and the leakage through other states is reduced, as expected from the Landau-Zener theory.

parameters as in the main text, except for a slightly different drive frequency corresponding to the mean of the pulled resonator frequencies associated with the ground and excited states of the transmon. This different drive frequency only leads to a small quantitative change with respect to Fig. 3. More importantly, Fig. 8(b) is obtained with that same drive frequency but now a larger decay rate $\kappa/2\pi = 80$ MHz. Here, the resonance around $N_r \sim 30$ mentioned in Sec. III A leads to smaller leakage. This is compatible with Landau-Zener theory, since with a larger κ the photon number rises more rapidly, corresponding to a faster sweep through this resonance and therefore a more diabatic transition.

In addition, we note that the numerically observed slow increase of the transmon population in Fig. 8(b) is overestimated by the semiclassical model. This is not the case for Fig. 3(a), where the drive frequency is set to the bare-resonator frequency. We suspect that the following mechanism is responsible for this discrepancy: for the drive frequency used in Fig. 8, we expect the cavity state to evolve to two very distinct coherent states depending on whether one initializes the transmon in the ground or the excited states. As has been highlighted in Leroux *et al.* [35], the distance between the resulting polaronic states is likely to diminish mixing of the states. This effect is, however, not accounted for by the semiclassical model, since the branches $\{i\}$ are derived from the static spectrum, where the drive frequency is not involved.

[1] A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, J. Majer, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Approaching Unit Visibility for Control of a Superconducting Qubit with Dispersive Readout, Phys. Rev. Lett. **95**, 060501 (2005).

[2] A. Blais, A. L. Grimsmo, S. M. Girvin, and A. Wallraff, Circuit quantum electrodynamics, Rev. Mod. Phys. **93**, 025005 (2021).

[3] 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, Phys. Rev. A **76**, 042319 (2007).

[4] T. Walter, P. Kurpiers, S. Gasparinetti, P. Magnard, A. Potočnik, Y. Salathé, M. Pechal, M. Mondal, M. Oppliger, C. Eichler, and A. Wallraff, Rapid High-Fidelity Single-Shot Dispersive Readout of Superconducting Qubits, Phys. Rev. Appl. **7**, 054020 (2017).

[5] Z. K. Minev, S. O. Mundhada, S. Shankar, P. Reinhold, R. Gutiérrez-Jáuregui, R. J. Schoelkopf, M. Mirrahimi, H. J. Carmichael, and M. H. Devoret, To catch and reverse a quantum jump mid-flight, Nature **570**, 200 (2019).

[6] M. Boissonneault, J. M. Gambetta, and A. Blais, Dispersive regime of circuit QED: Photon-dependent qubit dephasing and relaxation rates, Phys. Rev. A **79**, 013819 (2009).

[7] M. Malekakhlagh, A. Petrescu, and H. E. Türeci, Lifetime renormalization of weakly anharmonic superconducting qubits. I. Role of number nonconserving terms, Phys. Rev. B **101**, 134509 (2020).

[8] A. Petrescu, M. Malekakhlagh, and H. E. Türeci, Lifetime renormalization of driven weakly anharmonic superconducting qubits. II. The readout problem, Phys. Rev. B **101**, 134510 (2020).

[9] L. Verney, R. Lescanne, M. H. Devoret, Z. Leghtas, and M. Mirrahimi, Structural Instability of Driven Josephson Circuits Prevented by an Inductive Shunt, Phys. Rev. Appl. **11**, 024003 (2019).

[10] R. Lescanne, L. Verney, Q. Ficheux, M. H. Devoret, B. Huard, M. Mirrahimi, and Z. Leghtas, Escape of a Driven Quantum Josephson Circuit into Unconfined States, Phys. Rev. Appl. **11**, 014030 (2019).

[11] M. D. Reed, L. DiCarlo, B. R. Johnson, L. Sun, D. I. Schuster, L. Frunzio, and R. J. Schoelkopf, High-Fidelity Readout in Circuit Quantum Electrodynamics Using the Jaynes-Cummings Nonlinearity, Phys. Rev. Lett. **105**, 173601 (2010).

[12] L. S. Bishop, E. Ginossar, and S. M. Girvin, Response of the Strongly Driven Jaynes-Cummings Oscillator, Phys. Rev. Lett. **105**, 100505 (2010).

[13] M. Boissonneault, J. M. Gambetta, and A. Blais, Improved Superconducting Qubit Readout by Qubit-Induced Nonlinearities, Phys. Rev. Lett. **105**, 100504 (2010).

[14] D. Sank *et al.*, Measurement-Induced State Transitions in a Superconducting Qubit: Beyond the Rotating Wave Approximation, Phys. Rev. Lett. **117**, 190503 (2016).

[15] N. Jouppi, D. Yoon, G. Kurian, S. Li, N. Patil, J. Laudon, C. Young, and D. Patterson, A domain-specific supercomputer for training deep neural networks, Commun. ACM **63**, 67 (2020).

[16] T. Lu, T. Marin, Y. Zhuo, Y.-F. Chen, and C. Ma, in *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)* (IEEE, Nice, France, 2021), p. 783.

[17] T. Lu, Y.-F. Chen, B. Hechtman, T. Wang, and J. Anderson, Large-scale discrete Fourier transform on TPUs, IEEE Access **9**, 93422 (2021).

[18] F. Huot, Y.-F. Chen, R. Clapp, C. Boneti, and J. Anderson, High-resolution imaging on TPUs, arXiv:1912.08063 [cs.CE] (2019).

[19] F. Belletti, D. King, K. Yang, R. Nelet, Y. Shafi, Y.-F. Shen, and J. Anderson, in *Proceedings of the 2020 SIAM Conference on Parallel Processing for Scientific Computing (PP)* (SIAM, Seattle, Washington, USA, 2020), p. 12.

[20] T. Lu, T. Marin, Y. Zhuo, Y.-F. Chen, and C. Ma, in *2020 IEEE High Performance Extreme Computing Conference (HPEC)* (IEEE, Waltham, MA, USA, 2020), p. 1.

[21] Q. Wang, M. Ihme, Y.-F. Chen, and J. Anderson, A TensorFlow simulation framework for scientific computing of fluid flows on tensor processing units, Comput. Phys. Commun. **274**, 108292 (2022).

[22] A. G. M. Lewis, J. Beall, M. Ganahl, M. Hauru, S. B. Mallick, and G. Vidal, Large-scale distributed linear algebra with tensor processing units, Proc. Natl. Acad. Sci. **119**, e2122762119 (2022).

[23] R. Pederson, J. Kozlowski, R. Song, J. Beall, M. Ganahl, M. Hauru, A. G. M. Lewis, S. B. Mallick, V. Blum, and G. Vidal, Tensor processing units as quantum chemistry supercomputers, arXiv:2202.01255 [physics] (2022).

[24] A. Morningstar, M. Hauru, J. Beall, M. Ganahl, A. G. M. Lewis, V. Khemani, and G. Vidal, Simulation of Quantum Many-Body Dynamics with Tensor Processing Units: Floquet Prethermalization, PRX Quantum **3**, 020331 (2022).

[25] M. Hauru, A. Morningstar, J. Beall, M. Ganahl, A. Lewis, and G. Vidal, Simulation of quantum physics with tensor processing units: Brute-force computation of ground states and time evolution, arXiv:2111.10466 [quant-ph] (2021).

[26] E. Gustafson, B. Holzman, J. Kowalkowski, H. Lamm, A. C. Y. Li, G. Perdue, S. V. Isakov, O. Martin, R. Thomson, J. Beall, M. Ganahl, G. Vidal, and E, Peters, in *2021 IEEE/ACM Second International Workshop on Quantum Computing Software (QCS)* (2021), p. 72.

[27] M. Ganahl *et al.*, Tensor processing units for simulating quantum circuits (to be published).

[28] M. Ganahl, J. Beall, M. Hauru, A. G. M. Lewis, J. H. Yoo, Y. Zou, and G. Vidal, Density matrix renormalization group with tensor processing units, arXiv:2204.05693 [cond-mat.str-el] (2022).

[29] J. Cohen, A. Petrescu, R. Shillito, and A. Blais, Reminiscence of classical chaos in driven transmons, arXiv:2207.09361 [quant-ph] (2022).

[30] S. Boutin, D. M. Toyli, A. V. Venkatramani, A. W. Eddins, I. Siddiqi, and A. Blais, Effect of Higher-Order Nonlinearities on Amplification and Squeezing in Josephson Parametric Amplifiers, Phys. Rev. Appl. **8**, 054030 (2017).

[31] V. V. Sivak, N. E. Frattini, V. R. Joshi, A. Lingenfelter, S. Shankar, and M. H. Devoret, Kerr-Free Three-Wave Mixing in Superconducting Quantum Circuits, Phys. Rev. Appl. **11**, 054060 (2019).

[32] E. A. Sete, J. M. Gambetta, and A. N. Korotkov, Purcell effect with microwave drive: Suppression of qubit relaxation rate, Phys. Rev. B **89**, 104516 (2014).

[33] M. Malekakhlagh, W. Shanks, and H. Paik, Optimization of the resonator-induced phase gate for superconducting qubits, Phys. Rev. A **105**, 022607 (2022).

[34] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, JAX: Composable transformations of PYTHON+NumPy programs, (2018).

[35] C. Leroux, A. Di Paolo, and A. Blais, Superconducting Coupler with Exponentially Large On:Off Ratio, Phys. Rev. Appl. **16**, 064062 (2021).

5.4 Gate charge dependence

One of the intended goals of the above manuscript was to explore the dependence of gate charge on the ionization physics seen in the results. Unfortunately, due to time constraints and limited TPU runtime, this was not possible during the internship. Nevertheless, after this publication I had the opportunity to explore some of this physics in a modified parameter regime where hardware acceleration is not strictly needed. I provide some of these insights in this final section.

The transmon is referred to as a 'charge-insensitive' qubit since its energy spectrum becomes insensitive to the offset gate charge with increasing values of the ratio E_J/E_C. The dispersion of the m-th energy level is defined as [38]:

$$\epsilon_m \equiv E_m(n_g = 0.5) - E_m(n_g = 0). \tag{5.8}$$

For the transmon qubit, a ratio $E_J/E_c > 50$ results in negligible dispersion of the computational states – indeed, in the limit $E_J/E_c \gg 1$, Koch et. al. in Ref. [38] showed:

$$\epsilon_m \approx (-1)^m E_c \frac{2^{4m+5}}{m!}\sqrt{\frac{2}{\pi}}\left(\frac{E_J}{2E_c}\right)^{\frac{m}{2}+\frac{3}{4}} e^{-\sqrt{8E_J/E_c}}, \tag{5.9}$$

which shows the exponential suppression of the charge dispersion as a function of the ratio E_J/E_c. In the publication reproduced in Sect. 5.3, the master equation simulations of the transmon-resonator system were only conducted with an offset gate charge $n_g = 0$ due to limits on the computational time. In this section, I explore in more detail the effect of this offset gate charge, which was predicted in Ref. [63] to have significant effects in the presence of a strong periodic drive.

We begin by investigating how the eigenenergies of the transmon-resonator system depend on the offset gate charge. In Fig. 5.4 (a) I plot the energies of the dressed transmon eigenstates $|\overline{0,0}\rangle, |\overline{1,0}\rangle, |\overline{2,0}\rangle$ as a function of this offset gate charge. As expected, we observe almost negligible dispersion, since these states remain close to the charge-protected transmon state at the bottom of the cosine well. However, this story changes dramatically with (b) 22 photons and (c) 44 photons, where we observe significant changes in the energies of eigenstates with non-zero resonator populations. This indicates that the dressed states with large photon numbers have a significant hybridization with the higher energy transmon states close to and outside of the potential well, which are sensitive to charge, and that ionization may potentially occur at these photon numbers. We further note that the smallest and largest energies do not necessarily occur at $n_g = 0$ and $n_g = 0.5$. Motivated by these

observations, I adapt the expression for calculating the charge dispersion in Eq. (5.8):

$$\epsilon_{i,n} \equiv \max\{E_{\overline{i,n}}(n_g)\} - \min\{E_{\overline{i,n}}(n_g)\}, \tag{5.10}$$

where the maximium and minimum values are taken to be over the range of n_g. Using this new expression, I plot the charge dispersion as a function of the photon number n for the first branches in Fig. 5.4(d) (transparent lines). For comparison, I additionally plot with solid lines the charge dispersion calculated from the difference in energies at $n_g = 0.5$ and $n_g = 0$, similarly to Eq. (5.8), but now including the resonator index:

$$\epsilon'_{i,n} \equiv E_{\overline{i,n}}(n_g = 0.5) - E_{\overline{i,n}}(n_g = 0). \tag{5.11}$$

For comparison, the charge dispersion of the isolated transmon states (in the absence of any resonator coupling, $g = 0$) is plotted, see dashed lines. In agreement with Ref. [63], we see that the transmon-resonator system becomes *sensitive* to offset gate charge fluctuations even at modest photon numbers, potentially exposing the system to dephasing due to charge noise.

5.4.1 Master Equation Simulations – Gate Charge

A key feature of the charge dispersions in Fig. 5.4 are discontinuities as a function of the photon number – this means that the resonances between branches are dependent on this offset gate charge. Given these resonances can ultimately result in the qubit ionizing (see Ref. [2]), it would be informative to determine whether varying the gate charge significantly alters the conclusions reached so far.

Similarly to Sect. 5.2.4, we integrate the master equation in Eq. (5.7) for a duration of 100 ns, with a drive amplitude $\mathcal{E}/2\pi \approx 140$ MHz, yielding a steady-state resonator population of $\langle \hat{a}^\dagger \hat{a} \rangle \approx 22$. We then turn the drive off and allow the resonator to relax for another 100 ns. Two offset gate charges are considered – $n_g = 0$ and $n_g = 0.34$. Fig. 5.5(a) shows a parametric plot of the ground state branch population against the resonator population, beginning at $\langle \tilde{\Pi}_g \rangle = 1$ and with 0 resonator population, with arrows indicating the flow of time. For the $n_g = 0$ trajectory, there is a small decrease in the ground state branch population from unity as the resonator reaches the steady state. This drop-off appears to begin at a resonator population $\langle \hat{a}^\dagger \hat{a} \rangle \approx 19$, as indicated by the blue marker. This is followed by an abrupt drop-off corresponding to the drive being turned off. In stark contrast, for the $n_g = 0.34$ trajectory, there is a significant depopulation of the ground-state branch when the resonator reaches $\langle \hat{a}^\dagger \hat{a} \rangle \approx 11$, with up to 5% of the population lost to another branch. A

considerable amount of this population returns when the drive is turned off – nevertheless, once the resonator has returned to vacuum, the final population in the ground-state branch is only 99%. This deviation cannot be trivially explained by state identification, since the charge dispersion for the ground state $|\overline{i,0}\rangle$ should be exponentially small (see Fig. 5.4(a)), and thus virtually independent of gate charge when the resonator is in vacuum.

This difference can be explained by gate-charge-dependent resonances between branches. In Fig. 5.5(c), I plot the overlap metric $C_g(n)$ as defined in Eq. (5.6) as a function of the resonator population for both $n_g = 0$ and $n_g = 0.34$. At $n_g = 0$, this overlap remains very close to unity, with a small resonance occurring at $\langle \hat{a}^\dagger \hat{a} \rangle \approx 30$ photons. For $n_g = 0.34$, we note a strong resonance at $\langle \hat{a}^\dagger \hat{a} \rangle \approx 20$ photons, suggesting hybridization with another branch.

The impact of these resonances occurs well before the mean photon number in the resonator reaches the location of the resonance, which can be explained by the finite photon-number variance of the coherent state entering the resonator when the drive is applied. In Fig. 5.5(a), I plot the Poisson distribution corresponding to two coherent states with mean photon numbers $|\alpha|^2 = 11$ and 19, corresponding respectively to the photon numbers at which we see significant deviations from unity in the ground state branch population [see blue and orange circles in (a) and (b)]. The dashed lines indicate the 99% threshold for these distributions, – in other words, the higher energy states contribute to 1% of the population. These align almost perfectly with the resonances observed in (c), indicating that it is indeed these resonances likely causing the dramatic changes in behaviour and the deviation in final populations. This also suggests that only a small occupation probability of the resonator at inter-branch resonances is sufficient to significantly alter the branch population.

Drift in the external gate charge can occur over *long* time scales on the order of minutes or longer [96, 97]. Recent experimental data demonstrated that the readout fidelities of ostensibly identical experiments performed at different times had dramatically different dependence on photon number [60]. The numerical simulations provide a possible explanation for this behaviour.

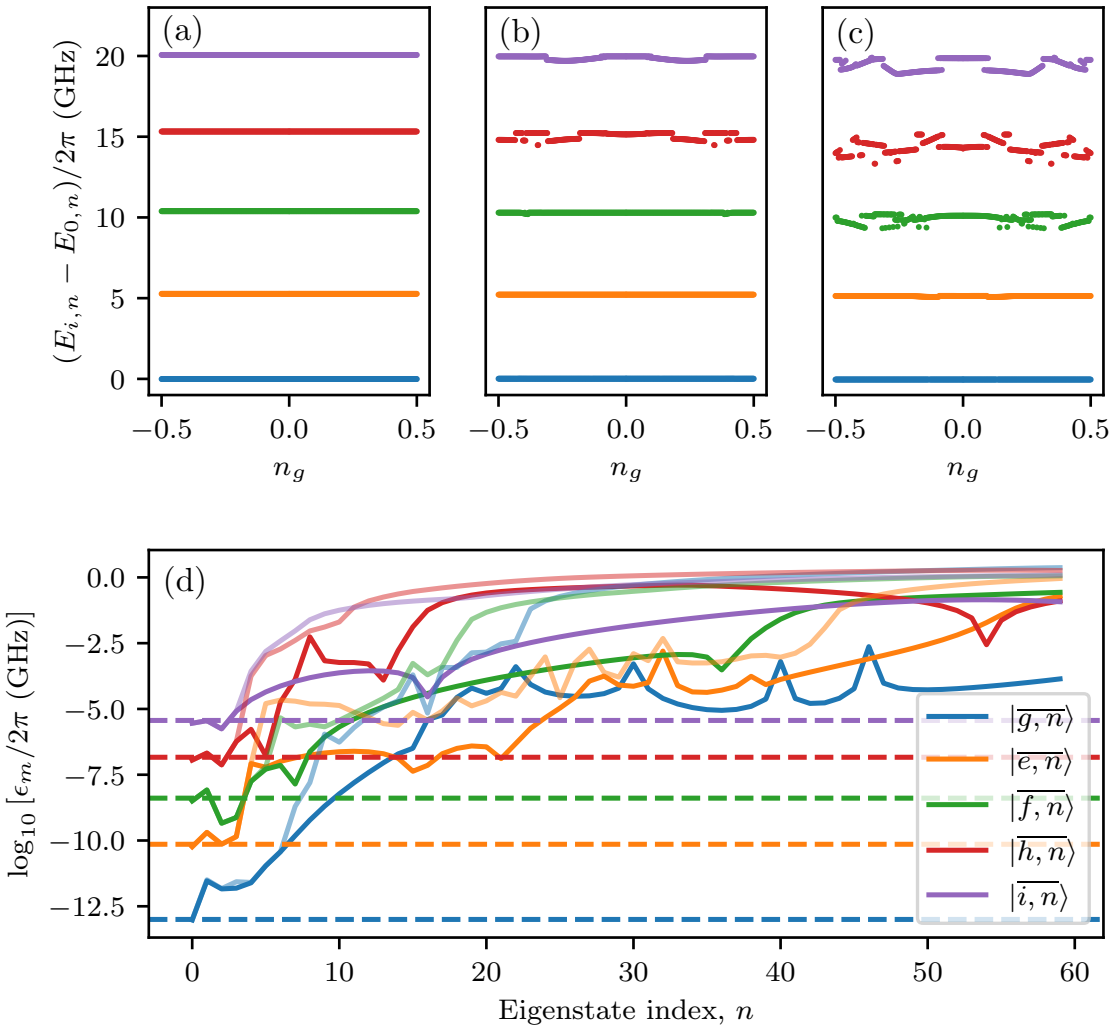

Figure 5.4 (a) The energies of the first five dressed transmon eigenstates as a function of the offset gate charge n_g with the resonator in vacuum. (b-c) The dispersion of the dressed eigenstates as in (a), but with 22 photons (b) and 44 photons (c) respectively. (d) The charge dispersion of the first five branches as a function of the eigenstate photon-number index n. The dashed lines indicate the charge dispersion of the states in the absence of the resonator, i.e. for the isolated transmon. The solid lines are calculated from Eq. (5.11), and the transparent lines using Eq. (5.10).

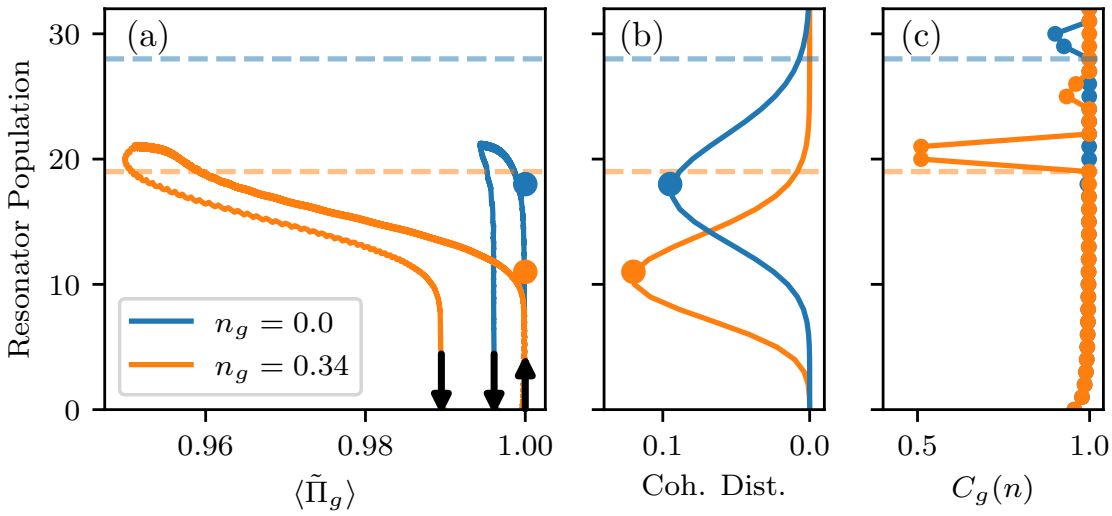

Figure 5.5 (a) A parametric plot of the dressed ground state branch population against the resonator population for two offset gate charges for an otherwise identical readout process. The different final populations of $\approx 99\%$ and $\approx 99.7\%$ indicate that the gate charge has had a significant impact on the dynamics. Arrows indicate the flow of time. (b) Poisson distributions corresponding to coherent states with amplitudes $|\alpha|^2 = 11$ and 19. The dashed lines indicate the 1% population threshold. (c) Overlap metrics for the ground state branch at the two different offset gate charges. We see a strong resonance for $n_g = 0.34$ with a resonator population of ≈ 20.

Chapter 6

Optimizing Superconducting Qubit Readout with Purcell Filters

"FYI, my noise-cancelling headphones proved ineffective last night."

– Sheldon Cooper, The Big Bang Theory

In this chapter, I overview the basics of the Purcell filter, how they can be modelled efficiently in a simulation, detail how the parameters of such a filter should be optimized for efficient qubit readout, and finish with a joint publication in collaboration with ETH Zürich.

6.1 Introduction

Chapters 2 and 5 detail a transmon-resonator system, where the state of the transmon can be determined by driving the resonator. Due to the hybridization of the qubit and resonator, we saw that this leads to a decay channel of the qubit state through the resonator, known as *Purcell decay* [56]:

$$\Gamma_P = \kappa \left|\langle \overline{g,0}|a|\overline{e,0}\rangle\right|^2 \approx \kappa \left(\frac{\tilde{g}}{\Delta}\right)^2. \tag{6.1}$$

Purcell decay can dramatically reduce the lifetime of a qubit – specifically when considering typical parameters such as $\tilde{g}/2\pi = 200$ MHz, $\Delta_{qr}/2\pi = 1$ GHz, and $\kappa/2\pi = 40$ MHz, the Purcell rate results in a mere 100 nanoseconds of qubit lifetime. In comparison, modern transmon qubits exhibit an intrinsic T_1 on the order of tens to hundreds of microseconds

[20, 28, 98], representing a reduction in its potential lifetime by more than two orders of magnitude.

