Quantum control of a cat-qubit with macroscopic bit-flip times

Ulysse Réglade

PSL
UNIVERSITÉ PARIS

THÈSE DE DOCTORAT
DE L'UNIVERSITÉ PSL

Préparée à l'École Normale Supérieure

Quantum control of a cat-qubit with macroscopic bit-flip times.

Soutenue par
Ulysse REGLADE
Le 17 juin 2024

École doctorale n°564
EDPIF
(École Doctorale Physique en Ile de France)

Spécialité
Physique

Composition du jury :

Julie GROLLIER Directeur de recherche, CNRS	*Présidente*
Denis VION Chargé de recherche, CEA Saclay	*Examinateur*
Antoine BROWAEYS Directeur de recherche, CNRS	*Examinateur*
Jonathan HOME Full professor, ETH Zürich	*Rapporteur*
Gary A. STEELE Full professor, TU Delft University	*Rapporteur*
Zaki LEGHTAS Professeur, Mines Paris-PSL	*Directeur de thèse*
Sébastien JEZOUIN Alice&Bob	*Membre Invité*

ENS
ÉCOLE NORMALE SUPÉRIEURE
PSL

RÉSUMÉ

Les interactions non contrôlées d'un système quantique avec son environnement provoquent de la décohérence, induisant des erreurs sur les qubits, l'unité fondamentale du traitement de l'information quantique. Cet obstacle freine aujourd'hui le développement de l'ordinateur quantique.

Une stratégie prometteuse est la réalisation de qubits naturellement protégés contre certains types d'erreurs. Ces qubits sont appelés "qubits à bruit biaisé". La promesse fondamentale est que: si un qubit est presque parfaitement protégé au niveau matériel contre un type d'erreur (par exemple, les bit-flips), alors les exigences matérielles pour protéger contre l'autre type d'erreur (par exemple, les phase-flips) sont considérablement réduites.
Cependant, le gain anticipé repose sur une condition sévère: les manipulations quantiques du qubit ne doivent pas rompre la protection méticuleusement conçue au niveau matériel. Si cette condition n'est pas respectée, l'avantage des qubits à bruit biaisé est perdu.

Dans ce travail, nous avons implémenté de tels qubits, connus sous le nom de qubits de chat. Dans une première expérience, le qubit est encodé dans une cavité 3D à haut facteur de qualité dotée de dissipation à deux photons. Nous démontrons la préparation de superpositions quantiques et mettons en évidence que l'appareil utilisé pour faire la tomographie d'état réduit la qualité du qubit implémenté. Dans une seconde expérience, nous utilisons un oscillateur supraconducteur 2D, et l'appareil de tomographie est retiré. Nous montrons que notre qubit de chat maintient un temps de bit-flip supérieur à dix secondes, et que nous pouvons contrôler la phase des états quantiques.

Nous démontrons la compatibilité du contrôle quantique et de la protection intrinsèque contre les bit-flips à un niveau sans précédent, prouvant la viabilité des qubits à bruit biaisé pour les technologies quantiques futures.

MOTS CLÉS

Physique quantique, information quantique, qubit de chat, ingénierie de la dissipation, jonctions Josephson

ABSTRACT

The uncontrolled interactions of a quantum system with its environment, known as decoherence, continue to induce too many errors on qubits, thereby blocking their deployment into scalable quantum machines.

A hopeful strategy is to build qubits that are inherently protected against certain types of errors. Such qubits are known as "biased noise qubits". The foundational promise is that if a qubit is almost perfectly protected at the hardware level against one type of error (for example bit-flips), then the hardware requirement to protect against the other type of error (for example phase-flips) is considerably less demanding.

However, the foreseen gain rests on a severe condition: quantum manipulations of the qubit must not break the protection that has been so carefully engineered at the hardware level. If this condition is not met, the advantage of biased noise qubits is lost.

In this work, we have implemented such qubits, known as cat qubits. We present two experiments. In the a first, the qubit is encoded in a high-quality factor 3D cavity equipped with two-photon dissipation. We demonstrate the preparation of quantum superpositions, and highlight that the device used for state tomography reduces the quality of the implemented qubit. In the second experiment, we use a 2D superconducting oscillator, and the tomography apparatus is removed. We show that our cat qubit maintains a bit-flip time of over ten seconds while demonstrating control of the phase of the quantum state.

We demonstrate the compatibility of quantum control and inherent bit-flip protection at an unprecedented level, showing the viability of biased noise qubits for future quantum technologies.

KEYWORDS

Quantum physics, quantum information, cat qubit, dissipation engineering, Josephson junctions

ACKNOWLEDGEMENTS

I would like to start by thanking my thesis jury for agreeing to be members of my defense committee. The rest of the acknowledgments will be written in French, a language closer to my heart.