This problem can be addressed with the use of a *Purcell filter* [35, 36, 37]. These filters are used extensively in modern superconducting architectures because of the vast improvement to qubit lifetimes, whilst simultaneously allowing for a large coupling κ to the environment, permitting fast qubit readout [19, 99, 100]. However, they introduce a degree of complication that must be treated carefully in analysis. This is achieved by engineering the electomagnetic environment at the qubit frequency, which is usually far detuned from the readout resonator. This can be achieved with different approaches, including the use of a 'notch' or band-rejection filter which almost entirely eliminates the spectrum of noise at the qubit's operating frequency [35]. Here, we consider the case of a 'band-pass' filter, where the Purcell filter can simply be considered as another resonator mode [101]. This reduces the spectrum of noise at all frequencies other than the resonator frequency, meaning that flux-tunable qubits will continue to be Purcell protected over their range of operation.

6.2 Circuit Hamiltonian and Rate Analysis

To understand how the Purcell filter protects the qubit from decaying into the readout circuitry and how we can readout the qubit state, we begin with the Hamiltonian describing the combined transmon-resonator-Purcell-filter system

$$\begin{aligned}\hat{H}_{trf} &= 4E_c(\hat{n}_t - n_g)^2 - E_J\cos\hat{\varphi}_t + \omega_r\hat{a}^\dagger\hat{a} \\ &\quad - ig(\hat{n}_t - n_g)(\hat{a} - \hat{a}^\dagger) + \omega_p\hat{f}^\dagger\hat{f} + J(\hat{f}^\dagger - \hat{f})(\hat{a}^\dagger - \hat{a}),\end{aligned} \tag{6.2}$$

where ω_p is the frequency of the Purcell filter mode and J is the strength of the capacitive coupling between the filter and resonator modes. An example of this circuit is given in Fig. 6.1. In Chapters 2 and 5, readout was performed by driving the resonator mode. Here, since the resonator is coupled to the feedline *via* the Purcell filter, we cannot drive it directly, only through a coherent drive on the filter mode:

$$\hat{H}_D(t) = -i\Omega(t)(\hat{f}^\dagger - \hat{f})\sin(\omega_d t), \tag{6.3}$$

where $\Omega(t)$ is an envelope function which varies slowly compared to the drive frequency ω_d. Perhaps most importantly, only the *Purcell* mode is coupled to the environment at rate κ. Incorporating this loss rate, we thus formulate a standard Lindblad master equation

$$\dot{\hat{\rho}} = -i[\hat{H}_{trf} + \hat{H}_D(t), \rho] + \kappa\mathcal{D}[\hat{f}]\hat{\rho}. \tag{6.4}$$

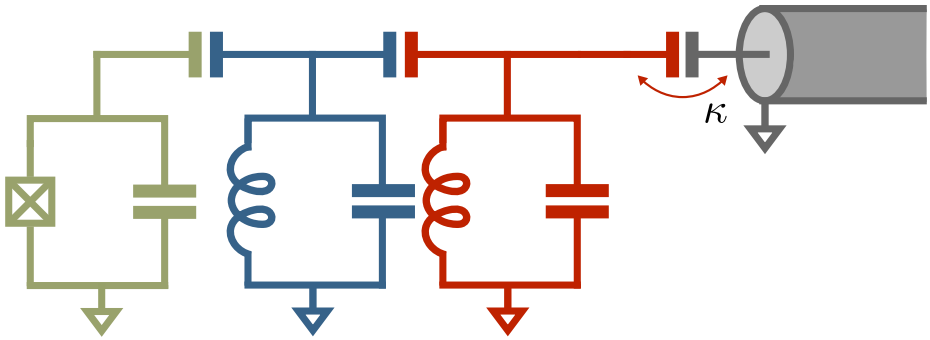

Figure 6.1 Example circuit diagram of the coupled transmon (green), resonator (blue) and Purcell filter (red) circuit. The Purcell filter is coupled to the feedline (grey) at rate κ.

To observe the properties of this master equation, we begin by analysing the static system in the absence of drive, i.e. $\Omega(t) = 0$.

6.2.1 Diagonalizing the resonator-Purcell subspace

To obtain some insight, we can first consider diagonalizing the two coupled linear modes – the resonator and Purcell filter – and find its normal modes before including the transmon. Later on in Sect. 6.2.2, we shall consider a more accurate approach. For simplicity we consider an ideal case where $\omega_p = \omega_r$, which yields two equally hybridized normal modes, $\hat{u}_\pm$ with frequencies $\omega_r \pm J$ respectively:

$$\hat{a} = \frac{1}{\sqrt{2}}(\hat{u}_+ + \hat{u}_-), \quad \hat{f} = \frac{1}{\sqrt{2}}(\hat{u}_+ - \hat{u}_-). \tag{6.5}$$

In this normal mode basis and treating the transmon as a Kerr oscillator, we find a Hamiltonian

$$\begin{aligned}\hat{H}_{tr} = \omega_q \hat{b}^\dagger \hat{b} - \frac{\alpha}{2}\hat{b}^\dagger\hat{b}^\dagger\hat{b}\hat{b} + (\omega_r + J)\hat{u}_+^\dagger\hat{u}_+ + (\omega_r - J)\hat{u}_-^\dagger\hat{u}_- \\ - \frac{i\tilde{g}}{\sqrt{2}}\left[\hat{b}^\dagger(\hat{u}_+ - \hat{u}_-) - \text{H.c.}\right],\end{aligned} \tag{6.6}$$

where $\tilde{g}$ is the renormalized coupling rate, $2\tilde{g} \approx g(E_J/2E_c)^{1/4}$. By employing the transformation $\hat{S} = \exp[-\frac{i\tilde{g}\hat{b}}{\sqrt{2}}(\frac{\hat{u}_+^\dagger}{\Delta_{qr}-J} - \frac{\hat{u}_-^\dagger}{\Delta_{qr}+J}) - \text{H.c.}]$, we can diagonalize Eq. (6.6) to first order in $\tilde{g}$. This transform reveals the approximate dressed eigenstates – in particular, the excited state of the qubit becomes

$$|\overline{e,0,0}\rangle \approx \hat{S}|e,0,0\rangle \approx |e,0,0\rangle + \frac{\tilde{g}}{(\Delta_{qr}-J)\sqrt{2}}|g,1,0\rangle + \frac{\tilde{g}}{(\Delta_{qr}+J)\sqrt{2}}|g,0,1\rangle. \tag{6.7}$$

We can use this form to approximately evaluate the Purcell decay rate

$$\begin{aligned}\Gamma_P &\approx \frac{\kappa}{2}|\langle\overline{g,0,0}|\hat{u}_+ - \hat{u}_-|\overline{e,0,0}\rangle|^2 \\ &= \frac{\kappa}{4}\left|\frac{\tilde{g}}{\Delta_{qr}-J} - \frac{\tilde{g}}{\Delta_{qr}+J}\right|^2 \approx \frac{\kappa\tilde{g}^2}{4}\left|\frac{2J}{\Delta_{qr}^2+J^2}\right|^2 \\ &\approx \kappa\frac{\tilde{g}^2J^2}{\Delta_{qr}^4},\end{aligned} \tag{6.8}$$

valid for small J. This is a dramatic improvement over the original Purcell rate – using the same parameters given in Sect. 6.2 with a typical value $J/2\pi = 20$MHz, the new Purcell-limited lifetime would be approximately 248 microseconds, putting it on par with the state-of-the-art intrinsic lifetimes of transmon qubits [98]. Corrections to this form can be made by looking at the effects of the non-Hermitian Hamiltonian in the single-excitation subspace [101], although for our purposes the differences are negligible.

Correlated Decay

We stress here that correlated decay is vital to explain the reduction in the Purcell rate. Let us consider the dissipator $\mathcal{D}[\hat{f}]\hat{\rho} = \frac{1}{2}\mathcal{D}[\hat{u}_+ - \hat{u}_-]\hat{\rho}$. Going into a rotating frame $\hat{U} = \exp(-i(\omega_r+J)t\hat{u}_+^\dagger\hat{u}_+ - i(\omega_r-J)t\hat{u}_-^\dagger\hat{u}_-)$, each mode picks up an oscillatory term, which leads to the dissipator

$$\mathcal{D}[\hat{u}_+e^{-iJt} - \hat{u}_-e^{iJt}]\hat{\rho}. \tag{6.9}$$

As we did in Eq. (2.28), it can be tempting to make the RWA, which is standard practice for dissipators [102, 103]. This would split the dissipator in two:

$$\mathcal{D}[\hat{u}_+e^{-iJt} - \hat{u}_-e^{iJt}]\hat{\rho} \approx \frac{1}{2}\mathcal{D}[\hat{u}_+]\hat{\rho} + \frac{1}{2}\mathcal{D}[\hat{u}_-]\hat{\rho}. \tag{6.10}$$

This procedure would however yield an *incorrect* Purcell rate:

$$\frac{\kappa}{2}|\langle\overline{g,0,0}|\hat{u}_+|\overline{e,0,0}\rangle|^2 + \frac{\kappa}{2}|\langle\overline{g,0,0}|\hat{u}_-|\overline{e,0,0}\rangle|^2 \approx \kappa\frac{g^2}{\Delta_{qr}^2}. \tag{6.11}$$

In other words, the Purcell suppression is lost. As such, for simulations performed in this dressed basis, it is vital to not make an RWA on this dissipator, since the difference in frequency ($2J$) between these two modes is not sufficiently large. Only in the *strong* coupling case $|J| \gg \kappa$ does the Purcell suppression rate indeed decrease and the RWA once again becomes valid. In comparison, the RWA made in Eq. (2.28) is generally valid, since the detuning between the resonator and qubit is on the order of GHz.

6.2.2 Diagonalizing the transmon-resonator subspace

The majority of devices have a transmon-resonator coupling g on the order of $100-200$ MHz, with the resonator-filter coupling J on the order of 20 MHz. Consequently, in contrast to what I did above, it is reasonable to diagonalize the transmon-resonator subsystem first, and later include the effects of the weakly coupled filter mode. Using the form in Eq. (2.22) and the dressed resonator mode defined in Eq. (2.27) we find the relevant Hamiltonian to be

$$\begin{aligned}\hat{H} &= \bar{\omega}_q \hat{b}^\dagger \hat{b} + \omega_p \hat{f}^\dagger \hat{f} + \omega_r^g \hat{a}^\dagger \hat{a} \\ &+ 2\chi \hat{a}^\dagger \hat{a} \hat{b}^\dagger \hat{b} - \frac{\alpha}{2} \hat{b}^{\dagger 2} \hat{b}^2 - \frac{E_c}{2} \lambda^4 \hat{a}^{\dagger 2} \hat{a}^2 - 2\lambda' \lambda^3 E_C \hat{a}^{\dagger 2} \hat{a}^2 \hat{b}^\dagger \hat{b} \\ &+ J \left(\left[1 - \lambda^2/2 - 2\lambda\lambda' \hat{b}^\dagger \hat{b} \right] \hat{a}^\dagger \hat{f} + \lambda \hat{b}^\dagger \hat{f} + \text{H.c.} \right).\end{aligned} \tag{6.12}$$

where $\lambda' = \lambda E_c / \left[\Delta_{qr} + E_c(1-\lambda^2) \right]$ and $\lambda = g/\Delta_{qr}$. The remaining parameters are defined in Eq. (2.23).

The form of Eq. (6.12) is almost diagonal in the qubit mode, with a small $J\lambda \hat{b}^\dagger \hat{f} + \text{H.c.}$ coupling between the transmon and filter modes – the same term that will lead to the Purcell decay derived in Eq. (6.8). Whilst this term is critical in understanding how the Purcell rate is suppressed, it only results in a very small renormalization of the system frequencies – for the example parameters given in Sect. 6.1 and $J/2\pi = 20$ MHz, we anticipate a shift in the qubit frequency of about 64 kHz. As such, we need not consider this term when analysing the readout dynamics, which is dominated by the other terms in $\hat{H}$, each having contributions at least on the order of MHz. Assuming the qubit is in one of the two computation states ($\langle \hat{b}^\dagger \hat{b} \rangle = 0$ or 1), we can write reduced Hamiltonians which help us understand how the resonator and filter modes will respond during qubit readout:

$$\begin{aligned}\hat{H} &\approx \hat{H}_L + \hat{H}_{NL}, \\ \hat{H}_L &= \omega_p \hat{f}^\dagger \hat{f} + \omega_r^g \hat{a}^\dagger \hat{a} + 2\chi \hat{a}^\dagger \hat{a} \langle \hat{b}^\dagger \hat{b} \rangle + J^{g,e} (\hat{a}^\dagger \hat{f} + \text{H.c.}), \\ \hat{H}_{NL} &= -\frac{E_c \lambda^3}{2} (\lambda + 4\lambda' \langle \hat{b}^\dagger \hat{b} \rangle) \hat{a}^{\dagger 2} \hat{a}^2,\end{aligned} \tag{6.13}$$

where we have defined a resonator-Purcell coupling $J^{g,e} = J \left[1 - \lambda^2/2 - 2\lambda\lambda' \langle \hat{b}^\dagger \hat{b} \rangle \right]$ which depends on the qubit state. We can further define qubit-state dependent Kerr nonlinearities, $K^g = -E_c \lambda^4/2$ and $K^e = -E_c \lambda^3 (\lambda + 4\lambda')/2$. Further, we note that for $\Delta_{qr} < 0$, it is simple to demonstrate $|\lambda + 4\lambda'| < |\lambda|$ and we consequently find $|K^e| < |K^g|$.

For small drive powers, we expect a small population of the readout resonator. As such, H_{NL} becomes negligible and the linear Hamiltonian $\hat{H}_L$ in Eq. (6.13) is sufficient to analyse

the readout. As such we can, to an excellent degree of approximation, assume that the resonator and filter modes both have coherent state responses. Consequently, we assume that the computational states $|g,0,0\rangle \rightarrow |g,\alpha^g(t),\beta^g(t)\rangle$ and $|e,0,0\rangle \rightarrow |e,\alpha^e(t),\beta^e(t)\rangle$ under the action of the Hamiltonian and loss channel, where $\alpha^{g,e}(t)$, $\beta^{g,e}(t)$ refer to the time-dependent coherent amplitude of the resonator and filter modes, respectively This approximation of course only remains valid at low photon numbers, which is the emphasis of the preprint in Sect. 6.4. I shall discuss the effects of nonlinearities in Sect. 6.3.5.

6.2.3 Normal Mode Hamiltonian

The assumption of a coherent response in both the resonator and filter modes allows us to dramatically simplify the analysis, turning the master equation into a description of two coupled linear oscillators. Following the results in Sect. 6.4, the master equation in Eq. (6.4) effectively reduces to the simple matrix equation

$$\begin{bmatrix} \dot{\alpha}^{g,e} \\ \dot{\beta}^{g,e} \end{bmatrix} = -i \begin{bmatrix} \Delta^{g,e} & J^{g,e} \\ J^{g,e} & \Delta_p - i\kappa/2 \end{bmatrix} \begin{bmatrix} \alpha^{g,e} \\ \beta^{g,e} \end{bmatrix} + \begin{bmatrix} 0 \\ \mathcal{E} \end{bmatrix}, \tag{6.14}$$

where $\Delta^{g,e} = \omega_r^{g,e} - \omega_d$ is the detuning between the *dressed* resonator conditioned on the qubit state and the drive, and $\Delta_p = \omega_p - \omega_d$ is the filter-drive detuning. Further, $\omega_r^e = \omega_r^g + 2\chi$ includes the χ shift imparted on the resonator by the qubit.

In this section, we are principally interested in the optimal choices of the Purcell filter parameters and coupling (J, ω_p, κ), rather than optimizing the χ-shift and other parameters of the transmon-resonator – a more detailed discussion of this aspect of the parameter optimization can be found in the preprint in Sect. 6.4.

6.3 Signal-to-Noise Ratio

As discussed in Chapter 2, the signal-to-noise ratio (or SNR for short) tells us information about how well the pointer states are resolved in the phase space of the resonator and will be one of the major predictors of the readout fidelity. However, unlike Chapter 2, we are measuring the output field of the *Purcell filter*, since only this resonator is directly connected to the feedline. Consequently, we must adapt the definition in Eq. (2.34) to the Purcell filter, which reduces to a simple description of the distance between the formed coherent states in the *internal filter mode*:

$$\mathrm{SNR}(t) = 2\kappa\eta \int_0^t |\beta^g(t') - \beta^e(t')|^2 dt'. \tag{6.15}$$

We seek to maximize the SNR for obtaining the largest upper bound on the readout fidelity – however, this is a non-trivial multivariable optimization, and care must be taken to understand the SNR landscape as a function of all of the relevant variables. For the rest of this chapter, I shall assume that the transmon-resonator subsystem is *fixed*, and only consider the parameters $\omega_d, \kappa, J, \omega_p$, the readout integration time t and the drive amplitude $\mathcal{E}$ in the optimization. Further details can be found in Appendix A.

The analytical expression for the time-dependent filter-mode responses $\beta^{g/e}(t)$ is provided in the preprint in Sect. 6.4. Of particular interest are the steady-state responses in both the filter and resonator modes, given by the expression

$$\begin{bmatrix} \alpha_{ss}^{g,e} \\ \beta_{ss}^{g,e} \end{bmatrix} = \frac{\mathcal{E}}{\Delta^{g,e}(\Delta_p - i\kappa_p/2) - J^2} \begin{bmatrix} -J \\ \Delta^{g,e} \end{bmatrix}. \tag{6.16}$$

From the above form, it is clear that the displacement between the two filter responses will increase linearly as a function of the drive amplitude $\mathcal{E}$ – however, from Chapter 5 and an abundance of experimental evidence [30, 61, 62, 95], large drive amplitudes result in a decrease in the readout fidelity due to breaking the weak-drive approximation, with a more significant contribution from the non-linear Hamiltonian. At even higher drive amplitudes, this model will also break down and ionization can occur, as in Chapter 5.

6.3.1 How strong is too strong?

In Chapter 5, I detailed how the states of the combined transmon-resonator system can be systematically identified, and that resonances occur at specific photon numbers which can

lead to the ionization of the transmon. Here, I wish to emphasize that, even in the transmon-resonator-Purcell system, the dominant indicator of ionization and strong nonlinearities will be the population of the *resonator*, and the Purcell filter population is irrelevant.[1] This is fairly intuitive –in Eq. (6.2), the transmon is only coupled to the resonator, not the filter. After the transmon-resonator subspace is diagonalized as in Eq. (6.12), the effective Purcell filter-qubit coupling is negligible. Consequently, I expect the Purcell mode to be predominantly linear, and its population alone should not govern nonlinear effects.

To demonstrate this further, I consider a representative set of parameters (see caption in Fig. 6.2) and diagonalize the entire transmon-resonator-Purcell system with methods detailed in Chapters 2 and 5. To perform this, the transmon-resonator subspace is first diagonalized with $t_{\text{dim}} = 18, r_{\text{dim}} = 125$, which is then projected onto the first 10 branches, preserving 35 resonator states per branch. I then add the Purcell mode with 35 levels, leading to the diagonalization of a Hilbert space size of 12250. Then, for each eigenstate, I calculate the transmon and resonator populations N_t and N_r as in Chapter 5, but in addition the average *filter* population of each state, $N_f = \langle \hat{f}^\dagger \hat{f} \rangle$. Then, in Fig. 6.2(a), I plot the transmon population against the resonator population, and in (b) the transmon population against the *filter* population. As can be seen in (a), there are strong correlations between the resonator and transmon populations – these are indeed closely related to the branches identified from the transmon-resonator population alone (see faint thick lines). In comparison, there is negligible relationship between the N_f and N_t populations seen in panel (b). This strengthens the conclusion that indeed, the number of photons in the *resonator*, not the filter, is the limiting factor before strong nonlinearities enter the system and ionization can occur.

6.3.2 Defining an optimization procedure

In our pursuit of optimizing the readout, we must take into account experimental considerations and limit the complexity. With this in mind, here I presume that

- the drive frequency ω_d is chosen independently of the drive amplitude and readout integration time;
- the number of photons that can be used in the readout process before spurious effects

[1] Due to the hybridization of the resonator and Purcell filter, their populations are inextricably linked, and the degree to which each is populated depends on which normal mode is driven. Nevertheless, the onset of ionization only depends explicitly on the resonator mode.

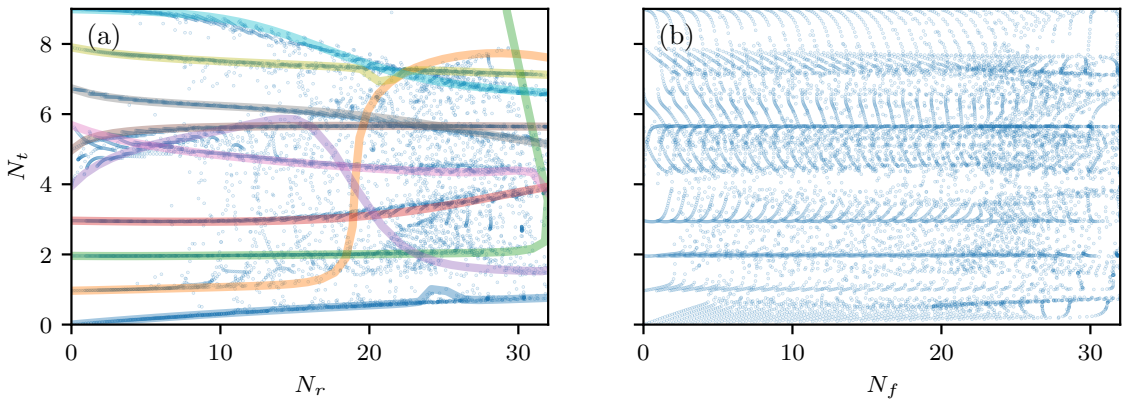

Figure 6.2 (a) Plot of the expected resonator population N_r against the transmon population N_t for all eigenstates of the combined transmon-resonator-Purcell system. The faint lines indicate the identified branches, obtained from diagonalizing the transmon-resonator alone. (b) A plot of the filter population N_f against the transmon population N_t. No strong correlations are observed between these two populations. The parameters used were: $E_c/\hbar = 200$ MHz, $E_J/E_c = 50$, $\omega_r/2\pi = 5.74$ GHz, $\omega_p/2\pi = 5.8$GHz, $J/2\pi = 20$ MHz, $\tilde{g}/2\pi = 200$ MHz.

enter is equal to n_{crit}:[2]

- the linear Hamiltonian remains a good approximation during the readout process;
- the Kerr nonlinearity limits the performance before the presence of spurious transitions.

Given the second assumption, we can begin by fixing the number of photons in the resonator $|\alpha_g^{ss}|^2 = n_{\text{crit}}$, which simply restricts $\mathcal{E}$:

$$n_{\text{crit}} = \left| \frac{-J\mathcal{E}}{\Delta^g(\Delta_p - i\kappa/2) - J^2} \right|^2 . \tag{6.17}$$

Given we can change the drive phase, we choose α_{ss}^g to be entirely real and set

$$\mathcal{E}_{\text{crit}} = -J^{-1}\sqrt{n_{\text{crit}}}[\Delta^g(\Delta_p - i\kappa/2) - J^2], \tag{6.18}$$

so that $\alpha_{ss}^g = \sqrt{n_{\text{crit}}}$. We note that there is now a *different* photon number response in the excited state,

$$|\alpha_{ss}^e|^2 = n_{\text{crit}} \left| 1 + \frac{2\chi(\Delta_p - i\kappa/2)}{(\Delta^e(\Delta_p - i\kappa/2) - J^2)} \right|^2 . \tag{6.19}$$

[2]From the previous chapters and Ref. [2], it is clear that n_{crit} is not a perfect metric. Nevertheless, it is useful in this context to help simplify the optimization procedure.

Depending on the choice of drive frequency and parameters, we commonly find $|\alpha_{ss}^e|^2 > |\alpha_{ss}^g|^2$, which may seem detrimental given the excited state branch generally ionizes before the ground state [2, 63]. However, drawing attention to the first and final assumptions, we cannot reoptimize the drive frequency based on the drive shift/induced Kerr nonlinearities and, more notably, the Kerr nonlinearity of the excited state branch is significantly lower than the ground state branch as predicted by the dispersive theory – see Eq. (6.13). As a result, for this optimization procedure, we assume that the excited state can tolerate a larger resonator response. Specifically, we can define a maximum photon number ratio:

$$\frac{|\alpha_{ss}^e|^2}{|\alpha_{ss}^g|^2} \leq \frac{K_g}{K_e} = r, \tag{6.20}$$

such that the nonlinearities for the excited state are, at most, equivalent to the ground state.

6.3.3 Choosing the drive frequency

As explained in the preprint in Sect. 6.4, the readout drive frequency can be chosen rather simply: the spectral function is measured at a range of relevant drive frequencies with a constant amplitude, and the frequency that exhibits the largest response is chosen as the readout frequency. Here, I'd like to stress that the choice of readout frequency can be further optimized by considering the resonator response. In Fig. 6.3, I plot in blue the 'performance', or the normalized SNR rate, taken to be $|\beta_{ss}^g - \beta_{ss}^e|^2/\max(|\beta_{ss}^g - \beta_{ss}^e|^2)$, as a function of the drive frequency ω_d at constant power. The peak here corresponds to one of the two normal mode responses – see Sect. 6.4 for more details. However, this metric does not consider the corresponding steady-state population in the *resonator*, which ultimately dictates the maximum drive power based upon the above assumptions. By renormalizing this performance metric by the population of the ground state, $|\alpha_{ss}^g|^2$, we find the optimal drive frequency to decrease slightly on the order of $1-2$ MHz. Finally, we must ensure that the excited state resonator response is not too large – consequently, we filter out a range of drive frequencies for which $\frac{|\alpha_{ss}^e|^2}{|\alpha_{ss}^g|^2} > \frac{K_g}{K_e}$, indicated by the dashed line. Finally, the optimal drive frequency is indicated by the arrow.

6.3.4 Readout Performance

Using the methodology above, we are now ready to plot the relative performance of the readout. For this, I consider a realistic transmon-resonator-Purcell system, with parameters: $E_J/2\pi = 22.65$ GHz, $E_J/E_C = 134$, $\omega_r/2\pi = 6.86$ GHz, and $g/2\pi = 210$ MHz, which

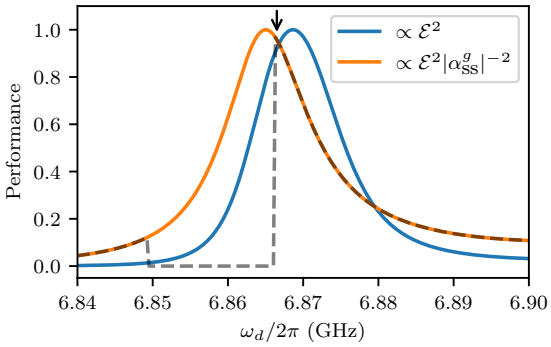

Figure 6.3 Performance of the readout as a function of the drive frequency, as indicated by constant drive (blue) and when renormalized by the resonator response (orange). The grey dashed line and orange lines overlap for drive frequencies giving an acceptable photon ratio response according to Eq. (6.20).

yields a qubit frequency $\omega_q/2\pi = 5.3$ GHz and a chi-shift $\chi/2\pi \approx 7$ MHz. Then, using the above optimization procedure, I perform sweeps of the Purcell-resonator coupling strength J, the decay rate κ, and the filter frequency ω_p, integrating for a readout integration time of up to 400 ns, and calculate the SNR in Eq. (6.15) assuming a measurement efficiency $\eta = 20\%$. In Fig. 6.4, I plot a representative example of these results, with panels (a-c) corresponding to a readout integration time of 50 ns and (d-f) to 400 ns, and the different columns corresponding to detunings of the Purcell filter away from the central frequency response of the resonator, $\omega_c = (\omega_r^e + \omega_r^g)/2$.

This data provides invaluable insights into parameter optimization. For example, at short readout times (a-c), we see that the readout performance is very poor for $\kappa/2\pi < 50$ MHz, which suggests that broadband Purcell filters are required for high-fidelity fast readout. Perhaps more surprisingly, we find that large SNRs can still be obtained at short readout times when the Purcell filter is detuned away from the central frequency, see Fig. 6.4(a), although this requires very precise engineering of the coupling strength J to compensate. For even larger detunings in Fig. 6.4(c), loss rates as high as 200 MHz are required for reasonable SNR.

At longer readout times of 400 ns (d-f), we see much broader features and can obtain excellent readout fidelities with small J and κ, in particular for (e), where $\omega_p = \omega_c$. We additionally note that good readout can be obtained for small κ and large J – however, in this limit, the two normal modes are strongly separated in frequency, leading to an *increased* Purcell decay rate.

Experimentally, the filter frequency is generally hard to precisely fabricate [104], with

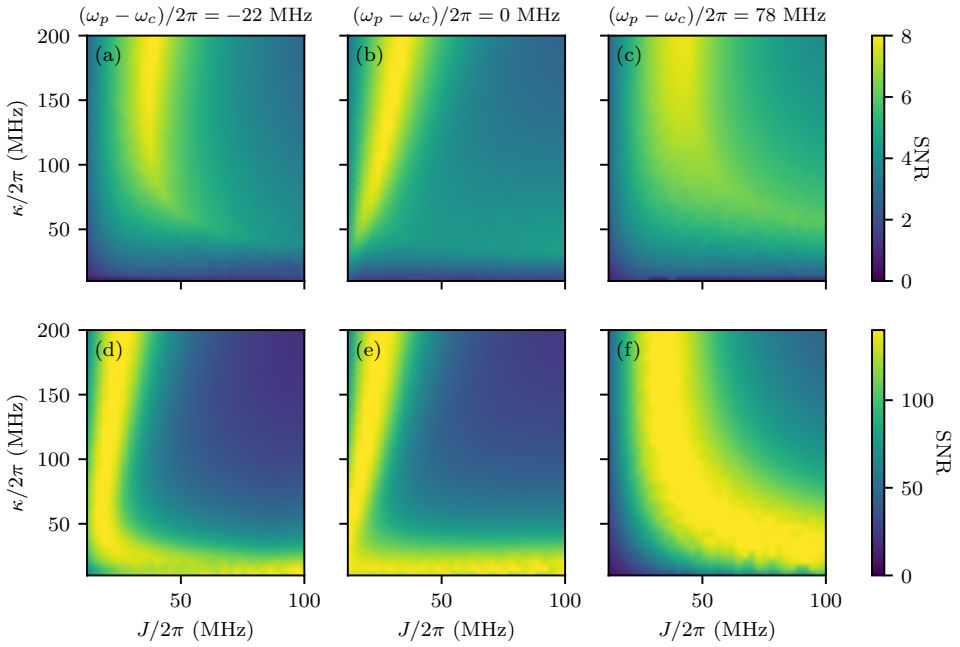

Figure 6.4 Performance of the readout as a function of the drive frequency at 50 ns (a-c) and 400 ns (d-f) for various parameters. Fast readout times require a large decay κ and more precise fabrication of the coupling strength J.

deviations on the order of 10s of MHz expected. This sort of analysis can provide a pathway for optimizing the relative parameters to ensure the fabricated couplings yield effective readout rates for various qubit setups. In Appendix A I provide additional details and analysis of the SNR function in different limits.

6.3.5 Nonlinear Effects – Semiclassical approach

As discussed in Chapter 5, the use of nonlinear semiclassical equations of motion proved useful in accurately describing important aspects of ionization dynamics. A suitable question therefore arises in this context – does accounting for the Kerr-nonlinearities in the semiclassical equations lead to better predictions of the dynamics at higher powers?

In Fig. 6.5, I plot in blue an example trajectory calculated from the *non-linear* equations of motion, i.e. where we preserve the Kerr nonlinearity in the resonator (see [2]) for the qubit initialized in the ground state:

$$\begin{aligned} \dot{\alpha}^{g,e} &= -i(\Delta^{g,e} - \frac{K^{g,e}}{2}|\alpha|^2)\alpha^{g,e} - iJ^{g,e}\beta^{g,e}, \\ \dot{\beta}^{g,e} &= -i(\Delta_p - i\kappa/2)\beta^{g,e} - iJ^{g,e}\alpha^{g,e} + \mathcal{E}, \end{aligned} \tag{6.21}$$

with parameters $J/2\pi = 30$ MHz, $\omega_p/2\pi = 6.88$ GHz, $K^g/2\pi \approx 600$ kHz, $K^e/2\pi \approx 100$ kHz, and $\kappa/2\pi = 35.3$ MHz, with the remaining parameters identified above. I additionally solve the linear set of equations Eq. (6.14), given in orange and compare these solutions to the predicted field from the master equation (see Sect. 6.4). The master equation, calculated using the voltsolve(alt) algorithm detailed in Chapter 4, was solved with 4 transmon levels, and 16 states in both the resonator and Purcell filter, corresponding to a Hilbert space size of 1024. The nonlinearity provides a small but important correction, bringing the dynamics closer to the true solution. This therefore suggests that the parameters for low-power qubit readout with Purcell filters, *including* the drive frequency and power, could be optimized without the need of simulating the entire master equation and Hamiltonian.

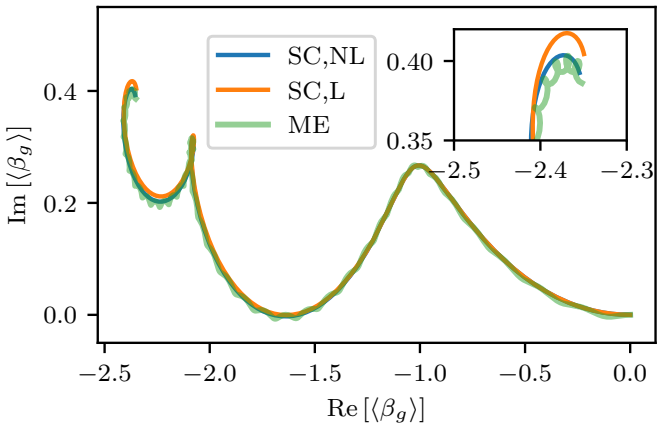

Figure 6.5 Parametric trajectory of the filter response $\beta_g(t)$ when the qubit is in the ground state, as calculated from the nonlinear coupled semiclassical equations of motion in Eq. (6.21) (blue), from the standard linear relation in Eq. (6.14)(orange) and from the solution to the master equation (green) for a duration of 400 ns. The inset more clearly shows the deviation at the system's steady state.

6.4 Publication – Enhancing Dispersive Readout of Superconducting Qubits Through Dynamic Control of the Dispersive Shift: Experiment and Theory [3]

In this publication, in which I worked as the theory support for the experimental results, we present a method by which the readout fidelity can be optimized by changing the qubit *frequency*. In doing so, we can increase the dispersive shift χ and the effective linewidth of the targeted mode, increasing the rate of SNR and optimizing the readout in a short time. As of writing, these results are state-of-the-art, reaching a readout fidelity of 99.75% in only 100

ns. This functions as a complement to the above work, where in this experimental setup, the resonator-qubit coupling J is fixed but the dispersive shift and the Purcell-resonator detuning varies. We additionally clearly see at which drive powers the dispersive measurement breaks down, which is indicated by a large drop in SNR and readout fidelity.

Enhancing Dispersive Readout of Superconducting Qubits Through Dynamic Control of the Dispersive Shift: Experiment and Theory

François Swiadek,[1,2,*] Ross Shillito,[3] Paul Magnard,[1] Ants Remm,[1,2] Christoph Hellings,[1,2] Nathan Lacroix,[1,2] Quentin Ficheux,[1] Dante Colao Zanuz,[1,2] Graham J. Norris,[1,2] Alexandre Blais,[3,4] Sebastian Krinner,[1,2] and Andreas Wallraff[1,2,5]

[1] *Department of Physics, ETH Zurich, 8093 Zurich, Switzerland*
[2] *Quantum Center, ETH Zurich, 8093 Zurich, Switzerland*
[3] *Institut Quantique and Département de Physique, Université de Sherbrooke, Sherbrooke J1K 2R1 QC, Canada*
[4] *Canadian Institute for Advanced Research, Toronto, Ontario M5G 1M1, Canada*
[5] *ETH Zurich - PSI Quantum Computing Hub, Paul Scherrer Institute, 5232 Villigen, Switzerland*

(Dated: July 18, 2023)

The performance of a wide range of quantum computing algorithms and protocols depends critically on the fidelity and speed of the employed qubit readout. Examples include gate sequences benefiting from mid-circuit, real-time, measurement-based feedback, such as qubit initialization, entanglement generation, teleportation, and perhaps most importantly, quantum error correction. A prominent and widely-used readout approach is based on the dispersive interaction of a superconducting qubit strongly coupled to a large-bandwidth readout resonator, frequently combined with a dedicated or shared Purcell filter protecting qubits from decay. By dynamically reducing the qubit-resonator detuning and thus increasing the dispersive shift, we demonstrate a beyond-state-of-the-art two-state-readout error of only 0.25 % in 100 ns integration time. Maintaining low readout-drive strength, we nearly quadruple the signal-to-noise ratio of the readout by doubling the readout mode linewidth, which we quantify by considering the hybridization of the readout-resonator and its dedicated Purcell-filter. We find excellent agreement between our experimental data and our theoretical model. The presented results are expected to further boost the performance of new and existing algorithms and protocols critically depending on high-fidelity, fast, mid-circuit measurements.

Realizing high-fidelity and fast single-shot readout of a qubit [1–3] is essential for quantum error correction protocols [4–8] in which qubit decoherence during readout and reset contributes significantly to the logical error. It is also key for algorithms requiring real-time feedback, such as teleportation [9–12], distillation [13, 14] and initialization [15–18].

In superconducting circuits, the most commonly used readout architecture employs the state-dependent dispersive shift of the resonance frequency of a resonator coupled to the qubit to infer the qubit state [19–21]. Whilst the frequency of the resonator is typically fixed, flux-tunable transmons allow to control the qubit-resonator detuning by modifying the transmon frequency [22], enable high-fidelity fast entangling gates [23–25] and avoid frequency collisions. Additionally, each qubit is often coupled to a microwave transmission line via a dedicated [6, 7, 26] or common Purcell filter [8, 27, 28] to protect the qubit from radiative decay [29–31]. Such measurements are usually performed with *weak* measurement tones to avoid nonlinearities and detrimental qubit state transitions, although high-power readout has been studied both theoretically [32] and experimentally [2].

In the past few years, significant improvements to the single-shot readout have been realized, reaching a two-level readout assignment fidelity of 4×10^{-3} in 88 ns [3]. Faster readout protocols have been realized, with a 9×10^{-3} fidelity readout achieved in 40 ns by utilizing the distributed-element, multimode nature of the readout resonator [33].

One of the critical parameters governing dispersive qubit readout is the detuning between the qubit and the readout resonator, which controls both the magnitude of the dispersive shift and the nonlinearities induced in the resonator. Different detuning regimes have been explored, including cases where the resonator frequency is *lower* than the qubit [3, 34]. Notably, the measurement fidelity has been shown to improve for smaller detunings [7, 35], although these observations were not fully explained.

In this work, we demonstrate an increase in the signal-to-noise ratio (SNR) and assignment fidelity by bringing the qubit frequency closer to the readout resonator's frequency using a flux pulse, see illustration in Fig. 1 (a,b), achieving a minimum two-level readout error of 2.5×10^{-3} in 100 ns. We accredit this remarkable performance not only to an increase in the dispersive shift χ imparted by the qubit on the cavity, but also to an increase in the effective linewidth of the targeted normal mode response, caused by bringing the Lamb-shifted readout resonator closer to resonance with the Purcell filter, see Fig. 1 (c).

* francois.swiadek@phys.ethz.ch

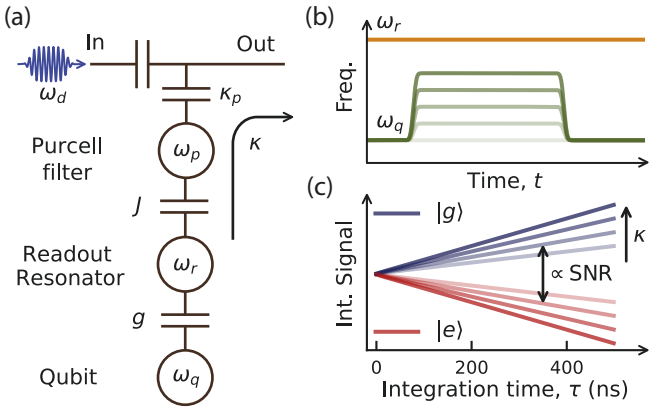

FIG. 1. **(a)** Schematic of a qubit coupled to a readout resonator-Purcell-filter system. The qubit of transition frequency ω_q is coupled capacitively at rate g to a readout resonator of frequency ω_r. The readout resonator in turn is coupled at rate J to a Purcell filter of frequency ω_p, which is coupled to a feedline at rate κ_p. We probe the system by measuring the transmission of a readout pulse at frequency ω_d through the feedline. The effective decay rate of the readout resonator is indicated as κ. **(b)** Schematic of time-dependence of qubit frequency ω_q relative to the readout resonator frequency ω_r. The qubit, initially idling at the lower flux sweet spot, is pulsed to a smaller detuning from the readout resonator using a fast Gaussian-filtered, rectangular flux pulse. **(c)** Illustration of the rise of the SNR of the qubit readout with integration time τ parameterized by the effective linewidth of the readout resonator κ at approximately constant dispersive shift χ. Increasing color saturation in panels (b) and (c) indicate increasing κ at reduced detuning between ω_q and ω_r.

I. READOUT PARAMETER CHARACTERIZATION

We perform the experiment with a transmon qubit of transition frequency $\omega_q/2\pi = 4.14\,\text{GHz}$ at the lower flux sweet spot [36] and anharmonicity $\alpha/2\pi = -181\,\text{MHz}$. It has a lifetime $T_1 = 30.4\,\mu\text{s}$ and is capacitively coupled to a readout resonator with a coupling strength $g/2\pi = 224\,\text{MHz}$. The readout resonator is coupled to a feedline used for multiplexed readout [26] via a dedicated Purcell filter of linewidth κ_p and with a coupling strength J, see Fig. 1 (a). The qubit is located on a device used to execute a distance-three surface code (see Fig. 7, Appendix A). Further information on the device properties and its fabrication can be found in Ref. [7].

To determine the readout parameters as a function of the frequency detuning between the qubit and its readout resonator, we perform pulsed spectroscopy experiments. We first prepare the qubit in the ground state $|g\rangle$ or excited state $|e\rangle$, pulse the qubit to a chosen readout frequency ω_q using a baseband flux pulse, and probe the readout circuit using a $2.2\,\mu\text{s}$ long microwave tone. This duration corresponds to the maximum integration time of our readout electronics (see Appendix A). The flux pulses are Gaussian-filtered rectangular pulses with short rising and falling edges minimizing coupling to two-level systems [7], see Fig. 1 (b).

We repeat the experiment for five different qubit—readout-resonator detunings $\Delta_{qr}/2\pi$, spanning $-2.7\,\text{GHz}$ to $-1.3\,\text{GHz}$, where $\Delta_{qr} = \omega_q - \omega_r^g$. We denote $\omega_r^{g/e}$ as the readout resonator frequency with the qubit prepared in the ground/excited state. The measured (light colored lines) and calculated (dark colored lines) transmission response is shown in Fig. 2 (a), with blue/red lines corresponding to the qubit prepared in the ground/excited state. From a fit to a coupled qubit—readout-resonator—Purcell-filter model (see Appendix A and solid blue and red lines in Fig. 2 (a)), we extract the relevant readout parameters for each value of Δ_{qr}. The measured (dots) and calculated (lines) dressed readout resonator frequencies ω_r^g and $\omega_r^e = \omega_r^g + 2\chi$ are shown in Fig. 2 (b) (blue and red circles) as a function of the detuning Δ_{qr}, along with the Purcell filter frequency $\omega_p/2\pi = 6.900\,\text{GHz}$, which remains constant. The variation in the resonator frequencies $\omega_r^{g/e}$ is due to the Lamb shift g^2/Δ_{qr} caused by the qubit [22]. Furthermore, we extract both a large intended Purcell filter linewidth $\kappa_p/2\pi = 34.5\,\text{MHz}$ and a large intended coupling strength between the readout resonator and the Purcell filter $J/2\pi = 27.5\,\text{MHz}$.

We consider a standard circuit-QED approach to model the transmon—resonator—Purcell-filter system, see Appendix B for details. In the case of a weak drive $\mathcal{E}$ applied to the filter mode, the readout resonator and the Purcell filter responses can be considered as linear. As such, the dynamics can be effectively mapped to the equations of motion [37]

$$\begin{bmatrix}\dot{\alpha}^{g/e}\\ \dot{\beta}^{g/e}\end{bmatrix} = -i\begin{bmatrix}\omega_r^{g/e} & J\\ J & \omega_p - i\kappa_p/2\end{bmatrix}\begin{bmatrix}\alpha^{g/e}\\ \beta^{g/e}\end{bmatrix} + \begin{bmatrix}0\\ \mathcal{E}e^{-i\omega_d t}\end{bmatrix}, \tag{1}$$

where α and β represent the coherent fields of the readout resonator and Purcell filter, respectively. In the regime $J \approx \kappa_p$, we observe two distinct hybridized readout-resonator—Purcell-filter modes, see Fig. 2 (a). We denote these as the *low* and *high* readout modes, respectively, the lowest and highest of the two modes in the transmission spectrum. The frequency and linewidth of these modes can be determined, respectively, from the real and the imaginary part of the eigenvalues of the equations of motion in Eq. (1) in the absence of a drive:

$$\begin{aligned}\omega_{l,h}^{g/e} &= \frac{\omega_r^{g/e}+\omega_p}{2} \pm \frac{1}{2}\,\text{Re}\sqrt{\left(\Delta_{rp}^{g/e} + \frac{i\kappa_p}{2}\right)^2 + 4J^2},\\ \kappa_{l,h}^{g/e} &= \frac{\kappa_p}{2} \mp \text{Im}\sqrt{\left(\Delta_{rp}^{g/e} + \frac{i\kappa_p}{2}\right)^2 + 4J^2}.\end{aligned} \tag{2}$$

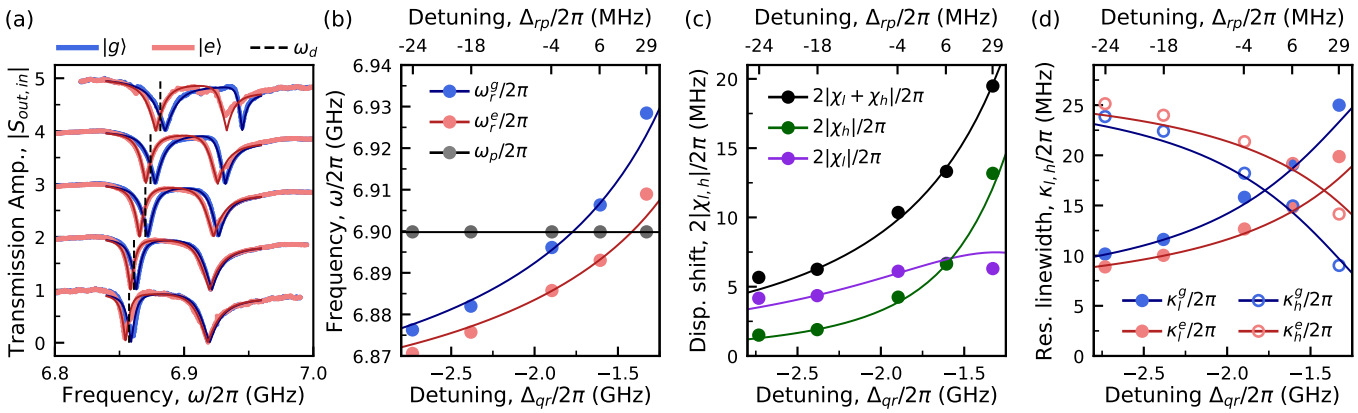

FIG. 2. **(a)** Readout circuitry transmission spectra measured for five qubit—readout-resonator detunings ($\Delta_{qr}/2\pi \in [-2.7, -2.4, -1.9, -1.6, -1.3]$ GHz from bottom to top, vertically shifted by increments of one). Spectra are shown for both the qubit prepared in the ground state $|g\rangle$ (blue) and in the excited state $|e\rangle$ (red). Solid lines are fits based on a coupled qubit—readout-resonator—Purcell-filter model (see Appendix A). Dashed black lines indicate the selected readout frequency, chosen such that the response in transmission $|S^{|e\rangle}_{\text{out,in}} - S^{|g\rangle}_{\text{out,in}}|$ is maximum. **(b)** Resonator frequency ω_r^g (ω_r^e) conditioned on the qubit being prepared in the ground (excited) state, and Purcell filter frequency ω_p as a function of Δ_{qr}. **(c)** Measured dispersive shift χ_l (χ_h) of the lower (higher) frequency hybridized readout mode (see green (purple) points), as a function of Δ_{qr}. The two contributions sum up to the bare readout resonator mode dispersive shift (solid black line) $\chi = \chi_l + \chi_h \approx \alpha g^2/\Delta_{qr}^2$. The solid lines (purple and green) are fits based on a qubit—readout-resonator—Purcell-filter model (see Appendix A). **(d)** Measured effective readout resonator linewidth κ_l (κ_h) for the lower (higher) frequency hybridized readout mode for the qubit prepared in the ground/excited (see blue/red), as a function of Δ_{qr}. Solid lines are fits based on a qubit—readout-resonator—Purcell-filter model (see Appendix A). In (b,c,d) the corresponding detunings $\Delta_{rp} = \omega_r^g - \omega_p$ between the Purcell filter and the readout resonator, indicating the degree of hybridization of the two resonator modes, are shown on the top axis.

The readout-resonator—Purcell-filter hybridization leads to a *distribution* of the total qubit-induced dispersive shift χ on the readout-resonator—Purcell-filter system. Using the model in Eq. (2), we can extract the dispersive shifts of the *low* and *high* modes respectively, $\chi_{l/h} = (\omega^e_{l/h} - \omega^g_{l/h})/2$, see purple (green) circles for the *low* (*high*) mode in Fig. 2 (c). While the total dispersive shift $\chi = \chi_l + \chi_h$ shows the expected $\alpha g^2/\Delta^2_{qr}$ dependence in Fig. 2 (c) (solid black line), with $2\chi/2\pi \in [-5.67, -19.49]$ MHz, the *low* mode dispersive shift only shows small variations in that range, staying between $2\chi_l/2\pi \in [-4.17, -6.69]$ MHz (see solid purple line in Fig. 2 (c)). In contrast, the *high* mode dispersive shift shows a similar scaling with Δ_{qr} as the total dispersive shift, with $2\chi_h/2\pi \in [-1.5, -13.18]$ MHz (solid green line in Fig. 2 (c)).

We observe that the dispersive shift of the *low* mode is dominant for qubit—readout-resonators detunings below -1.6 GHz, after which the dispersive shift of the *high* mode becomes larger. The crossing point where $\chi_l = \chi_h$, in the vicinity of the qubit—readout-resonator detuning $\Delta_{qr}/2\pi = -1.6$ GHz, corresponds to an equal hybridization of the two readout modes. It coincides with $\omega_p = \omega_r^g + \chi = \omega_r^e - \chi$ being equidistant to the ground and excited state responses of the readout resonator, see Fig. 2 (b).

Our model also gives us valuable information about the linewidth of the low and high modes, for the qubit prepared in the ground $|g\rangle$ or excited state $|e\rangle$, namely κ_l^g (κ_h^g) and κ_l^e (κ_h^e), as a function of the qubit—readout-resonator detuning Δ_{qr}. As shown in Fig. 2 (d), while $|\chi_l| > |\chi_h|$ for $\Delta_{qr}/2\pi \leq -1.6$ GHz, we have $\kappa_l^{g/e} < \kappa_h^{g/e}$. This is expected as for $\Delta_{qr}/2\pi \leq -1.6$ GHz, the *low* mode has a larger weight in the readout resonator. The difference between κ_l^g and κ_l^e for the *low* mode derives from the frequency detuning $\Delta_{rp}^{g/e} = \omega_r^{g/e} - \omega_p$, between the readout-resonator frequency and the Purcell-filter frequency. In particular, for the *low* mode, $\kappa_l^g > \kappa_l^e$ for all detunings while for the *high* mode $\kappa_h^g < \kappa_h^e$, which can be seen from the analysis of the normal mode Hamiltonian in Appendix B.

In the vicinity of the detuning leading to an equal hybridization of the *low* and *high* modes $\Delta_{qr}/2\pi \approx -1.6$ GHz, we further note that all $\kappa_h^e/2\pi \approx \kappa_l^g/2\pi \approx 19$ MHz and $\kappa_h^g/2\pi \approx \kappa_l^e/2\pi \approx 14$ MHz. After this crossing point, we observe that while $|\chi_l| > |\chi_h|$ for $\Delta_{qr}/2\pi \geq -1.6$ GHz, we find $\kappa_l^{g/e} > \kappa_h^{g/e}$, which we exploit in Sec. II. The detailed parameters are summarized in Table I.

II. SINGLE-SHOT READOUT

We perform single-shot readout for the qubit—readout-resonator detunings $\Delta_{qr}/2\pi \in [-2.7, -2.4, -1.9, -1.6, -1.3]$ GHz (see Fig. 2) as a function of the readout-pulse power and integration time $\tau \in [50, 100, 200, 300, 400]$ ns. Each experiment consists of 10^4 single-shot measurements with the qubit prepared

in the ground or excited state. The detuning is varied by tuning the qubit to a chosen frequency ω_q, using a flux pulse as described above. We use a rectangular readout pulse with a duration of 450 ns convolved with a Gaussian filter of width $\sigma = 0.5$ ns, and integrate the readout signal for a time τ using mode-matched weights [38] to discriminate the ground $|g\rangle$ and excited $|e\rangle$ qubit state responses. The flux pulse lasts longer than the readout pulse. In addition, we use a preselection readout to reduce residual excited state population of the qubit to below 0.1% [39].

We express the readout power as a function of the number of photons in the readout resonator n_g when the qubit is prepared in the ground state, relative to the critical number of photons in the resonator $n_{\mathrm{crit}} = \Delta_{qr}^2/4g^2$ [19] at a given qubit—readout-resonator detuning Δ_{qr}. By measuring the qubit-induced ac-Stark shift $\Delta_{\mathrm{ac}} = 2g^2/\Delta_{qr}$ on the readout resonator at $\Delta_{qr}/2\pi = -2.7$ GHz we can infer the number of photons n_g in the resonator when the qubit is prepared in the ground state, $n_g = \Delta_{\mathrm{ac}}/2(\chi_l+\chi_h)$ (see Appendix D). The photon number n_g at other detunings and the photon number n_e when the qubit is prepared in the excited state for all detunings are inferred using semi-classical analysis, see Appendix B.

The readout drive frequency ω_d is chosen such that the difference in the response in transmission when the qubit is prepared in the ground or excited state $|S_{\mathrm{out,in}}^{|e\rangle} - S_{\mathrm{out,in}}^{|g\rangle}|$, is maximum for the *low* mode, see vertical black dashed lines in Fig. 2 (a). In the theoretical model (see Appendix. C), this choice corresponds to selecting the drive frequency which leads to the largest steady-state displacement between the coherent g-state and e-state Purcell-filter-mode responses. This assumes a fixed weak drive power, such that the response is in the linear regime. We found that this choice consistently leads to a stronger resonator response for the qubit being in the excited state ($n_e > n_g$). We accredit this to a smaller effective linewidth for the excited state for the lower mode, $\kappa_l^e < \kappa_l^g$, see Fig. 2 (d). We find this to be an appropriate choice of drive frequency, as the Kerr nonlinearity imparted on the resonator is weaker for the excited state than the ground state (see Appendix B).

We extract the signal-to-noise ratio (SNR) in terms of power, of the acquired single-shot histograms (see Fig. 3) from a bimodal Gaussian distribution as [38]

$$\mathrm{SNR} \equiv \left| \frac{\mu_g - \mu_e}{(\sigma_g + \sigma_e)/2} \right|^2, \tag{3}$$

where $\mu_{g/e}$ and $\sigma_{g/e}$ are, respectively, the mean and the standard deviation of the Gaussian distributions of the g/e-state responses. In Fig. 3 the solid black line indicates the distance between the means μ_g and μ_e, and radii of the black circles are given by the square root of the diagonal covariance matrix elements of the bimodal Gaussian distribution.

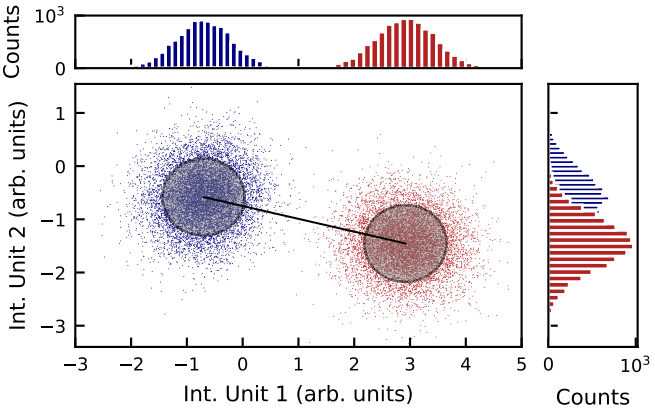

FIG. 3. Single-shot readout histogram for a qubit—readout-resonator detuning of $\Delta_{qr}/2\pi = -1.3$ GHz, a readout integration time of $\tau = 100$ ns, and $n_g/n_{\mathrm{crit}} = 0.93$. We assign the measured state using a bimodal Gaussian mixture model. The marginal distributions of this model along each axis are plotted along the corresponding axis. A solid black line indicates the distance between the means μ_g and μ_e of the Gaussian distributions of the ground and excited state responses. The square root of the covariance matrix diagonal elements of the Gaussian distributions σ_g and σ_e are used as the radii of the black circles.

We characterize the measurement by the average assignment error ε_a for two-state readout, limited by the overlap error between the Gaussian distributions and the qubit lifetime T_1, defined as [38]

$$\begin{aligned} \varepsilon_a &= 1 - \mathcal{F}_{g,e} \\ &= \left[P(e|g) + P(g|e)\right]/2 \\ &\gtrsim \frac{1}{2}\left[1 - \mathrm{erf}\left(\sqrt{\mathrm{SNR}/8}\right)\right] + \frac{\tau}{2T_1}, \end{aligned} \tag{4}$$

where $P(i|j)$ is the probability of measuring the state $|i\rangle$ when having prepared the state $|j\rangle$, and where the average two-state readout fidelity $\mathcal{F}_{g,e}$ characterizes the quality of the readout. The factor two present in the T_1 limit term arises from the fact that only $P(g|e)$ is affected by loss events.

In Fig. 4 (a) we present the lowest measured average assignment errors ε_a as a function of the qubit—readout-resonator detuning $\Delta_{qr}/2\pi \in [-2.7, -2.4, -1.9, -1.6, -1.3]$ GHz and as a function of the readout integration times $\tau \in [50,\ 100,\ 200,\ 300,\ 400]$ ns. We observe that $\varepsilon_a < 1 \times 10^{-2}$ for $\tau \geq 100$ ns. When $\tau \geq 100$ ns and for all qubit—readout-resonator detunings, the variations in the average assignment error are small and stay between $2.5 \times 10^{-3} \leq \varepsilon_a \leq 1 \times 10^{-2}$, except for $\Delta_{qr}/2\pi = -2.7$ GHz and $\tau = 100$ ns. The best assignment error $\varepsilon_a = 2.5 \times 10^{-3}$ is reached at $\tau = 100$ ns and $\Delta_{qr}/2\pi = -1.3$ GHz.

This observation suggests that beyond this integration time the assignment error is no longer limited by the SNR, which would continue to increase for longer readout times. This is further demonstrated in Fig. 4 (b), where

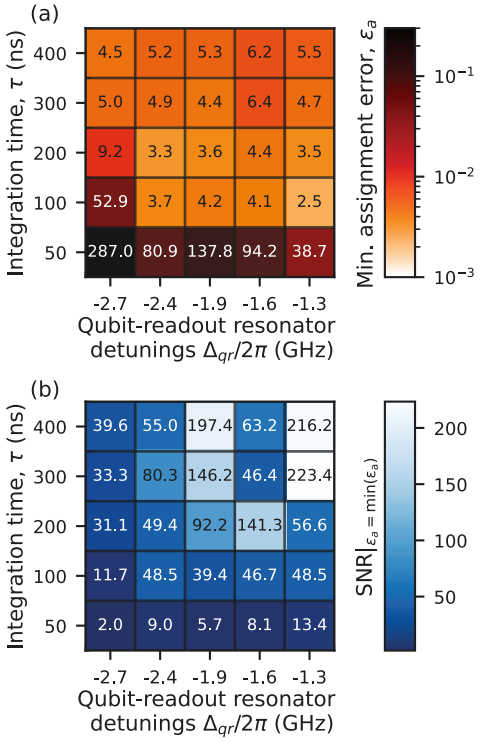

FIG. 4. **(a)** Minimum assignment error ε_a measured as a function of the qubit—readout-resonator detuning $\Delta_{qr}/2\pi \in [-2.7, -2.4, -1.9, -1.6, -1.3]$ GHz and the readout integration time $\tau \in [50,\ 100,\ 200,\ 300,\ 400]$ ns. Annotated values are in *per mille* unit of probability: the lowest assignment error $\varepsilon_a = 2.5 \times 10^{-3}$ is reached at a qubit—readout-resonator detuning of $\Delta_{qr}/2\pi = -1.3$ GHz, for an integration time of $\tau = 100$ ns. **(b)** Measured SNR corresponding to the minimum assignment error in (a).

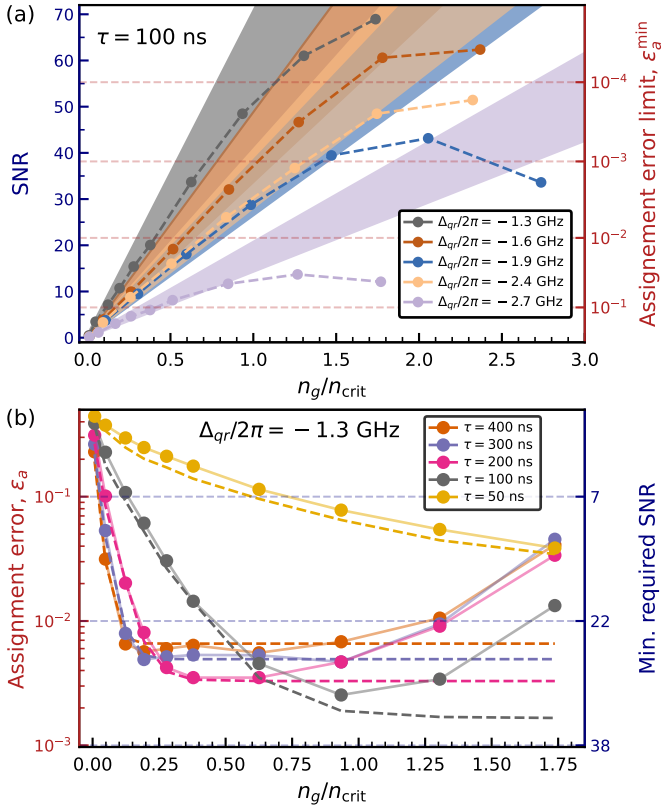

FIG. 5. **(a)** SNR as a function of n_g/n_{crit} for the indicated qubit—readout-resonator detunings Δ_{qr} and at a fixed readout-integration time of $\tau = 100$ ns. The shaded regions provide estimates from the analytical solution in the linear regime, detailed in Appendix C. **(b)** Average assignment error ε_a as a function of n_g/n_{crit} for the indicated readout integration times τ at a fixed qubit—readout-resonator detuning $\Delta_{qr}/2\pi = -1.3$ GHz. Solid lines are plotted for ease of visualization. Dashed lines correspond to the theoretical limit imposed by the relaxation time of the qubit given by $\varepsilon_a^{\mathrm{min}} = 0.5[1 - \mathrm{erfc}(\sqrt{\mathrm{SNR}/8})] + \tau/2T_1$.

we indicate the measured SNR corresponding to each lowest measured average assignment error in Fig. 4 (a) as a function of the same qubit—readout-resonator detuning and readout-integration-time range. We observe that a SNR ≥ 30 leads to $2.5 \times 10^{-3} \leq \varepsilon_a \leq 1 \times 10^{-2}$ for $\tau \geq 100$ ns and for all detunings Δ_{qr}. On the other hand, SNR ≤ 14 leads to a larger assignment error $3.87 \times 10^{-2} \leq \varepsilon_a \leq 2.87 \times 10^{-1}$. In particular, we find that the best assignment error $\varepsilon_a = 2.5 \times 10^{-3}$ is reached for SNR $= 48.5$. An SNR ≥ 200 leads to assignment errors on the same order as an SNR ≈ 50 (see for example at $\Delta_{qr}/2\pi = -1.6$ GHz compared to at $\Delta_{qr}/2\pi = -1.3$ GHz, with $\tau \in [300,\ 400]$ ns).

We next compare the readout performance in terms of SNR at different qubit—readout-resonator detunings as a function of the readout power n_g/n_{crit}, for a fixed integration time $\tau = 100$ ns, see Fig. 5 (a). The shaded regions indicate the theoretical SNR prediction from the linear response to the readout drive power [40]

$$\mathrm{SNR}(t) = 2\eta\kappa_p \int_0^t |\beta_e(t') - \beta_g(t')|^2 dt', \qquad (5)$$

where η is the measurement efficiency. This expression is in line with Eq. (3). We note that the SNR for the smallest detunings $\Delta_{qr} \in [-1.6, -1.3]$ GHz is significantly higher, which we accredit to the increase in the linewidth of the targeted lower mode $\kappa_l^{g/e}$, see Fig. 2 (d). This increased linewidth results in the pointer states $\beta_{g/e}(t)$ reaching the steady state faster, thus maximizing the SNR rate. The shaded region contains the upper- and lower- bound estimates of the SNR based on uncertainties in the model parameters, see Appendix C.

In all instances we observe a saturation of the SNR at a readout power $n_g \gtrsim n_{\mathrm{crit}}$, where the dispersive approximation is known to break down [32, 41, 42]. This is in part due to the broadening of the pointer states caused by the qubit-induced Kerr nonlinearity of the resonator (see Fig. 9), measurement-induced state transitions [43] and ionization [44].

Finally, we compare in Fig. 5 (b) the average assignment error ε_a at different readout integration times $\tau \in$

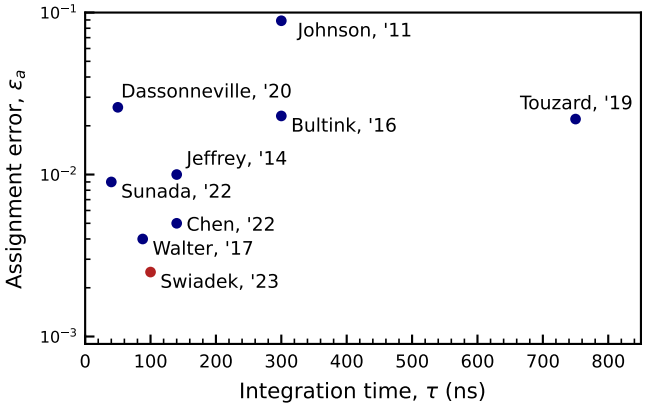

FIG. 6. Two-level average assignment error reported in Johnson *et al.* [41], Jeffrey *et al.* [30], Bultink *et al.* [45], Walter *et al.* [3], Dassonneville *et al.* [46], Touzard *et al.* [47] (blue circles) and in this work (red circle) as a function of the readout integration time. Jurcevic *et al.* [48] reached a two-level assignment error of 3.5×10^{-2} using the excited state promotion technique [49] and is not plotted here.

[50, 100, 200, 300, 400] ns for a fixed qubit—readout-resonator detuning $\Delta_{qr}/2\pi = -1.3\,\text{GHz}$ as a function of the readout power n_g/n_{crit}. For $n_g/n_{\text{crit}} < 1$, we find excellent agreement between the experimental data (dots) and the approximate theoretical limit (dashed lines) in Eq. (4). Here, we note that the 50 ns measurement is clearly limited by the SNR, which consistently improves at higher drive powers. For 100 ns, the minimum assignment error $\varepsilon_a^{\text{min}} = 2.5 \times 10^{-3}$ is limited by the intrinsic lifetime of the qubit T_1 rather than the SNR, which can be seen by comparison to the calculated assignment error (gray dashed line), which plateaus at higher readout powers. We notice a distinct upturn in the 100 ns measurement at higher drive powers, which we attribute to non-linearities and measurement-induced transitions.

For the longer readout times $\tau \in [200, 300, 400]$ ns, the minimum assignment error is obtained at *lower* drive powers, since these measurements reach a larger SNR. Given that the assignment fidelity is limited by the qubit lifetime, the increase in SNR by increasing drive power has little impact on the final assignment error ε_a, as indicated by the plateaus (dashed lines).

III. CONCLUSION

We have demonstrated beyond-state-of-the-art single shot readout reaching a minimum assignment error of 2.5×10^{-3} in only 100 ns when reducing the qubit detuning from the resonator by applying a flux pulse, see our work in perspective with other techniques in Fig. 6. We provided new insights on dispersive readout for a qubit—readout-resonator—Purcell-filter system, in a strongly hybridized regime where the coupling strength between the readout resonator and the Purcell filter J is comparable to the Purcell filter linewidth κ_p. We showed that by probing the dispersive regime *via* flux pulses we can increase the effective decay rate of the targeted readout mode, thus allowing us to reach larger SNR in a shorter integration time.

Our findings open opportunities to study other regimes and help optimize the readout parameters in the design stage of quantum processors in order to adjust the effective decay rate of the readout mode depending on the applications. For instance, we expect this work to help reduce the readout contribution to the quantum error correction cycle time on superconducting qubit platforms [7], without compromising on the readout fidelity constraints. Such techniques combined with machine learning methods as in Ref. [50] for the optimization of pulse shapes, could continue to decrease readout times while maintaining low readout errors.

ACKNOWLEDGEMENTS

The team in Zurich thanks Johannes Herrmann and Stefania Lažar for contributions to the experimental setup. The team in Sherbrooke thanks Cristóbal Lledó and Catherine Leroux for insightful discussions.

The team in Zurich acknowledges financial support by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), through the U.S Army Research Office grant W911NF-16-1-0071, by the EU Flagship on Quantum Technology H2020-FETFLAG-2018-03 project 820363 OpenSuperQ, by the National Center of Competence in Research 'Quantum Science and Technology' (NCCR QSIT), a reseach instrument of the Swiss National Science Foundation (SNSF, grant number 51NF40-185902), by the SNSF R'Equip grant 206021-170731, by the EU-programme H2020-FETOPEN project 828826 Quromorphic and by ETH Zurich. S.K acknowledges financial support from Fondation Jean-Jacques et Félicia Lopez-Loreta and the ETH Zurich Foundation. The team in Sherbrooke acknowledges the financial support by NSERC, the Canada First Research Excellence Fund, and the Ministère de l'Économie et de l'Innovation du Québec. Support is also acknowledged from the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Quantum Systems Accelerator. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the ODNI, IARPA or the U.S Government.

The authors declare no competing interests.

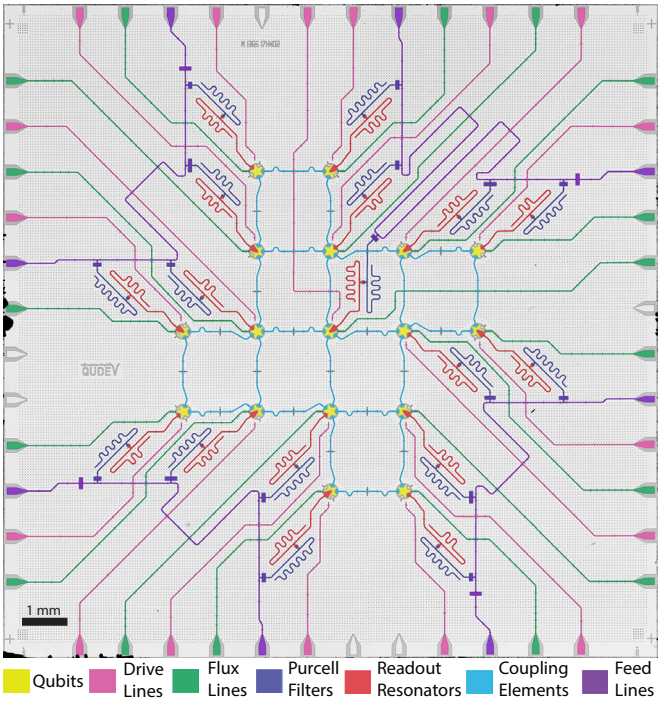

FIG. 7. False-colour micrograph of the 17-qubit device used for the experiment, adapted from Ref. [7]; the scale bar denotes 1 mm. The experiment was realized using the qubit in the feedline at the bottom left corner, as presented in Fig. 1.

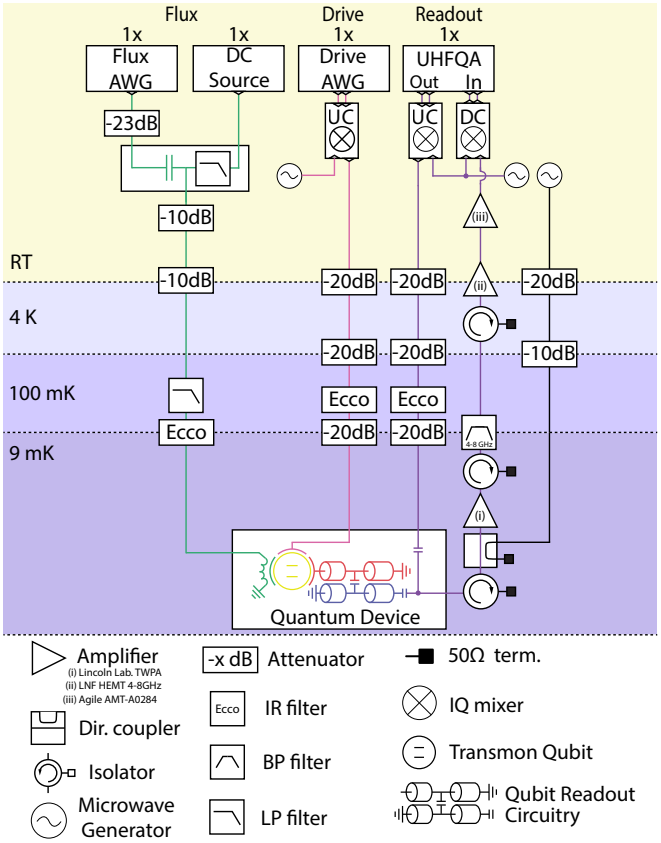

FIG. 8. Schematic of the wiring and control electronics. The qubit (yellow) on the quantum device is connected to the room-temperature electronics *via* flux lines (green), drive lines (pink) and readout lines (purple) through its readout resonator (red) and Purcell filter (blue). The background colors indicate to the temperature stages of the experimental setup.

AUTHOR CONTRIBUTIONS

S.K., P.M. and F.S. planned the experiments with support from all co-authors, and S.K., F.S. performed the experiments. F.S., R.S., and S.K. analyzed the data. R.S. and F.S. worked on the theory. F.S. and A.R. designed the device, and A.R., S.K., D.C.Z. and G.J.N. fabricated the device. C.H., N.L. and A.R. developed the experimental software framework. S.K., A.R., F.S., C.H. and N.L. contributed to the experimental setup and maintained it. F.S., R.S. and S.K. prepared the figures for the manuscript and S.K., A.W., A.B. and Q.F. provided feedback. F.S. and R.S. wrote the manuscript with inputs from all co-authors. A.W., S.K. and A.B. supervised the work.

Appendix A: Experimental setup and device characterization

We used a qubit of a 17-qubit quantum device, shown in Fig. 7, to perform the experiment. We fabricated the 17-qubit quantum processor by sputtering a niobium 150-nm-thin film onto a high-resistivity intrinsic silicon substrate. All coplanar waveguides, capacitors and qubit islands were patterned using photolithography and reactive-ion etching. The aluminium-titanium-aluminium trilayer airbridges establish a well-connected ground plane on the device and connect signal lines split by crossings. We fabricated aluminium-based Josephson junctions using shadow evaporation of aluminium through a resist mask defined by electron-beam lithography.

We characterized the properties of the qubit using spectroscopy and standard time-domain measurements. The qubit has an idling frequency $\omega_q/2\pi = 4.144\,\mathrm{GHz}$, an anharmonicity $\alpha/2\pi = -181\,\mathrm{MHz}$, a lifetime $T_1 = 30.4\,\mu\mathrm{s}$, a Ramsey decay time $T_2^* = 29.2\,\mu\mathrm{s}$, and an echo decay time $T_2^e = 33.9\,\mu\mathrm{s}$.

Following the method in Ref. [26], we fit the transmission amplitude of the readout signal through a feed-line to the function

$$|S_{\mathrm{out,in}}|(\omega) = (A + k(\omega - \omega_0))\times \left|\cos(\phi) - e^{i\phi}\frac{\kappa_p(-2i\Delta_r^{g/e})}{4J^2 + (\kappa_p - 2i\Delta_p)(-2i\Delta_r^{g/e})}\right|, \quad \text{(A1)}$$

where A is the amplitude, k describes a tilt in the spectrum centered at ω_0, ϕ is a phase rotation induced by the

qubit—readout-resonator detuning	$\Delta_{qr}/2\pi$	GHz	-2.7	-2.4	-1.9	-1.6	-1.3
Qubit frequency during readout	$\omega_q/2\pi$	MHz	4144	4500	5000	5300	5600
Bare readout resonator frequency	$\omega_{r,b}/2\pi$	MHz	6854.63	6858.02	6857.98	6859.74	6864.86
Dressed readout resonator frequency	$\omega_r^g/2\pi$	MHz	6876.27	6881.98	6896.09	6906.33	6928.43
Purcell filter frequency	$\omega_p/2\pi$	MHz	6899.86	6899.86	6899.86	6899.86	6899.86
Readout drive frequency	$\omega_d/2\pi$	MHz	6857.4	6861.2	6870.0	6874.0	6881.6
Qubit readout resonator coupling	$g_b/2\pi$	MHz	224.32	205.61	211.49	204.2	205.53
Qubit charge-readout resonator coupling	$g/2\pi$	MHz	284.01	271.40	293.71	292.27	302.34
Readout resonator-Purcell filter coupling	$J/2\pi$	MHz	27.9	27.9	27.9	27.9	27.9
Low mode linewidth, qubit in $\|g\rangle$ state	$\kappa_l^g/2\pi$	MHz	10.16	11.61	15.81	19.07	25.00
Low mode linewidth, qubit in $\|e\rangle$ state	$\kappa_l^e/2\pi$	MHz	8.88	10.03	12.66	14.84	19.87
High mode linewidth, qubit in $\|g\rangle$ state	$\kappa_h^g/2\pi$	MHz	23.86	22.41	18.21	14.95	9.02
High mode linewidth, qubit in $\|e\rangle$ state	$\kappa_h^e/2\pi$	MHz	25.14	23.99	21.36	19.18	14.15
Low mode dispersive shift	$2\chi_l/2\pi$	MHz	-4.17	-4.35	-6.11	-6.69	-6.31
High mode dispersive shift	$2\chi_h/2\pi$	MHz	-1.50	-1.90	-4.24	-6.64	-13.18
Critical readout resonator photon number	n_{crit}		23.14	19.26	10.42	7.55	4.83

TABLE I. List of readout parameters extracted for the qubit—readout-resonator detunings $\Delta_{qr}/2\pi$ spanning -2.7 GHz to -1.3 GHz using pulsed-spectroscpy measurements.

capacitive couplings to other lines, κ_p is the external coupling rate of the Purcell filter, $\Delta_p = \omega - \omega_p$ is the detuning between the drive frequency ω and the Purcell-filter frequency ω_p, and $\Delta_r^{g/e} = \omega - \omega_r^{g/e}$ is the detuning between the drive frequency and the resonator frequency conditioned on the state of the qubit. The relevant parameters for the studied qubit at the indicated qubit-resonator detunings are provided in Table I.

We installed the device in a magnetically-shielded sample holder mounted at the base plate (9 mK) of a cryogenic measurement setup [51] and connected it to the control and measurement electronics as shown in Fig. 8. We use a DC signal to generate a current inducing a magnetic flux in the SQUID-loop of the transmon qubit, to control its idle frequency. We use arbitrary waveform generators to apply voltage pulses (2.4 GSa/s sampling rate) to the qubit to tune its frequency for readout. The DC and AWG signals are combined using a bias-tee. A precompensation of distortions in the flux line is applied, as in Ref. [7].

We perform the single-shot readout experiments with an ultra-high frequency quantum analyzer (UHFQA) by using an IQ-mixer to upconvert the frequency-multiplexed readout pulses from an intermediate frequency signal sampled at 1.8 GSa/s to the gigahertz frequency range of the readout circuitry. At the output of the device feedline, the readout signal passes through a wide-bandwidth near-quantum limited traveling-wave parametric amplifier (TWPA) [52], a high-electron-mobility transistor (HEMT) amplifier, and room-temperature amplifiers. It is then down-converted with an IQ-mixer and digitally demodulated and integrated in the UHFQA. Further details on the device fabrication, characterization, and the experimental setup, can be found in Ref. [7].

Appendix B: Model

To model the system, we use the Hamiltonian

$$\begin{aligned}\hat{H}_{trp} &= 4E_c\hat{n}_t^2 - E_J(\Phi)\cos\hat{\varphi}_t \\ &+ \omega_{r,b}\hat{a}^\dagger\hat{a} - ig(\hat{n}_t - n_g)(\hat{a} - \hat{a}^\dagger) \\ &+ \omega_p\hat{f}^\dagger\hat{f} + J(\hat{f}^\dagger - \hat{f})(\hat{a}^\dagger - \hat{a}) \\ &+ 2i\mathcal{E}\sin(\omega_d t)(\hat{f}^\dagger - \hat{f}),\end{aligned} \tag{B1}$$

where $\hat{n}_t$ is the charge operator of the transmon, $\hat{a}$ the readout resonator mode creation operator and $\hat{f}$ the Purcell filter mode creation operator. E_c is the charging energy of the transmon, $E_J(\Phi)$ the flux-tunable Josephson energy of the transmon, $\omega_{r,b}$, ω_p the bare resonator and Purcell filer frequencies, and g, J the transmon-resonator and resonator-Purcell coupling rates respectively. $\mathcal{E}, \omega_d$ are the the drive amplitude and drive frequency. Further, we consider a master equation

$$\dot{\hat{\rho}} = -i[\hat{H}_{trp}, \hat{\rho}] + \kappa_p\mathcal{D}[\hat{f}], \tag{B2}$$

where κ_p is the coupling rate between the Purcell filter and the feedline. We first diagonalize the transmon-resonator subsystem. We follow the notation of Ref. [21], and assume a Kerr-nonlinear oscillator model for the transmon, valid in the low readout power regime. A Schrieffer-Wolff transformation yields an effective Hamiltonian

$$\begin{aligned}\hat{H} &= \bar{\omega}_q\hat{b}^\dagger\hat{b} + \omega_p\hat{f}^\dagger\hat{f} + \omega_r^g\hat{a}^\dagger\hat{a} + 2\chi\hat{a}^\dagger\hat{a}\hat{b}^\dagger\hat{b} \\ &- 2\lambda'\lambda^3E_C\hat{a}^{\dagger 2}\hat{a}^2\hat{b}^\dagger\hat{b} - \frac{\alpha}{2}\hat{b}^{\dagger 2}\hat{b}^2 - \frac{E_c}{2}\lambda^4\hat{a}^{\dagger 2}\hat{a}^2 \\ &+ J\left(\left[1 - 2\lambda\lambda'\hat{b}^\dagger\hat{b}\right]\hat{a}^\dagger\hat{f} + \lambda\hat{b}^\dagger\hat{f} + \text{H.c.}\right),\end{aligned} \tag{B3}$$

where

$$\chi = -g^2E_C/(\Delta_{qr}(\Delta_{qr} - E_C)), \tag{B4}$$

$$\lambda' = \lambda E_c/\left[\Delta_{qr} + E_c(1 - 2\lambda^2)\right] \tag{B5}$$

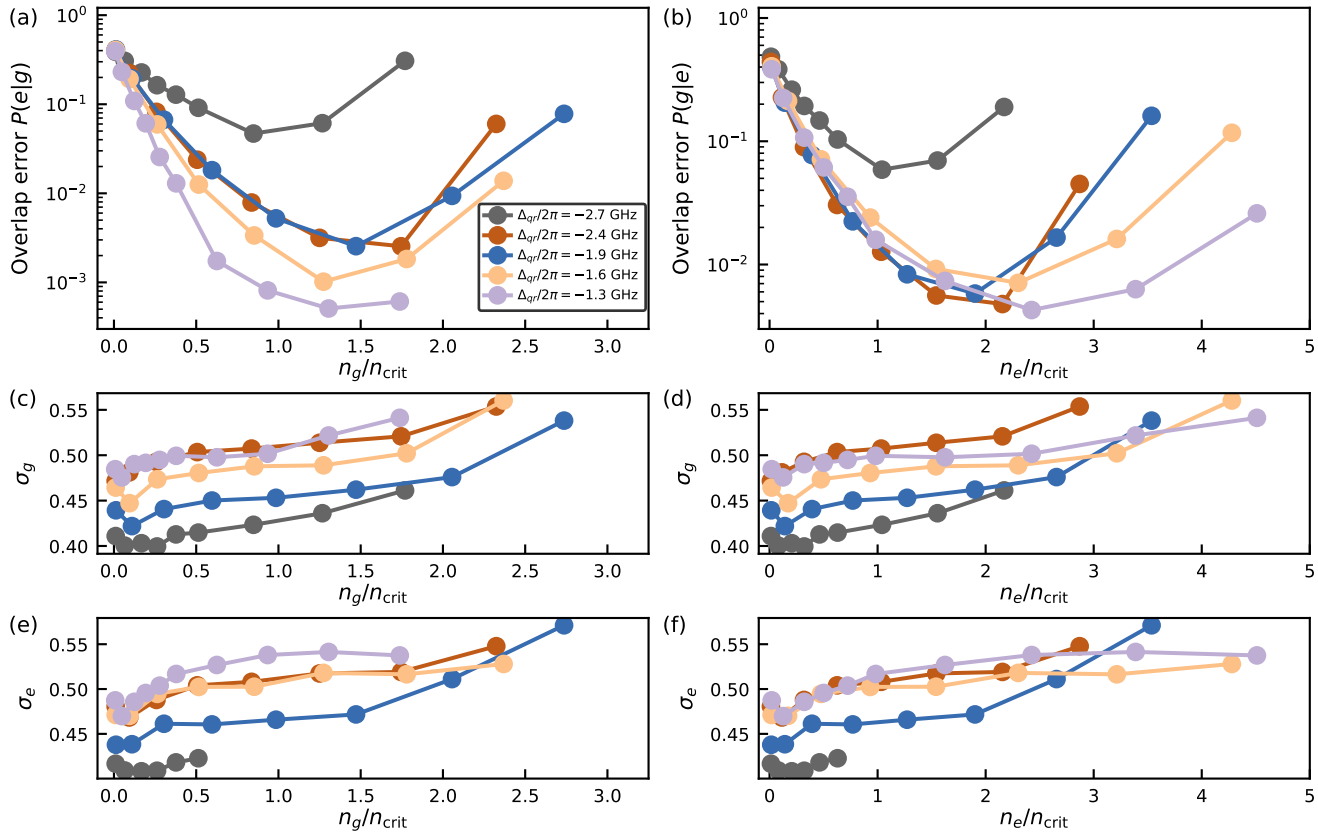

FIG. 9. **(a), (b)** Overlap error $P(e|g)$ and $P(g|e)$ for the indicated qubit—readout-resonator frequency detunings Δ_{qr} at a fixed integration time of $\tau = 100\,\text{ns}$. **(c), (d)** Gaussian width σ_g of the ground state single-shot histogram as a function of n_g/n_{crit} and n_e/n_{crit}. **(e), (f)** Gaussian width σ_e of the ground state single-shot histogram as a function of n_g/n_{crit} and n_e/n_{crit}.

with $\lambda = g/\Delta_{qr}$ and $\Delta_{qr} = \omega_q - \omega_{r,b}$. Further, the qubit and resonator frequencies become Lamb-shifted, with $\bar{\omega}_q \approx \omega_q + g^2/\Delta_{qr}$, $\omega_r^g \approx \omega_{r,b} - g^2/\Delta_{qr}$. The contribution $-2\lambda'\lambda^3 E_C \hat{a}^{\dagger 2}\hat{a}^2\hat{b}^\dagger\hat{b}$ normalizes down the effective Kerr nonlinearity when the qubit is in the excited state, and we note that $K_e/K_g \approx 1 + 4\lambda'/\lambda$, where $4\lambda'/\lambda < 0$ for $\Delta_{qr} < 0$. This nonlinearity leads to a significantly larger increase in the Gaussian width of the ground state response than the excited state response, as seen in Fig. 9 (c,d,e,f) [53, 54]. For this reason, we quote the drive power in the main text as a function of n_g/n_{crit} as opposed to n_e/n_{crit}.

We also note a correlation between the broadening of the ground state response and an increase in the overlap errors $P(g|e)$ and $P(e|g)$ for $n_g/n_{\text{crit}} \gtrsim 1$ where nonlinear effects are expected to be more important, see Fig. 9 (a,b). Frequency renormalizations from the effective coupling of the filter to the qubit are on the order of $J^2\lambda^2/\Delta_{qr}^2$ and can be safely ignored in this regime, which was corroborated by numerical diagonalization of Eq. (B1) and simulation of the master equation. More importantly, we note that the effective coupling strength between the resonator and Purcell filter, $J[1 - 2\lambda\lambda'\hat{b}^\dagger\hat{b}]$, only weakly depends on the qubit state.

Appendix C: Linear response

For sufficiently small drive amplitudes, we can assume negligible impact from the Kerr nonlinearity and take the resonator and filter responses to be linear. As such, we can use the relation

$$\begin{bmatrix}\dot{\alpha}^{g/e}\\ \dot{\beta}^{g/e}\end{bmatrix} = -i\begin{bmatrix}\omega_r^{g/e} - \omega_d & J^{g/e}\\ J^{g/e} & \omega_p - \omega_d - i\frac{\kappa_p}{2}\end{bmatrix}\begin{bmatrix}\alpha^{g/e}\\ \beta^{g/e}\end{bmatrix} + \begin{bmatrix}0\\ \mathcal{E}\end{bmatrix}, \tag{C1}$$

where α and β represent the coherent fields of the readout resonator and Purcell filter respectively, $\mathcal{E}$ is the drive amplitude and $J^{g/e} = J\left[1 - \lambda\lambda'(\langle\hat{\sigma}_z\rangle + 1)\right]$ is the effective readout-resonator—Purcell-filter coupling, and $\omega_r^e = \omega_r^g + 2\chi$. Diagonalizing the equation of motion in the absence of a drive ($\mathcal{E} = \omega_d = 0$) yields eigenvalues

$$\lambda_{l,h}^{g/e} = \frac{\omega_r^{g/e} + \omega_p - i\kappa_p/2}{2} \pm \frac{1}{2}\sqrt{\left(\Delta_{rp}^{g/e} + \frac{i\kappa_p}{2}\right)^2 + 4J^{2,g/e}}. \tag{C2}$$

For $4J \gg \kappa$, the eigenvalues approximately correspond

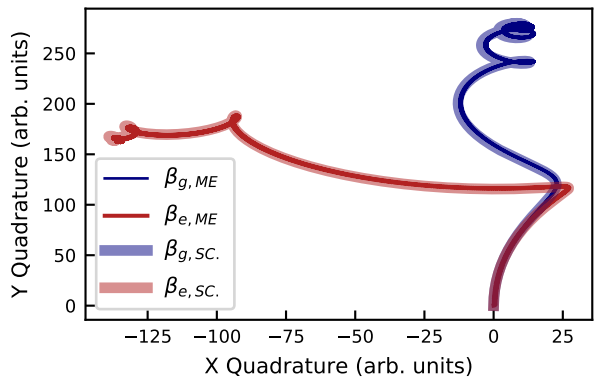

FIG. 10. Example comparison of the master equation $\beta_{g/e,ME}$, against the semi-classical trajectories $\beta_{g/e,SC}$, plotted in the phase space of the Purcell filter mode for $\Delta_{qr}/2\pi = -1.3$ GHz and $\mathcal{E}/2\pi = 10$ MHz. Solid lines correspond to the master equation expectations of the resonator and filter modes, $\langle a\rangle$, and $\langle f\rangle$, respectively, where $\langle X\rangle = \mathrm{Re}\{a, f\}$, $\langle Y\rangle = \mathrm{Im}\{a, f\}$. Transparent lines correspond to the semi-classical solution.

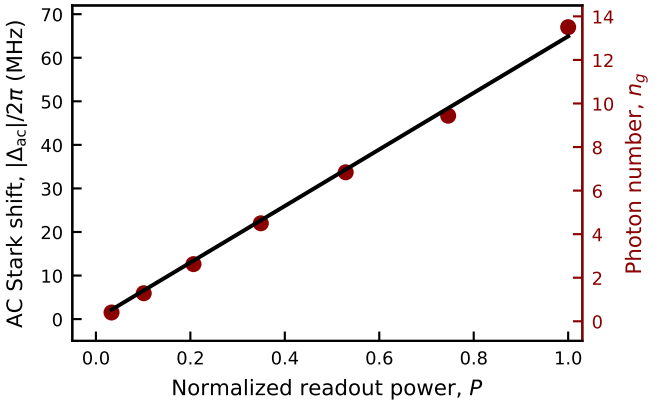

FIG. 11. Measured (dots) ac-Stark shift Δ_{ac} of the qubit prepared in the ground state at $\omega_q/2\pi = 4.14$ GHz as a function of the normalized readout power (real range spans 2.0 µV to 12.5 µV). The corresponding inferred resonator photon number $n_g = \Delta_{ac}/2(\chi_l + \chi_h)$ (solid line) is shown on the right axis.

to a normal mode, where the indices l, h corresponds to the *low* and *high* mode respectively. In this fashion, the real and imaginary components of $\lambda_{l,h}^{g/e}$ correspond to the frequency and linewidth of these *low* and *high* modes

$$\omega_{l,h}^{g/e} = \mathrm{Re}[\lambda_{l,h}^{g/e}], \quad \kappa_{l,h}^{g/e} = -2\mathrm{Im}[\lambda_{l,h}^{g/e}]. \tag{C3}$$

To obtain a qualitative understanding of the eigenvalues, we perform an expansion of the square root in $(\Delta_{rp}^{g/e} + i\kappa/2)$. Assuming $J^{g/e} \approx J$, this yields

$$\begin{aligned}\lambda_{l,h}^{g/e} = &\frac{\omega_r^{g/e} + \omega_p - i\kappa_p/2}{2} \\ &\pm \left(J + \frac{\Delta_{rp}^{2ge} - i\Delta_{rp}^{g/e}\kappa_p - \kappa^2/4}{8J}\right).\end{aligned} \tag{C4}$$

Consequently, we see that the frequency $\omega_{l,h}^{g/e}$ of the two sets of normal modes are approximately separated by $2J$, with the relative dispersive shift of each mode $\chi_{l,h}$ between the low and high mode being

$$\begin{aligned}2\chi_l &= \omega_l^e - \omega_l^g \approx \chi - \frac{\Delta_{rp}^g\chi + \chi^2}{2J}, \\ 2\chi_h &= \omega_h^e - \omega_h^g \approx \chi + \frac{\Delta_{rp}^g\chi + \chi^2}{2J}, \\ \kappa_l^{g/e} &\approx \frac{\kappa_p}{2} + \frac{\Delta_{rp}^{g/e}\kappa_p}{4J}, \\ \kappa_h^{g/e} &\approx \frac{\kappa_p}{2} - \frac{\Delta_{rp}^{g/e}\kappa_p}{4J}.\end{aligned} \tag{C5}$$

Noting that $\Delta_{rp}^e < \Delta_{rp}^g$, we see that $\kappa_h^g < \kappa_h^e$ and $\kappa_l^g > \kappa_l^e$ for $\Delta_{rp}^g < 0$, and vice versa for $\Delta_{rp}^g > 0$.

The relevant frequencies and linewidths extracted from the normal-mode Hamiltonian in Eq. (C1) are plotted in Fig. 2.

The steady state expressions are found to be

$$\begin{bmatrix}\alpha_{ss}^{g/e} \\ \beta_{ss}^{g/e}\end{bmatrix} = \frac{\mathcal{E}}{\Delta^{g/e}(\Delta_p - i\kappa_p/2) - J^2}\begin{bmatrix}-J \\ \Delta^{g/e}\end{bmatrix}. \tag{C6}$$

Using this expression, the full time-dependent response takes the form

$$\begin{aligned}\beta^{g/e}(t) = \beta_{ss}^{g/e} &- \mathcal{E}\frac{\lambda_h - (\omega_r + \chi\langle\sigma_z\rangle)}{d}\frac{e^{-i(\lambda_h - \omega_d)t}}{(\lambda_h - \omega_d)} \\ &+ \mathcal{E}\frac{\lambda_l - (\omega_r + \chi\langle\sigma_z\rangle)}{d}\frac{e^{-i(\lambda_l - \omega_d)t}}{(\lambda_l - \omega_d)},\end{aligned} \tag{C7}$$

where $d = \sqrt{(-\Delta_{rp} - i\kappa_p/2 - \chi\langle\sigma_z\rangle)^2 + 4J^2}$.

We then use Eq. (C7) to express the SNR as [40]

$$\mathrm{SNR}(t) = 2\eta\kappa_p\int_0^t |\beta_e(t') - \beta_g(t')|^2 dt', \tag{C8}$$

where η is the measurement efficiency. We note that this expression is the square of the often used expression but is in line with Eq. (3). We then plot these results for Fig. 5 (a) allowing for a ±1 MHz deviation in the calculated values of g, J, ω_r and κ_p to allow for uncertainties in the fitted parameters and nonidealities caused by spurious couplings to two-level systems, alongside a variation of up to 5% in the measurement efficiency at different frequencies. The shaded region contains the upper and low bound estimates of the SNR based on these uncertainties.

Finally, we verify the validity of the semiclassical approximation. Negligible difference was noted in the trajectories in phase space between the expected internal

coherent fields $\langle\hat{a}\rangle$, $\langle\hat{f}\rangle$, calculated by solving the master equation (B2), and the corresponding semiclassical predictions α and β of the resonator and Purcell filter respectively, confirming that the semiclassical model (in Appendix C) captures the state separation at low powers. Example trajectories at low power for the Purcell filter mode are plotted for the $\Delta_{qr}/2\pi = -1.3$ GHz case in Fig. 10.

Appendix D: Photon number and drive power calibration

We measure the ac-Stark shift $\Delta_{\rm ac}$ caused on the qubit prepared in the ground state by the readout resonator as a function of power, see Fig. 11. To this mean we simultaneously apply a readout tone with a length of 0.6 µs of variable power and a π-pulse of variable frequency. We measure the excited state population as a function of the drive pulse frequency for each readout power for the $\omega_q/2\pi = 4.14\,$GHz qubit frequency. The frequency at which the excited state population is maximum corresponds to the instantaneous and ac-Stark shifted qubit frequency.

We determine this frequency using a Gaussian fit. We infer and then calibrate the steady state readout resonator photon number n_g with the qubit prepared in the ground state from $n_g = \Delta_{\rm ac}/2(\chi_l + \chi_h)$ for the chosen drive powers [55]. Then, using the steady state resonator response from Eq. (C1), this allows us to extract the effective drive amplitudes $\mathcal{E}$. The steady state resonator responses for the qubit-resonator detunings $\Delta_{qr}/2\pi \in [-4.5, -5.0, -5.3, -5.6]\,$GHz are then inferred from Eq. (C6) at the same drive powers.

[1] F. Mallet, F. R. Ong, A. Palacios-Laloy, F. Nguyen, P. Bertet, D. Vion, and D. Esteve, Single-shot qubit readout in circuit quantum electrodynamics, Nat. Phys. **5**, 791 (2009).

[2] M. D. Reed, L. DiCarlo, B. R. Johnson, L. Sun, D. I. Schuster, L. Frunzio, and R. J. Schoelkopf, High-fidelity readout in circuit quantum electrodynamics using the Jaynes-Cummings nonlinearity, Phys. Rev. Lett. **105**, 173601 (2010).

[3] T. Walter, P. Kurpiers, S. Gasparinetti, P. Magnard, A. Potočnik, Y. Salathé, M. Pechal, M. Mondal, M. Oppliger, C. Eichler, and A. Wallraff, Rapid, high-fidelity, single-shot dispersive readout of superconducting qubits, Phys. Rev. Appl. **7**, 054020 (2017).

[4] A. Y. Kitaev, Fault-tolerant quantum computation by anyons, Annals of Physics **303**, 2 (2003).

[5] D. P. DiVincenzo, Fault-tolerant architectures for superconducting qubits, Phys. Scr. **2009**, 014020 (2009).

[6] C. K. Andersen, A. Remm, S. Lazar, S. Krinner, N. Lacroix, G. J. Norris, M. Gabureac, C. Eichler, and A. Wallraff, Repeated quantum error detection in a surface code, Nature Physics **16**, 875 (2020).

[7] S. Krinner, N. Lacroix, A. Remm, A. D. Paolo, E. Genois, C. Leroux, C. Hellings, S. Lazar, F. Swiadek, J. Herrmann, G. J. Norris, C. K. Andersen, M. Müller, A. Blais, C. Eichler, and A. Wallraff, Realizing repeated quantum error correction in a distance-three surface code, Nature **605**, 669 (2022).

[8] R. Acharya, I. Aleiner, R. Allen, T. I. Andersen, M. Ansmann, F. Arute, K. Arya, A. Asfaw, J. Atalaya, R. Babbush, D. Bacon, J. C. Bardin, J. Basso, A. Bengtsson, S. Boixo, G. Bortoli, A. Bourassa, J. Bovaird, L. Brill, M. Broughton, B. B. Buckley, D. A. Buell, T. Burger, B. Burkett, N. Bushnell, Y. Chen, Z. Chen, B. Chiaro, J. Cogan, R. Collins, P. Conner, W. Courtney, A. L. Crook, B. Curtin, D. M. Debroy, A. Del Toro Barba, S. Demura, A. Dunsworth, D. Eppens, C. Erickson, L. Faoro, E. Farhi, R. Fatemi, L. Flores Burgos, E. Forati, A. G. Fowler, B. Foxen, W. Giang, C. Gidney, D. Gilboa, M. Giustina, A. Grajales Dau, J. A. Gross, S. Habegger, M. C. Hamilton, M. P. Harrigan, S. D. Harrington, O. Higgott, J. Hilton, M. Hoffmann, S. Hong, T. Huang, A. Huff, W. J. Huggins, L. B. Ioffe, S. V. Isakov, J. Iveland, E. Jeffrey, Z. Jiang, C. Jones, P. Juhas, D. Kafri, K. Kechedzhi, J. Kelly, T. Khattar, M. Khezri, M. Kieferová, S. Kim, A. Kitaev, P. V. Klimov, A. R. Klots, A. N. Korotkov, F. Kostritsa, J. M. Kreikebaum, D. Landhuis, P. Laptev, K.-M. Lau, L. Laws, J. Lee, K. Lee, B. J. Lester, A. Lill, W. Liu, A. Locharla, E. Lucero, F. D. Malone, J. Marshall, O. Martin, J. R. McClean, T. McCourt, M. McEwen, A. Megrant, B. Meurer Costa, X. Mi, K. C. Miao, M. Mohseni, S. Montazeri, A. Morvan, E. Mount, W. Mruczkiewicz, O. Naaman, M. Neeley, C. Neill, A. Nersisyan, H. Neven, M. Newman, J. H. Ng, A. Nguyen, M. Nguyen, M. Y. Niu, T. E. O'Brien, A. Opremcak, J. Platt, A. Petukhov, R. Potter, L. P. Pryadko, C. Quintana, P. Roushan, N. C. Rubin, N. Saei, D. Sank, K. Sankaragomathi, K. J. Satzinger, H. F. Schurkus, C. Schuster, M. J. Shearn, A. Shorter, V. Shvarts, J. Skruzny, V. Smelyanskiy, W. C. Smith, G. Sterling, D. Strain, M. Szalay, A. Torres, G. Vidal, B. Villalonga, C. Vollgraff Heidweiller, T. White, C. Xing, Z. J. Yao, P. Yeh, J. Yoo, G. Young, A. Zalcman, Y. Zhang, N. Zhu, and G. Q. AI, Suppressing quantum errors by scaling a surface code logical qubit, Nature **614**, 676 (2023).

[9] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters, Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels, Phys. Rev. Lett. **70**, 1895 (1993).

[10] D. Gottesman and I. L. Chuang, Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations, Nature **402**, 390 (1999).

[11] L. Steffen, Y. Salathe, M. Oppliger, P. Kurpiers, M. Baur, C. Lang, C. Eichler, G. Puebla-Hellmann, A. Fedorov, and A. Wallraff, Deterministic quantum teleportation with feed-forward in a solid state system, Nature **500**, 319 (2013).

[12] J. Qiu, Y. Liu, J. Niu, L. Hu, Y. Wu, L. Zhang, W. Huang, Y. Chen, J. Li, S. Liu, Y. Zhong, L. Duan, and D. Yu, Deterministic quantum teleportation between distant superconducting chips, arXiv:2302.08756

10.48550/arXiv:2302.08756 (2023).
[13] C. H. Bennett, H. J. Bernstein, S. Popescu, and B. Schumacher, Concentrating partial entanglement by local operations, Phys. Rev. A **53**, 2046 (1996).
[14] S. Bravyi and A. Kitaev, Universal quantum computation with ideal clifford gates and noisy ancillas, Phys. Rev. A **71**, 022316 (2005).
[15] J. E. Johnson, C. Macklin, D. H. Slichter, R. Vijay, E. B. Weingarten, J. Clarke, and I. Siddiqi, Heralded state preparation in a superconducting qubit, Phys. Rev. Lett. **109**, 050506 (2012).
[16] D. Ristè, J. G. van Leeuwen, H.-S. Ku, K. W. Lehnert, and L. DiCarlo, Initialization by measurement of a superconducting quantum bit circuit, Phys. Rev. Lett. **109**, 050507 (2012).
[17] Y. Salathé, M. Mondal, M. Oppliger, J. Heinsoo, P. Kurpiers, A. Potočnik, A. Mezzacapo, U. Las Heras, L. Lamata, E. Solano, S. Filipp, and A. Wallraff, Digital quantum simulation of spin models with circuit quantum electrodynamics, Phys. Rev. X **5**, 021027 (2015).
[18] J. Herrmann, S. M. Llima, A. Remm, P. Zapletal, N. A. McMahon, C. Scarato, F. Swiadek, C. K. Andersen, C. Hellings, S. Krinner, N. Lacroix, S. Lazar, M. Kerschbaum, D. C. Zanuz, G. J. Norris, M. J. Hartmann, A. Wallraff, and C. Eichler, Realizing quantum convolutional neural networks on a superconducting quantum processor to recognize quantum phases, Nature Comm. **13**, 4144 (2022).
[19] A. Blais, R.-S. Huang, A. Wallraff, S. M. Girvin, and R. J. Schoelkopf, Cavity quantum electrodynamics for superconducting electrical circuits: An architecture for quantum computation, Phys. Rev. A **69**, 062320 (2004).
[20] A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, J. Majer, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Approaching unit visibility for control of a superconducting qubit with dispersive readout, Phys. Rev. Lett. **95**, 060501 (2005).
[21] A. Blais, A. L. Grimsmo, S. M. Girvin, and A. Wallraff, Circuit quantum electrodynamics, Rev. Mod. Phys. **93**, 025005 (2021).
[22] 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, Phys. Rev. A **76**, 042319 (2007).
[23] V. Negirneac, H. Ali, N. Muthusubramanian, F. Battistel, R. Sagastizabal, M. S. Moreira, J. F. Marques, W. J. Vlothuizen, M. Beekman, C. Zachariadis, N. Haider, A. Bruno, and L. DiCarlo, High-fidelity controlled-z gate with maximal intermediate leakage operating at the speed limit in a superconducting quantum processor, Phys. Rev. Lett. **126**, 220502 (2021).
[24] L. DiCarlo, M. D. Reed, L. Sun, B. R. Johnson, J. M. Chow, J. M. Gambetta, L. Frunzio, S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf, Preparation and measurement of three-qubit entanglement in a superconducting circuit, Nature **467**, 574 (2010).
[25] F. W. Strauch, P. R. Johnson, A. J. Dragt, C. J. Lobb, J. R. Anderson, and F. C. Wellstood, Quantum logic gates for coupled superconducting phase qubits, Phys. Rev. Lett. **91**, 167005 (2003).
[26] J. Heinsoo, C. K. Andersen, A. Remm, S. Krinner, T. Walter, Y. Salathé, S. Gasparinetti, J.-C. Besse, A. Potočnik, A. Wallraff, and C. Eichler, Rapid high-fidelity multiplexed readout of superconducting qubits, Phys. Rev. Appl. **10**, 034040 (2018).
[27] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, A. Dunsworth, E. Farhi, B. Foxen, A. Fowler, C. Gidney, M. Giustina, R. Graff, K. Guerin, S. Habegger, M. P. Harrigan, M. J. Hartmann, A. Ho, M. Hoffmann, T. Huang, T. S. Humble, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, J. Kelly, P. V. Klimov, S. Knysh, A. Korotkov, F. Kostritsa, D. Landhuis, M. Lindmark, E. Lucero, D. Lyakh, S. Mandrà, J. R. McClean, M. McEwen, A. Megrant, X. Mi, K. Michielsen, M. Mohseni, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Y. Niu, E. Ostby, A. Petukhov, J. C. Platt, C. Quintana, E. G. Rieffel, P. Roushan, N. C. Rubin, D. Sank, K. J. Satzinger, V. Smelyanskiy, K. J. Sung, M. D. Trevithick, A. Vainsencher, B. Villalonga, T. White, Z. J. Yao, P. Yeh, A. Zalcman, H. Neven, and J. M. Martinis, Quantum supremacy using a programmable superconducting processor, Nature **574**, 505 (2019).
[28] Z. Chen, K. J. Satzinger, J. Atalaya, A. N. Korotkov, A. Dunsworth, D. Sank, C. Quintana, M. McEwen, R. Barends, P. V. Klimov, S. Hong, C. Jones, A. Petukhov, D. Kafri, S. Demura, B. Burkett, C. Gidney, A. G. Fowler, A. Paler, H. Putterman, I. Aleiner, F. Arute, K. Arya, R. Babbush, J. C. Bardin, A. Bengtsson, A. Bourassa, M. Broughton, B. B. Buckley, D. A. Buell, N. Bushnell, B. Chiaro, R. Collins, W. Courtney, A. R. Derk, D. Eppens, C. Erickson, E. Farhi, B. Foxen, M. Giustina, A. Greene, J. A. Gross, M. P. Harrigan, S. D. Harrington, J. Hilton, A. Ho, T. Huang, W. J. Huggins, L. B. Ioffe, S. V. Isakov, E. Jeffrey, Z. Jiang, K. Kechedzhi, S. Kim, A. Kitaev, F. Kostritsa, D. Landhuis, P. Laptev, E. Lucero, O. Martin, J. R. McClean, T. McCourt, X. Mi, K. C. Miao, M. Mohseni, S. Montazeri, W. Mruczkiewicz, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Newman, M. Y. Niu, T. E. O'Brien, A. Opremcak, E. Ostby, B. Pató, N. Redd, P. Roushan, N. C. Rubin, V. Shvarts, D. Strain, M. Szalay, M. D. Trevithick, B. Villalonga, T. White, Z. J. Yao, P. Yeh, J. Yoo, A. Zalcman, H. Neven, S. Boixo, V. Smelyanskiy, Y. Chen, A. Megrant, J. Kelly, and A. I. Google Quantum, Exponential suppression of bit or phase errors with cyclic error correction, Nature **595**, 383 (2021).
[29] M. D. Reed, B. R. Johnson, A. A. Houck, L. DiCarlo, J. M. Chow, D. I. Schuster, L. Frunzio, and R. J. Schoelkopf, Fast reset and suppressing spontaneous emission of a superconducting qubit, Appl. Phys. Lett. **96**, 203110 (2010).
[30] E. Jeffrey, D. Sank, J. Y. Mutus, T. C. White, J. Kelly, R. Barends, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. Megrant, P. J. J. O'Malley, C. Neill, P. Roushan, A. Vainsencher, J. Wenner, A. N. Cleland, and J. M. Martinis, Fast accurate state measurement with superconducting qubits, Phys. Rev. Lett. **112**, 190504 (2014).
[31] N. T. Bronn, Y. Liu, J. B. Hertzberg, A. D. Córcoles, A. A. Houck, J. M. Gambetta, and J. M. Chow, Broadband filters for abatement of spontaneous emission in circuit quantum electrodynamics, Applied Physics Letters **107**, 172601 (2015).
[32] M. Boissonneault, J. M. Gambetta, and A. Blais, Improved superconducting qubit readout by qubit-induced

nonlinearities, Phys. Rev. Lett. **105**, 100504 (2010).
[33] Y. Sunada, S. Kono, J. Ilves, S. Tamate, T. Sugiyama, Y. Tabuchi, and Y. Nakamura, Fast readout and reset of a superconducting qubit coupled to a resonator with an intrinsic purcell filter, Phys. Rev. Applied **17**, 044016 (2022).
[34] M. Khezri, A. Opremcak, Z. Chen, A. Bengtsson, T. White, O. Naaman, R. Acharya, K. Anderson, M. Ansmann, F. Arute, K. Arya, A. Asfaw, J. C. Bardin, A. Bourassa, J. Bovaird, L. Brill, B. B. Buckley, D. A. Buell, T. Burger, B. Burkett, N. Bushnell, J. Campero, B. Chiaro, R. Collins, A. L. Crook, B. Curtin, S. Demura, A. Dunsworth, C. Erickson, R. Fatemi, V. S. Ferreira, L. F. Burgos, E. Forati, B. Foxen, G. Garcia, W. Giang, M. Giustina, R. Gosula, A. G. Dau, M. C. Hamilton, S. D. Harrington, P. Heu, J. Hilton, M. R. Hoffmann, S. Hong, T. Huang, A. Huff, J. Iveland, E. Jeffrey, J. Kelly, S. Kim, P. V. Klimov, F. Kostritsa, J. M. Kreikebaum, D. Landhuis, P. Laptev, L. Laws, K. Lee, B. J. Lester, A. T. Lill, W. Liu, A. Locharla, E. Lucero, S. Martin, M. McEwen, A. Megrant, X. Mi, K. C. Miao, S. Montazeri, A. Morvan, M. Neeley, C. Neill, A. Nersisyan, J. H. Ng, A. Nguyen, M. Nguyen, R. Potter, C. Quintana, C. Rocque, P. Roushan, K. Sankaragomathi, K. J. Satzinger, C. Schuster, M. J. Shearn, A. Shorter, V. Shvarts, J. Skruzny, W. C. Smith, G. Sterling, M. Szalay, D. Thor, A. Torres, B. W. K. Woo, Z. J. Yao, P. Yeh, J. Yoo, G. Young, N. Zhu, N. Zobrist, D. Sank, A. Korotkov, Y. Chen, and V. Smelyanskiy, Measurement-induced state transitions in a superconducting qubit: Within the rotating wave approximation, arXiv 10.48550/arXiv:2212.05097 (2022).
[35] M. Malekakhlagh, W. Shanks, H. Paik, M. Malekakhlagh, W. Shanks, and H. Paik, Optimization of the resonator-induced phase gate for superconducting qubits, Phys. Rev. A **105**, 022607 (2022).
[36] M. D. Hutchings, J. B. Hertzberg, Y. Liu, N. T. Bronn, G. A. Keefe, M. Brink, J. M. Chow, and B. L. T. Plourde, Tunable superconducting qubits with flux-independent coherence, Phys. Rev. Appl. **8**, 044003 (2017).
[37] C. W. Gardiner and M. J. Collett, Input and output in damped quantum systems: Quantum stochastic differential equations and the master equation, Phys. Rev. A **31**, 3761 (1985).
[38] J. Gambetta, W. A. Braff, A. Wallraff, S. M. Girvin, and R. J. Schoelkopf, Protocols for optimal readout of qubits using a continuous quantum nondemolition measurement, Phys. Rev. A **76**, 012325 (2007).
[39] P. Magnard, P. Kurpiers, B. Royer, T. Walter, J.-C. Besse, S. Gasparinetti, M. Pechal, J. Heinsoo, S. Storz, A. Blais, and A. Wallraff, Fast and unconditional all-microwave reset of a superconducting qubit, Phys. Rev. Lett. **121**, 060502 (2018).
[40] C. C. Bultink, B. Tarasinski, N. Haandbæk, S. Poletto, N. Haider, D. J. Michalak, A. Bruno, and L. DiCarlo, General method for extracting the quantum efficiency of dispersive qubit readout in circuit qed, Appl. Phys. Lett. **112**, 092601 (2018).
[41] J. E. Johnson, E. M. Hoskinson, C. Macklin, D. H. Slichter, I. Siddiqi, and J. Clarke, Dispersive readout of a flux qubit at the single-photon level, Phys. Rev. B **84**, 220503 (2011).
[42] Z. K. Minev, S. O. Mundhada, S. Shankar, P. Reinhold, R. Gutiérrez-Jáuregui, R. J. Schoelkopf, M. Mirrahimi, H. J. Carmichael, and M. H. Devoret, To catch and reverse a quantum jump mid-flight, Nature **570**, 200 (2019).
[43] D. Sank, Z. Chen, M. Khezri, J. Kelly, R. Barends, B. Campbell, Y. Chen, B. Chiaro, A. Dunsworth, A. Fowler, and et al., Measurement-induced state transitions in a superconducting qubit: Beyond the rotating wave approximation, Physical Review Letters **117**, 190503 (2016).
[44] R. Shillito, A. Petrescu, J. Cohen, J. Beall, M. Hauru, M. Ganahl, A. G. Lewis, G. Vidal, and A. Blais, Dynamics of transmon ionization, Phys. Rev. Appl. **18**, 034031 (2022).
[45] C. C. Bultink, M. A. Rol, T. E. O'Brien, X. Fu, B. C. S. Dikken, C. Dickel, R. F. L. Vermeulen, J. C. de Sterke, A. Bruno, R. N. Schouten, and L. DiCarlo, Active resonator reset in the nonlinear dispersive regime of circuit QED, Phys. Rev. Applied **6**, 034008 (2016).
[46] R. Dassonneville, T. Ramos, V. Milchakov, L. Planat, E. Dumur, F. Foroughi, J. Puertas, S. Leger, K. Bharadwaj, J. Delaforce, C. Naud, W. Hasch-Guichard, J. J. García-Ripoll, N. Roch, and O. Buisson, Fast high-fidelity quantum nondemolition qubit readout via a nonperturbative cross-kerr coupling, Phys. Rev. X **10**, 011045 (2020).
[47] S. Touzard, A. Kou, N. E. Frattini, V. V. Sivak, S. Puri, A. Grimm, L. Frunzio, S. Shankar, and M. H. Devoret, Gated conditional displacement readout of superconducting qubits, Phys. Rev. Lett. **122**, 080502 (2019).
[48] P. Jurcevic, A. Javadi-Abhari, L. S. Bishop, I. Lauer, D. F. Bogorin, M. Brink, L. Capelluto, O. Günlük, T. Itoko, N. Kanazawa, A. Kandala, G. A. Keefe, K. Krsulich, W. Landers, E. P. Lewandowski, D. T. McClure, G. Nannicini, A. Narasgond, H. M. Nayfeh, E. Pritchett, M. B. Rothwell, S. Srinivasan, N. Sundaresan, C. Wang, K. X. Wei, C. J. Wood, J.-B. Yau, E. J. Zhang, O. E. Dial, J. M. Chow, and J. M. Gambetta, Demonstration of quantum volume 64 on a superconducting quantum computing system, Quantum Science and Technology **6**, 025020 (2021).
[49] S. S. Elder, C. S. Wang, P. Reinhold, C. T. Hann, K. S. Chou, B. J. Lester, S. Rosenblum, L. Frunzio, L. Jiang, and R. J. Schoelkopf, High-fidelity measurement of qubits encoded in multilevel superconducting circuits, Phys. Rev. X **10**, 011001 (2020).
[50] B. Lienhard, A. Vepsäläinen, L. C. G. Govia, C. R. Hoffer, J. Y. Qiu, D. Ristè, M. Ware, D. Kim, R. Winik, A. Melville, B. Niedzielski, J. Yoder, G. J. Ribeill, T. A. Ohki, H. K. Krovi, T. P. Orlando, S. Gustavsson, and W. D. Oliver, Deep-neural-network discrimination of multiplexed superconducting-qubit states, Phys. Rev. Applied **17**, 014024 (2022).
[51] S. Krinner, S. Storz, P. Kurpiers, P. Magnard, J. Heinsoo, R. Keller, J. Lütolf, C. Eichler, and A. Wallraff, Engineering cryogenic setups for 100-qubit scale superconducting circuit systems, EPJ Quantum Technology **6**, 2 (2019).
[52] C. Macklin, K. O'Brien, D. Hover, M. E. Schwartz, V. Bolkhovsky, X. Zhang, W. D. Oliver, and I. Siddiqi, A near-quantum-limited Josephson traveling-wave parametric amplifier, Science **350**, 307 (2015).
[53] N. Bartolo, F. Minganti, W. Casteels, and C. Ciuti, Exact steady state of a kerr resonator with one- and two-photon driving and dissipation: Controllable wigner-function multimodality and dissipative phase transitions,

Phys. Rev. A **94**, 033841 (2016).
[54] D. Roberts and A. A. Clerk, Driven-dissipative quantum kerr resonators: New exact solutions, photon blockade and quantum bistability, Phys. Rev. X **10**, 021022 (2020).
[55] D. I. Schuster, A. Wallraff, A. Blais, L. Frunzio, R.-S. Huang, J. Majer, S. M. Girvin, and R. J. Schoelkopf, ac Stark shift and dephasing of a superconducting qubit strongly coupled to a cavity field, Phys. Rev. Lett. **94**, 123602 (2005).

Conclusion

"Farewell, King under the Mountain! This is a bitter adventure, if it must end so; and not a mountain of gold can amend it. Yet I am glad that I have shared in your perils – that has been more than any Baggins deserves."

– Bilbo Baggins, *The Hobbit*

The journey to fault-tolerant quantum computing is a long and treacherous one, filled with unwelcome surprises. Some of the greatest challenges – including improving gate performance, reducing spurious effects during readout, and decreasing leakage – will inevitably require increasingly accurate simulations to help obtain the next order of magnitude in the fidelity of operations. In this thesis, I have presented concepts and methodologies which I hope will contribute to the field of circuit quantum electrodynamics and lead to a better understanding of the limiting factors in qubit control and readout.

In Chapter 3, I introduced a tailored solver for systems with a fast-oscillating drive, enabling a significant speedup over traditional numerical integrators. Furthermore, I extended this solver to master equations in Chapter 4. Such a solver could be used for much larger system sizes than I demonstrated in this thesis – for example, the simulation of a surface code – without the necessity of rotating-wave approximations, which could lead to a better understanding of the leakage processes and cross-talk. Further, the solver could be applied to simulate dynamics of different quantum architectures such as spin qubits, so long as a matrix formulation of the system exists. Many questions still remain, such as whether there exists a more effective approach to approximate the Dyson operators, how to optimally select step sizes, and how to best incorporate additional corrections to the Volterra integrals. In my future work, I hope to answer some of these problems and push forward toward highly efficient integrators that can tackle the most challenging quantum simulations.

Qubit readout remains one of the greatest challenges in circuit QED – in Chapter 5, I detailed how the readout process can be corrupted due to resonances between the low and

high energy states, and provided more insight into the role of the external gate charge. As the location of these resonances becomes better understood, this could eventually lead to more tailored choices of the system parameters to help minimize the non-QND nature of the readout process. Moreover, the methods developed in this thesis can be extended to other types of qubits, which would again help understand the limiting factors for qubit readout.

While spurious resonances remain problematic, there are still other ways we can improve upon readout – namely, optimization of the Purcell filter parameters, which I addressed in Chapter 6. Whilst the properties of this filter are well understood in the broadband and strong-coupling regimes, there are still many unanswered questions for the intermediate and weak coupling cases which only I touched on – for example, how do we understand spurious resonances of the transmon-resonator-Purcell filter system in the context of the resonator branches and ionization? Is there a more efficient method of simulating the dynamics without resorting to semiclassical approaches? How does having different resonator photon populations conditioned on the qubit state impact the readout?

The field of circuit quantum electrodynamics is undeniably rich, fascinating, and of great promise in the years to come. Whilst there are a multitude of challenges remaining on the road to useful quantum computation, I am thrilled to be part of and contribute to this area of research, and see their resolutions in the years to come.

appendix A

SNR Analysis

"You'll have some tea... are you sure you don't want any? Aw go on, you'll have some. Go on go on go on go on go on go on go on go on GO ON!"

– Mrs Doyle, Father Ted

In this appendix, I provide some additional context to the SNR and the performance of readout in different parameter regimes.

In order to understand how to optimize the SNR function, we must first consider the steady state separation between the filter pointer states:

$$\begin{aligned} |\beta_{ss}^g - \beta_{ss}^e|^2 &= J^{-2} n_{\text{crit}} \left| \Delta^g - \Delta^e \frac{(\Delta^g(\Delta_p - i\kappa/2) - J^2)}{(\Delta^e(\Delta_p - i\kappa/2) - J^2)} \right|^2 \\ &= 4\chi^2 J^2 n_{\text{crit}} \left| \frac{1}{(\Delta^e(\Delta_p - i\kappa/2) - J^2)} \right|^2 . \end{aligned} \tag{A.1}$$

If we assume that the system is always in the steady state, then the SNR scales as

$$\text{SNR}(t) = 2\kappa t |\beta_{ss}^g - \beta_{ss}^e|^2. \tag{A.2}$$

Thus, the SNR rate at steady state is simply given by the time derivative,

$$\text{SNR}' = \frac{d\text{SNR}}{dt} = 2\kappa |\beta_{ss}^g - \beta_{ss}^e|^2. \tag{A.3}$$

Proceeding we must make assumptions about the drive frequency. Since we wish to obtain analytic expressions, the optimization procedures in Sect. 6.3.3 are not particularly useful.

Instead, we estimate that the optimal drive frequency is the average frequency of one of the two normal modes, translating to:

$$\omega_d \approx \frac{1}{2}\,\mathrm{Re}\,[\lambda^g_{l,h} + \lambda^e_{l,h}], \tag{A.4}$$

where the eigenvalues are given by

$$\lambda^{g/e}_{l,h} = \frac{\omega_r^{g/e} + \omega_p - i\kappa_p/2}{2} \pm \frac{1}{2}\sqrt{\left(\Delta^{g/e}_{rp} + \frac{i\kappa_p}{2}\right)^2 + 4J^{2,g/e}}, \tag{A.5}$$

where $\Delta^{g/e}_{rp} = \omega_r^{g/e} - \omega_p$.

A.1 Strong Coupling Regime – $4J \gg \kappa_p$

In the strong coupling regime, we expect two distinct low and high normal mode responses, separated in frequency space on the order of $2J$. We begin by assuming that the renormalization to the real part of the eigenvalues $\lambda^{g/e}_{l,h}$ from the imaginary κ_p term present in the square root is negligible, yielding the optimal drive frequency for each normal mode,

$$\omega_{d,l,h} = \frac{\omega_r^g + \chi + \omega_p}{2} \pm \frac{1}{4}\left(\sqrt{\Delta^{2,g}_{rp} + 4J^{2,g}} + \sqrt{\Delta^{2,e}_{rp} + 4J^{2,e}}\right). \tag{A.6}$$

We then assume $4J \gg \Delta^{g,e}_{rp}$ and $J^{g,e} \approx J$, yielding the expansion

$$\omega_{d,l,h} \approx \frac{\omega_r^g + \chi + \omega_p}{2} \pm (J + C), \tag{A.7}$$

where

$$C = \frac{(\Delta^{g2}_{rp} + 2\chi\Delta^g_{rp} + 2\chi^2)}{8J}. \tag{A.8}$$

Let's assume we're targeting the lower mode, choosing the $-$ term. Thus, we find

$$\begin{aligned} \Delta^e &= \frac{\Delta^g_{rp} + 3\chi}{2} - (J + C), \\ \Delta_p &= \frac{-\Delta^g_{rp} - \chi}{2} - (J + C). \end{aligned} \tag{A.9}$$

Then, differentiating the SNR function at steady state by κ and setting this result to zero, we find

$$\frac{d\mathrm{SNR'}}{d\kappa} = 0 \longrightarrow \kappa_{\mathrm{opt}} = -2\left(\frac{J^2}{\Delta^e} - \Delta_p\right). \tag{A.10}$$

Then, assuming $(\Delta_{rp} + 3\chi)/2 - C$ is small, we find

$$\frac{J^2}{\Delta^e} \approx -J + C - \frac{\Delta^g_{rp} + 3\chi}{2}, \tag{A.11}$$

thus yielding

$$\kappa_{\text{opt}} = -2\left(\chi + 2C\right). \tag{A.12}$$

This is only a correction to the well known optimal decay rate for a transmon-resonator system, $\kappa = -2\chi$ (here we use the notation $\chi < 0$.)

A.2 Broadband Filter Regime – $\kappa_p \gg 4J, \Delta^{g/e}_{rp}$

In the broadband filter regime, we cannot ignore frequency renormalizations from the broadband Purcell filter. Instead, we treat the imaginary part under the square root perturbatively, yielding

$$\sqrt{-\kappa^2/4 + 4J^2 + \Delta^{2,g/e}_{rp} - i\Delta^{g/e}_{rp}\kappa} \approx \frac{-i}{2}\kappa' + \frac{\kappa\Delta^{g/e}_{rp}}{\kappa'}, \tag{A.13}$$

where

$$\kappa' = \sqrt{\kappa^2 - 16J^2 + 4\Delta^{2,g/e}_{rp}} \approx \kappa - \frac{8J^2 - 2\Delta^{2,g/2}_{r,p}}{\kappa}. \tag{A.14}$$

.Thus, we find the approximate pairs of eigenvalues for the 'low' and 'high' modes are

$$\begin{aligned}
\lambda^{g/e}_l =& \frac{\omega^{g/e}_r + \omega_p}{2} - \frac{\kappa(\omega^{g/e}_r - \omega_p)}{2\kappa'} - \frac{i}{4}(\kappa + \kappa'), \\
\approx& \omega_p - \frac{i\kappa}{2}, \\
\lambda^{g/e}_h =& \frac{\omega^{g/e}_r + \omega_p}{2} + \frac{\kappa(\omega^{g/e}_r - \omega_p)}{2\kappa'} - \frac{i}{4}(\kappa - \kappa'), \\
\approx& \omega^{g/e}_r - i\frac{4J^2 - \Delta^{g/e}_{rp}}{2\kappa}.
\end{aligned} \tag{A.15}$$

Here, one set of normal mode responses $\lambda^{g/e}_l$ has virtually no dependence on the qubit state – thus, we must target the other mode, with optimal drive frequency $\omega_d = \omega_r$. This simplifies the formula for the optimization of the SNR rate,

$$\text{SNR}' = 8\chi^2 J^2 \kappa n_{\text{crit}} \left| \frac{1}{\chi(\omega_p - \omega_r - i\kappa/2) - J^2} \right|^2. \tag{A.16}$$

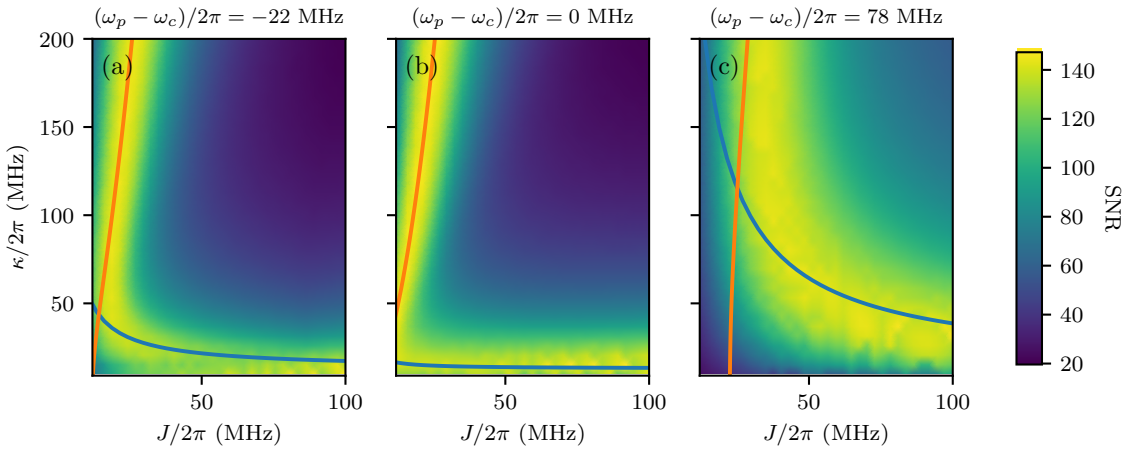

Figure A.1 Plot of the SNR functions at after 400ns as in Chapter 6, including the optimal parameters κ_{opt}, J_{opt} defined in the limits $4J \gg \kappa$ (blue lines) and $\kappa \gg 4J$ (orange lines), respectively.

Now taking the derivative with respect to J and setting the result to 0, we obtain an equation

$$J_{\text{opt}}^4 = \chi^2(\omega_p - \omega_r)^2 + \chi^2\kappa^2/4. \tag{A.17}$$

To see how well these approximations match the predictions from the optimization in Chapter 6, I replot the figures of the SNR at steady state, this time including the curves indicating the optimal parameters κ_{opt} (blue lines) and J_{opt} (orange lines) in the two distinct parameter regimes, $4J \gg \kappa$ and $\kappa \gg 4J$, respectively. These well capture the behaviour at small resonator-Purcell detunings, but show some more significant deviations for large detunings, see Fig. A.1(c). Using this methodology may help with the optimization of the targeted parameters for readout during the fabrication and design process.

Bibliography

[1] Ross Shillito, Jonathan A. Gross, Agustin Di Paolo, Élie Genois, and Alexandre Blais. Fast and differentiable simulation of driven quantum systems. *Phys. Rev. Res.*, 3:033266, Sep 2021. doi: 10.1103/PhysRevResearch.3.033266. URL https://link.aps.org/doi/10.1103/PhysRevResearch.3.033266. (Cited on pages ix, 24, 25, 26, 29, and 34.)

[2] Ross Shillito, Alexandru Petrescu, Joachim Cohen, Jackson Beall, Markus Hauru, Martin Ganahl, Adam G.M. Lewis, Guifre Vidal, and Alexandre Blais. Dynamics of transmon ionization. *Phys. Rev. Appl.*, 18:034031, Sep 2022. doi: 10.1103/PhysRevApplied.18.034031. URL https://link.aps.org/doi/10.1103/PhysRevApplied.18.034031. (Cited on pages ix, 5, 24, 62, 63, 65, 69, 82, 94, 95, and 97.)

[3] François Swiadek, Ross Shillito, Paul Magnard, Ants Remm, Christoph Hellings, Nathan Lacroix, Quentin Ficheux, Dante Colao Zanuz, Graham J. Norris, Alexandre Blais, Sebastian Krinner, and Andreas Wallraff. Enhancing dispersive readout of superconducting qubits through dynamic control of the dispersive shift: Experiment and theory, 2023. (Cited on pages x, 19, and 98.)

[4] Richard P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21:467–488, 1981. (Cited on page 1.)

[5] John Preskill. Quantum computing 40 years later, 2021. (Cited on page 1.)

[6] 40 years of quantum computing. *Nature Reviews Physics*, 4(1):1–1, Jan 2022. ISSN 2522-5820. doi: 10.1038/s42254-021-00410-6. URL https://doi.org/10.1038/s42254-021-00410-6. (Cited on page 1.)

[7] Maud Vinet. The path to scalable quantum computing with silicon spin qubits. *Nature Nanotechnology*, 16(12):1296–1298, Dec 2021. ISSN 1748-3395. doi: 10.

1038/s41565-021-01037-5. URL https://doi.org/10.1038/s41565-021-01037-5. (Cited on page 1.)

[8] W. Huang, C. H. Yang, K. W. Chan, T. Tanttu, B. Hensen, R. C. C. Leon, M. A. Fogarty, J. C. C. Hwang, F. E. Hudson, K. M. Itoh, A. Morello, A. Laucht, and A. S. Dzurak. Fidelity benchmarks for two-qubit gates in silicon. *Nature*, 569 (7757):532–536, May 2019. ISSN 1476-4687. doi: 10.1038/s41586-019-1197-0. URL https://doi.org/10.1038/s41586-019-1197-0. (Cited on page 1.)

[9] Peter Stano and Daniel Loss. Review of performance metrics of spin qubits in gated semiconducting nanostructures. *Nature Reviews Physics*, 4(10):672–688, Oct 2022. ISSN 2522-5820. doi: 10.1038/s42254-022-00484-w. URL https://doi.org/10.1038/s42254-022-00484-w. (Cited on page 1.)

[10] Xiao Xue, Maximilian Russ, Nodar Samkharadze, Brennan Undseth, Amir Sammak, Giordano Scappucci, and Lieven M. K. Vandersypen. Quantum logic with spin qubits crossing the surface code threshold. *Nature*, 601(7893):343–347, Jan 2022. ISSN 1476-4687. doi: 10.1038/s41586-021-04273-w. URL https://doi.org/10.1038/s41586-021-04273-w. (Cited on page 1.)

[11] J. I. Cirac and P. Zoller. Quantum computations with cold trapped ions. *Phys. Rev. Lett.*, 74:4091–4094, May 1995. doi: 10.1103/PhysRevLett.74.4091. URL https://link.aps.org/doi/10.1103/PhysRevLett.74.4091. (Cited on page 1.)

[12] H. Häffner, C.F. Roos, and R. Blatt. Quantum computing with trapped ions. *Physics Reports*, 469(4):155–203, 2008. ISSN 0370-1573. doi: https://doi.org/10.1016/j.physrep.2008.09.003. URL https://www.sciencedirect.com/science/article/pii/S0370157308003463. (Cited on page 1.)

[13] Rainer Blatt and David Wineland. Entangled states of trapped atomic ions. *Nature*, 453(7198):1008–1015, Jun 2008. ISSN 1476-4687. doi: 10.1038/nature07125. URL https://doi.org/10.1038/nature07125. (Cited on page 1.)

[14] J M Amini, H Uys, J H Wesenberg, S Seidelin, J Britton, J J Bollinger, D Leibfried, C Ospelkaus, A P VanDevender, and D J Wineland. Toward scalable ion traps for quantum information processing. *New Journal of Physics*, 12(3):033031, mar 2010. doi: 10.1088/1367-2630/12/3/033031. URL https://dx.doi.org/10.1088/1367-2630/12/3/033031. (Cited on page 1.)

[15] Alexandre Blais, Ren-Shou Huang, Andreas Wallraff, S. M. Girvin, and R. J. Schoelkopf. Cavity quantum electrodynamics for superconducting electrical circuits: An architecture for quantum computation. *Phys. Rev. A*, 69:062320, Jun

2004. doi: 10.1103/PhysRevA.69.062320. URL https://link.aps.org/doi/10.1103/PhysRevA.69.062320. (Cited on pages 1, 2, and 18.)

[16] Alexandre Blais, Arne L. Grimsmo, S. M. Girvin, and Andreas Wallraff. Circuit quantum electrodynamics. *Rev. Mod. Phys.*, 93:025005, May 2021. doi: 10.1103/RevModPhys.93.025005. URL https://link.aps.org/doi/10.1103/RevModPhys.93.025005. (Cited on pages 1, 5, 13, 16, and 17.)

[17] J. Q. You and Franco Nori. Superconducting Circuits and Quantum Information. *Physics Today*, 58(11):42–47, 11 2005. ISSN 0031-9228. doi: 10.1063/1.2155757. URL https://doi.org/10.1063/1.2155757. (Cited on page 1.)

[18] J. Q. You and Franco Nori. Atomic physics and quantum optics using superconducting circuits. *Nature*, 474(7353):589–597, Jun 2011. ISSN 1476-4687. doi: 10.1038/nature10122. URL https://doi.org/10.1038/nature10122. (Cited on page 1.)

[19] Frank Arute et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, Oct 2019. ISSN 1476-4687. doi: 10.1038/s41586-019-1666-5. URL https://doi.org/10.1038/s41586-019-1666-5. (Cited on pages 1 and 87.)

[20] Sebastian Krinner, Nathan Lacroix, Ants Remm, Agustin Di Paolo, Elie Genois, Catherine Leroux, Christoph Hellings, Stefania Lazar, Francois Swiadek, Johannes Herrmann, Graham J. Norris, Christian Kraglund Andersen, Markus Müller, Alexandre Blais, Christopher Eichler, and Andreas Wallraff. Realizing repeated quantum error correction in a distance-three surface code. *Nature*, 605(7911):669–674, May 2022. ISSN 1476-4687. doi: 10.1038/s41586-022-04566-8. URL https://doi.org/10.1038/s41586-022-04566-8. (Cited on pages 1, 58, 62, 63, and 87.)

[21] Rajeev Acharya et al. Suppressing quantum errors by scaling a surface code logical qubit. *Nature*, 614(7949):676–681, Feb 2023. ISSN 1476-4687. doi: 10.1038/s41586-022-05434-1. URL https://doi.org/10.1038/s41586-022-05434-1. (Cited on pages 1 and 62.)

[22] B.D. Josephson. Possible new effects in superconductive tunnelling. *Physics Letters*, 1(7):251–253, 1962. ISSN 0031-9163. doi: https://doi.org/10.1016/0031-9163(62)91369-0. URL https://www.sciencedirect.com/science/article/pii/0031916362913690. (Cited on pages 1 and 4.)

[23] Michel H. Devoret, John M. Martinis, and John Clarke. Measurements of macroscopic quantum tunneling out of the zero-voltage state of a current-biased Josephson junction.

Phys. Rev. Lett., 55:1908–1911, Oct 1985. doi: 10.1103/PhysRevLett.55.1908. URL https://link.aps.org/doi/10.1103/PhysRevLett.55.1908. (Cited on page 1.)

[24] A. J. Leggett. Essays in theoretical physics. in honour of dirk ter haar. In W. E. Parry, editor, *Essays in theoretical physics. In honour of Dirk ter Haar.*, pages 95–127. Pergamon, Oxford, 1984. (Cited on page 1.)

[25] Morten Kjaergaard, Mollie E. Schwartz, Ami Greene, Gabriel O. Samach, Andreas Bengtsson, Michael O'Keeffe, Christopher M. McNally, Jochen Braumüller, David K. Kim, Philip Krantz, Milad Marvian, Alexander Melville, Bethany M. Niedzielski, Youngkyu Sung, Roni Winik, Jonilyn Yoder, Danna Rosenberg, Kevin Obenland, Seth Lloyd, Terry P. Orlando, Iman Marvian, Simon Gustavsson, and William D. Oliver. Programming a quantum computer with quantum instructions, 2020. (Cited on page 1.)

[26] Austin G. Fowler, Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland. Surface codes: Towards practical large-scale quantum computation. *Phys. Rev. A*, 86: 032324, Sep 2012. doi: 10.1103/PhysRevA.86.032324. URL https://link.aps.org/doi/10.1103/PhysRevA.86.032324. (Cited on page 1.)

[27] Earl T. Campbell, Barbara M. Terhal, and Christophe Vuillot. Roads towards fault-tolerant universal quantum computation. *Nature*, 549(7671):172–179, Sep 2017. ISSN 1476-4687. doi: 10.1038/nature23460. URL https://doi.org/10.1038/nature23460. (Cited on page 2.)

[28] Alexander P. M. Place, Lila V. H. Rodgers, Pranav Mundada, Basil M. Smitham, Mattias Fitzpatrick, Zhaoqi Leng, Anjali Premkumar, Jacob Bryon, Andrei Vrajitoarea, Sara Sussman, Guangming Cheng, Trisha Madhavan, Harshvardhan K. Babla, Xuan Hoang Le, Youqi Gang, Berthold Jäck, András Gyenis, Nan Yao, Robert J. Cava, Nathalie P. de Leon, and Andrew A. Houck. New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds. *Nature Communications*, 12(1):1779, Mar 2021. ISSN 2041-1723. doi: 10.1038/s41467-021-22030-5. URL https://doi.org/10.1038/s41467-021-22030-5. (Cited on pages 2 and 87.)

[29] Alexandre Blais, Jay Gambetta, A. Wallraff, D. I. Schuster, S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf. Quantum-information processing with circuit quantum electrodynamics. *Phys. Rev. A*, 75:032329, Mar 2007. doi: 10.1103/PhysRevA.75.032329. URL https://link.aps.org/doi/10.1103/PhysRevA.75.032329. (Cited on page 2.)

[30] T. Walter, P. Kurpiers, S. Gasparinetti, P. Magnard, A. Potočnik, Y. Salathé, M. Pechal, M. Mondal, M. Oppliger, C. Eichler, and A. Wallraff. Rapid high-fidelity single-shot dispersive readout of superconducting qubits. *Phys. Rev. Appl.*, 7:054020, May 2017. doi: 10.1103/PhysRevApplied.7.054020. URL https://link.aps.org/doi/10.1103/PhysRevApplied.7.054020. (Cited on pages 2, 18, 19, 63, and 92.)

[31] Y. Sunada, S. Kono, J. Ilves, S. Tamate, T. Sugiyama, Y. Tabuchi, and Y. Nakamura. Fast readout and reset of a superconducting qubit coupled to a resonator with an intrinsic Purcell filter. *Phys. Rev. Applied*, 17:044016, Apr 2022. doi: 10.1103/PhysRevApplied.17.044016. URL https://link.aps.org/doi/10.1103/PhysRevApplied.17.044016. (Cited on page 2.)

[32] Maxime Boissonneault, J. M. Gambetta, and Alexandre Blais. Nonlinear dispersive regime of cavity qed: The dressed dephasing model. *Phys. Rev. A*, 77:060305, Jun 2008. doi: 10.1103/PhysRevA.77.060305. URL https://link.aps.org/doi/10.1103/PhysRevA.77.060305. (Cited on page 2.)

[33] Alexandru Petrescu, Moein Malekakhlagh, and Hakan E. Türeci. Lifetime renormalization of driven weakly anharmonic superconducting qubits. ii. the readout problem. *Phys. Rev. B*, 101:134510, Apr 2020. doi: 10.1103/PhysRevB.101.134510. URL https://link.aps.org/doi/10.1103/PhysRevB.101.134510. (Cited on page 2.)

[34] Daniel Sank, Zijun Chen, Mostafa Khezri, J. Kelly, R. Barends, B. Campbell, Y. Chen, B. Chiaro, A. Dunsworth, A. Fowler, E. Jeffrey, E. Lucero, A. Megrant, J. Mutus, M. Neeley, C. Neill, P. J. J. O'Malley, C. Quintana, P. Roushan, A. Vainsencher, T. White, J. Wenner, Alexander N. Korotkov, and John M. Martinis. Measurement-induced state transitions in a superconducting qubit: Beyond the rotating wave approximation. *Phys. Rev. Lett.*, 117:190503, Nov 2016. doi: 10.1103/PhysRevLett.117.190503. URL https://link.aps.org/doi/10.1103/PhysRevLett.117.190503. (Cited on page 2.)

[35] M. D. Reed, B. R. Johnson, A. A. Houck, L. DiCarlo, J. M. Chow, D. I. Schuster, L. Frunzio, and R. J. Schoelkopf. Fast reset and suppressing spontaneous emission of a superconducting qubit. *Applied Physics Letters*, 96(20), 05 2010. ISSN 0003-6951. doi: 10.1063/1.3435463. URL https://doi.org/10.1063/1.3435463. 203110. (Cited on pages 3 and 87.)

[36] Evan Jeffrey, Daniel Sank, J. Y. Mutus, T. C. White, J. Kelly, R. Barends, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. Megrant, P. J. J. O'Malley, C. Neill, P. Roushan, A. Vainsencher, J. Wenner, A. N. Cleland, and John M. Martinis. Fast accurate state

measurement with superconducting qubits. *Phys. Rev. Lett.*, 112:190504, May 2014. doi: 10.1103/PhysRevLett.112.190504. URL https://link.aps.org/doi/10.1103/PhysRevLett.112.190504. (Cited on pages 3 and 87.)

[37] Nicholas T. Bronn, Easwar Magesan, Nicholas A. Masluk, Jerry M. Chow, Jay M. Gambetta, and Matthias Steffen. Reducing spontaneous emission in circuit quantum electrodynamics by a combined readout/filter technique. *IEEE Transactions on Applied Superconductivity*, 25(5):1–10, 2015. doi: 10.1109/TASC.2015.2456109. (Cited on pages 3 and 87.)

[38] 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. *Phys. Rev. A*, 76:042319, Oct 2007. doi: 10.1103/PhysRevA.76.042319. URL https://link.aps.org/doi/10.1103/PhysRevA.76.042319. (Cited on pages 4, 5, 7, and 81.)

[39] Uri Vool and Michel Devoret. Introduction to quantum electromagnetic circuits. *International Journal of Circuit Theory and Applications*, 45(7):897–934, 2017. doi: https://doi.org/10.1002/cta.2359. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/cta.2359. (Cited on page 4.)

[40] V. Bouchiat, D. Vion, P. Joyez, D. Esteve, and M. H. Devoret. Quantum coherence with a single Cooper pair. *Physica Scripta*, T76(1):165–170, 1998. ISSN 0031-8949. doi: 10.1238/physica.topical.076a00165. (Cited on page 4.)

[41] Y. Nakamura, Yu. A. Pashkin, and J. S. Tsai. Coherent control of macroscopic quantum states in a single-Cooper-pair box. *Nature*, 398(6730):786–788, 1999. ISSN 0028-0836. doi: 10.1038/19718. (Cited on page 4.)

[42] Felix Bloch. Über die Quantenmechanik der Elektronen in Kristallgittern. *Zeitschrift für Physik*, 52(7):555–600, July 1929. ISSN 0044-3328. doi: 10.1007/BF01339455. URL https://doi.org/10.1007/BF01339455. (Cited on page 7.)

[43] A. Cottet. Implementation of a quantum bit in a superconducting circuit, 2002. (Cited on page 7.)

[44] Michel Devoret and John M. Martinis. Quantum information processing with superconducting circuits. In Domenico J.W. Giulini, Erich Joos, Claus Kiefer, Joachim Kupsch, Ion-Olimpiu Stamatescu, and H. Dieter Zeh, editors, *Quantum Entanglement and Information Processing (Les Houches Session LXXIX)*, volume 79 of *Les Houches - École d'Été de Physique Théorique*, pages 443–485. Oxford University Press, 2004. (Cited on page 7.)

[45] Bengt Fornberg. Generation of finite difference formulas on arbitrarily spaced grids. *Mathematics of Computation*, 51:699–706, 1988. (Cited on page 8.)

[46] Vladimir E. Manucharyan, Jens Koch, Leonid I. Glazman, and Michel H. Devoret. Fluxonium: Single Cooper-pair circuit free of charge offsets. *Science*, 326(5949): 113–116, 2009. doi: 10.1126/science.1175552. URL https://www.science.org/doi/abs/10.1126/science.1175552. (Cited on page 9.)

[47] Susanne Richer and David DiVincenzo. Circuit design implementing longitudinal coupling: A scalable scheme for superconducting qubits. *Physical Review B*, 93(13): 134501, Apr 2016. doi: 10.1103/PhysRevB.93.134501. (Cited on page 9.)

[48] Charles R. Harris et al. Array programming with NumPy. *Nature*, 585(7825):357–362, September 2020. doi: 10.1038/s41586-020-2649-2. URL https://doi.org/10.1038/s41586-020-2649-2. (Cited on page 11.)

[49] P. Krantz, M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, and W. D. Oliver. A quantum engineer's guide to superconducting qubits. *Applied Physics Reviews*, 6 (2), 06 2019. ISSN 1931-9401. doi: 10.1063/1.5089550. URL https://doi.org/10.1063/1.5089550. 021318. (Cited on pages 13 and 17.)

[50] E.T. Jaynes and F.W. Cummings. Comparison of quantum and semiclassical radiation theories with application to the beam maser. *Proceedings of the IEEE*, 51(1):89–109, 1963. doi: 10.1109/PROC.1963.1664. (Cited on pages 13 and 64.)

[51] J. R. Luttinger and P. A. Kohn. Motion of electrons and holes in perturbed periodic fields. *Physical Review*, 97(4):869–883, 1955. doi: 10.1103/PhysRev.97.869. (Cited on page 13.)

[52] J. R. Schrieffer and P. A. Wolff. Relation between the Anderson and Kondo hamiltonians. *Physical Review*, 149(2):491–492, 1966. doi: 10.1103/PhysRev.149.491. (Cited on page 13.)

[53] D. Walls and G. Milburn. *Quantum Optics*. Springer, Berlin, 1994. (Cited on page 14.)

[54] H. Carmichael. *An Open Systems Approach to Quantum Optics: Lectures Presented at the Université Libre de Bruxelles, October 28 to November 4, 1991*. Number v. 18 in An Open Systems Approach to Quantum Optics: Lectures Presented at the Université Libre de Bruxelles, October 28 to November 4, 1991. Springer Berlin Heidelberg, 1993. ISBN 9783540566342. URL https://books.google.ca/books?id=El5gxgxWhpgC. (Cited on pages 15 and 51.)

[55] Mickey M. A useful tool for later use, 2021. URL https://tenor.com/view/surprise-tool-mickey-mouse-mickey-clubhouse-gif-18764497. (Cited on pages 15 and 20.)

[56] Edward M. Purcell. Spontaneous emission probabilities at radio frequencies. *Physical Review*, 69(9-10):681, 1946. (Cited on pages 15 and 86.)

[57] D. I. Schuster, A. Wallraff, A. Blais, L. Frunzio, R.-S. Huang, J. Majer, S. M. Girvin, and R. J. Schoelkopf. ac Stark shift and dephasing of a superconducting qubit strongly coupled to a cavity field. *Phys. Rev. Lett.*, 94:123602, Mar 2005. doi: 10.1103/PhysRevLett.94.123602. URL https://link.aps.org/doi/10.1103/PhysRevLett.94.123602. (Cited on page 17.)

[58] C. C. Bultink, B. Tarasinski, N. Haandbæk, S. Poletto, N. Haider, D. J. Michalak, A. Bruno, and L. DiCarlo. General method for extracting the quantum efficiency of dispersive qubit readout in circuit QED. *Applied Physics Letters*, 112(9), 03 2018. ISSN 0003-6951. doi: 10.1063/1.5015954. URL https://doi.org/10.1063/1.5015954. 092601. (Cited on pages 17 and 18.)

[59] A. A. Clerk, M. H. Devoret, S. M. Girvin, F. Marquardt, and R. J. Schoelkopf. Introduction to quantum noise, measurement, and amplification. *Reviews of Modern Physics*, 82(2):1155–1208, 2010. doi: 10.1103/RevModPhys.82.1155. (Cited on page 17.)

[60] Mostafa Khezri et al. Measurement-induced state transitions in a superconducting qubit: Within the rotating wave approximation, 2022. (Cited on pages 18, 19, and 83.)

[61] Z. K. Minev, S. O. Mundhada, S. Shankar, P. Reinhold, R. Gutiérrez-Jáuregui, R. J. Schoelkopf, M. Mirrahimi, H. J. Carmichael, and M. H. Devoret. To catch and reverse a quantum jump mid-flight. *Nature*, 570(7760):200–204, Jun 2019. doi: 10.1038/s41586-019-1287-z. Epub 2019 Jun 3. (Cited on pages 18 and 92.)

[62] J. E. Johnson, E. M. Hoskinson, C. Macklin, D. H. Slichter, I. Siddiqi, and John Clarke. Dispersive readout of a flux qubit at the single-photon level. *Phys. Rev. B*, 84:220503, Dec 2011. doi: 10.1103/PhysRevB.84.220503. URL https://link.aps.org/doi/10.1103/PhysRevB.84.220503. (Cited on pages 18 and 92.)

[63] Joachim Cohen, Alexandru Petrescu, Ross Shillito, and Alexandre Blais. Reminiscence of classical chaos in driven transmons. *PRX Quantum*, 4:020312, Apr 2023. doi: 10.1103/PRXQuantum.4.020312. URL https://link.aps.org/doi/10.1103/PRXQuantum.4.020312. (Cited on pages 18, 81, 82, and 95.)

[64] John C. Butcher. *Numerical Methods for Ordinary Differential Equations*. John Wiley & Sons, New York, 2003. ISBN 978-0-471-96758-3. (Cited on page 19.)

[65] Carl David Tolmé Runge. Über die numerische auflösung von differentialgleichungen. *Mathematische Annalen*, 46(2):167–178, 1895. doi: 10.1007/BF01446807. (Cited on page 19.)

[66] Wilhelm Kutta. Beitrag zur näherungsweisen integration totaler differentialgleichungen. *Zeitschrift für Mathematik und Physik*, 46:435–453, 1901. (Cited on pages 19 and 60.)

[67] Francis Bashforth and J. C. Adams. *An Attempt to test the Theories of Capillary Action by comparing the theoretical and measured forms of drops of fluid. With an explanation of the method of integration employed in constructing the tables which give the theoretical forms of such drops*. Cambridge, 1883. (Cited on page 19.)

[68] C. F. Curtiss and J. O. Hirschfelder. Integration of stiff equations. *Proceedings of the National Academy of Sciences*, 38(3):235–243, 1952. (Cited on page 19.)

[69] Wilhelm Magnus. On the exponential solution of differential equations for a linear operator. *Communications on Pure and Applied Mathematics*, 7(4):649–673, 1954. doi: https://doi.org/10.1002/cpa.3160070404. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/cpa.3160070404. (Cited on page 23.)

[70] Navin Khaneja, Timo Reiss, Cindie Kehlet, Thomas Schulte-Herbrüggen, and Steffen J. Glaser. Optimal control of coupled spin dynamics: design of nmr pulse sequences by gradient ascent algorithms. *Journal of Magnetic Resonance*, 172(2):296 – 305, 2005. ISSN 1090-7807. doi: https://doi.org/10.1016/j.jmr.2004.11.004. URL http://www.sciencedirect.com/science/article/pii/S1090780704003696. (Cited on page 23.)

[71] Hugo Ribeiro, Alexandre Baksic, and Aashish A. Clerk. Systematic Magnus-based approach for suppressing leakage and nonadiabatic errors in quantum dynamics. *Phys. Rev. X*, 7:011021, Feb 2017. doi: 10.1103/PhysRevX.7.011021. URL https://link.aps.org/doi/10.1103/PhysRevX.7.011021. (Cited on page 24.)

[72] Daniel Puzzuoli, Sophia Fuhui Lin, Moein Malekakhlagh, Emily Pritchett, Benjamin Rosand, and Christopher J. Wood. Algorithms for perturbative analysis and simulation of quantum dynamics. *Journal of Computational Physics*, 489:112262, 2023. ISSN 0021-9991. doi: https://doi.org/10.1016/j.jcp.2023.112262. URL https://www.sciencedirect.com/science/article/pii/S0021999123003571. (Cited on pages 24, 34, and 61.)

[73] F. J. Dyson. The radiation theories of Tomonaga, Schwinger, and Feynman. *Phys. Rev.*, 75:486–502, Feb 1949. doi: 10.1103/PhysRev.75.486. URL https://link.aps.org/doi/10.1103/PhysRev.75.486. (Cited on pages 24 and 51.)

[74] E. T. Whittaker and G. Robinson. *The Calculus of Observations: A Treatise on Numerical Mathematics*. Dover, New York, 1967. (Cited on page 25.)

[75] Carl de Boor. Divided differences. *Surveys in Approximation Theory.*, 1:46–69, 2005. (Cited on page 25.)

[76] Ledyard R. Tucker. Implications of factor analysis of three-way matrices for measurement of change. In C. W. Harris, editor, *Problems in Measuring Change*, pages 122–137. University of Wisconsin Press, Madison, WI, 1963. (Cited on page 32.)

[77] Stephan Rabanser, Oleksandr Shchur, and Stephan Günnemann. Introduction to tensor decompositions and their applications in machine learning, 2017. (Cited on page 32.)

[78] Tamara G. Kolda and Brett W. Bader. Tensor decompositions and applications. *SIAM Review*, 51(3):455–500, 2009. doi: 10.1137/07070111X. URL https://doi.org/10.1137/07070111X. (Cited on page 32.)

[79] Timothy F. Havel. Robust procedures for converting among Lindblad, Kraus and matrix representations of quantum dynamical semigroups. *Journal of Mathematical Physics*, 44(2):534–557, 01 2003. ISSN 0022-2488. doi: 10.1063/1.1518555. URL https://doi.org/10.1063/1.1518555. (Cited on pages 49 and 50.)

[80] Jean Dalibard, Yvan Castin, and Klaus Mølmer. Wave-function approach to dissipative processes in quantum optics. *Phys. Rev. Lett.*, 68:580–583, Feb 1992. doi: 10.1103/PhysRevLett.68.580. URL https://link.aps.org/doi/10.1103/PhysRevLett.68.580. (Cited on page 51.)

[81] Andrei D. Polyanin and Alexander V. Manzhirov. *Handbook of Integral Equations*. Chapman and Hall/CRC, Boca Raton, FL, 2nd edition, 2008. ISBN 978-1584885078. (Cited on page 52.)

[82] Carl Friedrich Gauss. *Methodus Nova Integralium Valores Per Approximationem Inveniendi*, volume 3 of *Cambridge Library Collection - Mathematics*, page 165–196. Cambridge University Press, 2011. doi: 10.1017/CBO9781139058247.008. (Cited on page 53.)

[83] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004. ISBN 978-1-107-39400-1. (Cited on page 54.)

[84] Carl Friedrich Gauss. Methodus nova integralium valores per approximationem inveniendi. *Comment. Soc. Reg. Scient. Gotting. Recent.*, 1814. (Cited on page 55.)

[85] I. Bogaert. Iteration-free computation of Gauss-Legendre quadrature nodes and weights. *SIAM J. Sci. Comput.*, 36:C1008–C1026, 2014. (Cited on page 55.)

[86] Nagaraja Kallur and Hanamantappa Rathod. Symmetric gauss legendre quadrature rules for numerical integration over an arbitrary linear tetrahedra in euclidean three-dimensional space. *International Journal of Mathematical Analysis (Ruse)*, 4, 01 2010. (Cited on page 55.)

[87] Jiří Kosinka and Michael Bartoň. Gaussian quadrature for c1 cubic Clough–Tocher macro-triangles. *Journal of Computational and Applied Mathematics*, 351:6–13, 2019. ISSN 0377-0427. doi: https://doi.org/10.1016/j.cam.2018.10.036. URL https://www.sciencedirect.com/science/article/pii/S0377042718306447. (Cited on page 55.)

[88] D.M. Williams, L. Shunn, and A. Jameson. Symmetric quadrature rules for simplexes based on sphere close packed lattice arrangements. *Journal of Computational and Applied Mathematics*, 266:18–38, 2014. ISSN 0377-0427. doi: https://doi.org/10.1016/j.cam.2014.01.007. URL https://www.sciencedirect.com/science/article/pii/S0377042714000211. (Cited on page 55.)

[89] C. T. Reddy and D. J. Shippy. Alternative integration formulae for triangular finite elements. *International Journal for Numerical Methods in Engineering*, 17(1):133–139, 1981. doi: https://doi.org/10.1002/nme.1620170111. (Cited on page 55.)

[90] D. A. Dunavant. High degree efficient symmetrical Gaussian quadrature rules for the triangle. *International Journal for Numerical Methods in Engineering*, 21(6):1129–1148, 1985. doi: https://doi.org/10.1002/nme.1620210612. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/nme.1620210612. (Cited on page 56.)

[91] G. R. Cowper. Gaussian quadrature formulas for triangles. *International Journal for Numerical Methods in Engineering*, 7(3):405–408, 1973. doi: https://doi.org/10.1002/nme.1620070316. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/nme.1620070316. (Cited on page 56.)

[92] Preston C. Hammer and Arthur H. Stroud. Numerical integration over simplexes. *Mathematical Tables and Other Aids to Computation*, 10(55):137–139, 1956. ISSN 08916837. URL http://www.jstor.org/stable/2002484. (Cited on page 56.)

[93] N. Walkington. *Quadrature on Simplices of Arbitrary Dimension*. Research report (Carnegie Mellon University. Department of Mathematical Sciences. Center for Nonlinear Analysis). Carnegie Mellon University, Department of Mathematical Sciences, Center for Nonlinear Analysis, 2000. URL https://books.google.ca/books?id=nS1dGwAACAAJ. (Cited on page 56.)

[94] J.R. Johansson, P.D. Nation, and Franco Nori. Qutip 2: A python framework for the dynamics of open quantum systems. *Computer Physics Communications*, 184(4): 1234–1240, 2013. ISSN 0010-4655. doi: https://doi.org/10.1016/j.cpc.2012.11.019. URL https://www.sciencedirect.L/science/article/pii/S0010465512003955. (Cited on page 60.)

[95] Mostafa Khezri et al. Measurement-induced state transitions in a superconducting qubit: Within the rotating wave approximation, 2022. (Cited on pages 63 and 92.)

[96] K. Serniak, S. Diamond, M. Hays, V. Fatemi, S. Shankar, L. Frunzio, R.J. Schoelkopf, and M.H. Devoret. Direct dispersive monitoring of charge parity in offset-charge-sensitive transmons. *Phys. Rev. Appl.*, 12:014052, Jul 2019. doi: 10.1103/PhysRevApplied.12.014052. URL https://link.aps.org/doi/10.1103/PhysRevApplied.12.014052. (Cited on page 83.)

[97] B. G. Christensen, C. D. Wilen, A. Opremcak, J. Nelson, F. Schlenker, C. H. Zimonick, L. Faoro, L. B. Ioffe, Y. J. Rosen, J. L. DuBois, B. L. T. Plourde, and R. McDermott. Anomalous charge noise in superconducting qubits. *Phys. Rev. B*, 100:140503, Oct 2019. doi: 10.1103/PhysRevB.100.140503. URL https://link.aps.org/doi/10.1103/PhysRevB.100.140503. (Cited on page 83.)

[98] Chenlu Wang, Xuegang Li, Huikai Xu, Zhiyuan Li, Junhua Wang, Zhen Yang, Zhenyu Mi, Xuehui Liang, Tang Su, Chuhong Yang, Guangyue Wang, Wenyan Wang, Yongchao Li, Mo Chen, Chengyao Li, Kehuan Linghu, Jiaxiu Han, Yingshan Zhang, Yulong Feng, Yu Song, Teng Ma, Jingning Zhang, Ruixia Wang, Peng Zhao, Weiyang Liu, Guangming Xue, Yirong Jin, and Haifeng Yu. Towards practical quantum computers: transmon qubit with a lifetime approaching 0.5 milliseconds. *npj Quantum Information*, 8(1):3, Jan 2022. ISSN 2056-6387. doi: 10.1038/s41534-021-00510-2. URL https://doi.org/10.1038/s41534-021-00510-2. 2022/01/13. (Cited on pages 87 and 89.)

[99] Johannes Heinsoo, Christian Kraglund Andersen, Ants Remm, Sebastian Krinner, Theodore Walter, Yves Salathé, Simone Gasparinetti, Jean-Claude Besse, Anton Potočnik, Andreas Wallraff, and Christopher Eichler. Rapid high-fidelity multiplexed readout of superconducting qubits. *Phys. Rev. Appl.*, 10:034040, Sep 2018. doi:

10.1103/PhysRevApplied.10.034040. URL https://link.aps.org/doi/10.1103/PhysRevApplied.10.034040. (Cited on page 87.)

[100] Eric J. Zhang, Srikanth Srinivasan, Neereja Sundaresan, Daniela F. Bogorin, Yves Martin, Jared B. Hertzberg, John Timmerwilke, Emily J. Pritchett, Jeng-Bang Yau, Cindy Wang, William Landers, Eric P. Lewandowski, Adinath Narasgond, Sami Rosenblatt, George A. Keefe, Isaac Lauer, Mary Beth Rothwell, Douglas T. McClure, Oliver E. Dial, Jason S. Orcutt, Markus Brink, and Jerry M. Chow. High-performance superconducting quantum processors via laser annealing of transmon qubits. *Science Advances*, 8(19):eabi6690, 2022. doi: 10.1126/sciadv.abi6690. URL https://www.science.org/doi/abs/10.1126/sciadv.abi6690. (Cited on page 87.)

[101] Eyob A. Sete, John M. Martinis, and Alexander N. Korotkov. Quantum theory of a bandpass Purcell filter for qubit readout. *Phys. Rev. A*, 92:012325, Jul 2015. doi: 10.1103/PhysRevA.92.012325. URL https://link.aps.org/doi/10.1103/PhysRevA.92.012325. (Cited on pages 87 and 89.)

[102] R. Blümel, A. Buchleitner, R. Graham, L. Sirko, U. Smilansky, and H. Walther. Dynamical localization in the microwave interaction of Rydberg atoms: The influence of noise. *Phys. Rev. A*, 44:4521–4540, Oct 1991. doi: 10.1103/PhysRevA.44.4521. URL https://link.aps.org/doi/10.1103/PhysRevA.44.4521. (Cited on page 89.)

[103] Heinz-Peter Breuer and Francesco Petruccione. *The Theory of Open Quantum Systems*. Oxford University Press, 01 2007. ISBN 9780199213900. doi: 10.1093/acprof:oso/9780199213900.001.0001. URL https://doi.org/10.1093/acprof:oso/9780199213900.001.0001. (Cited on page 89.)

[104] S. Vallés-Sanclemente, S. L. M. van der Meer, M. Finkel, N. Muthusubramanian, M. Beekman, H. Ali, J. F. Marques, C. Zachariadis, H. M. Veen, T. Stavenga, N. Haider, and L. DiCarlo. Post-fabrication frequency trimming of coplanar-waveguide resonators in circuit qed quantum processors, 2023. (Cited on page 96.)