Je remercie Zaki Leghtas pour les nombreuses et profondes discussions de physique que nous avons eues, ainsi que pour sa patience et sa capacité à canaliser les ardeurs parfois impulsives de son étudiant qui, souvent, ne savait pas expliquer ses raisonnements intuitifs. L'écriture reste encore aujourd'hui un exercice difficile pour moi, et les nombreuses heures passées ensemble à côté du clavier m'ont beaucoup aidé.

Je remercie Sébastien Jézouin, Raphaël Lescanne et Anil Murani pour leur acharnement et leur dévouement, des qualités indispensables à la pratique de la physique expérimentale de pointe, et pour m'avoir montré et expliqué en profondeur en quoi consistaient nos expériences. Je les remercie également pour leur soutien moral.

Je remercie Philippe Campagne-Ibarcq, Mazyar Mirrahimi et Jérémie Guillaud pour avoir pris le temps de répondre à mes questions, de m'expliquer des concepts physiques qui me passionnent, même si parfois ils me dépassaient un peu.

Je remercie Théau Perronin et Raphaël Lescanne de m'avoir permis de faire partie de l'aventure humaine qu'ils ont entreprise. Quoi qu'il advienne, je suis honoré d'avoir été un acteur important d'Alice&Bob dès ses débuts. Plus largement, je souhaite étendre mes remerciements à toute l'équipe qui s'est constituée durant ces trois années, pour les rencontres avec les gens formidables qui la composent.

Je remercie Felix Rautschke, qui a passé un temps infini à simuler et améliorer le design micro-ondes des expériences au centre de ce manuscrit, ainsi que Natalia Pankratova et Emanuele Albertinale pour les journées passées en salle blanche à fabriquer les centaines d'échantillons qui ont été nécessaires à la réalisation de cette thèse.

Je remercie Adrien Bocquet pour la confiance qu'il m'a accordée et pour la détermination dont il a fait preuve tout au long de notre collaboration. J'ai énormément apprécié les moments que nous avons passés ensemble, que

ce soit à travailler sur les expériences ou à discuter de physique en allant acheter des gâteaux pour reprendre des forces. Monter une expérience aussi ambitieuse dès le début de ta thèse témoigne de ton talent, continue comme ça !

Je remercie Benjamin Huard et Antoine Marquet pour les échanges riches et profonds que nous avons eus sur la physique des qubits de chat de Schrödinger.

Je remercie Takis Kontos pour toutes les discussions de physique que nous avons eues autour de sa machine à café dans son laboratoire. Et Dieu sait qu'il me fallait du café !

Je remercie Diego Ruiz, Matthieu Hennekinne et Michaël de Marliave de m'avoir donné l'opportunité de participer à un événement de vulgarisation scientifique sur le plateau de Underscore.

Je remercie Samuel Deléglise et Emmanuel Flurin pour leur attention et leur soutien durant ces trois années.

Je remercie Marius Villiers, Alvise Borgognoni, Aron Vanselow, Erwan Roverc'h, Matthieu Praquin, Brieuc Beauseigneur, Anthony Giraudo, Armelle Celarier, Louis Lattier, ainsi que tous les membres du laboratoire que je n'aurais pas cités pour l'ambiance formidable qui règne dans le groupe.

Je remercie Matthieu Cornet, Maxence Leroy, Flora Ferrari, Audrey Fourniès, Elsa Vizier, Jérémie Osele, Antonin Diaz, Fares Mahjoub, Gabriel Ozouf, Gaël Macherel, Pierrick Pascal, Rayane Beddek, Thibaut Witt, Valentin Le Guen et Vassily Derouette pour leur soutien et leur amitié indéfectible. Avec une mention spéciale pour Florian Fesquet pour le temps passé à jouer aux jeux vidéo pendant d'interminables soirées, tout en parlant d'informatique quantique.

Je tiens à remercier Camille Berdou d'avoir pris le temps de discuter de physique quantique avec un étudiant désabusé en première année à l'École des Mines. Ton enthousiasme et ton énergie m'ont profondément marqué. Merci de m'avoir accueilli dans ta colocation rue Vandrezanne, où, grâce à toi, Alan Sérafin, Léa Sanchez, Julie Chalony et Lorène Imperor j'ai pu vivre à Paris le temps de ma thèse entouré de gens bienveillants et chaleureux.

Enfin, je remercie Papa, Maman, Alizé et Naïs de la tribu Réglade pour leur amour inconditionnel. En écrivant ces lignes, j'ai une pensée pour Anne-Marie, ma mamie, qui est décédée quelques jours avant que je n'envoie ce manuscrit.

RÉSUMÉ DÉTAILLÉ EN FRANÇAIS

Cette thèse est le fruit d'une collaboration entre l'École Normale Supérieure et la startup Alice&Bob dont l'objectif est de concevoir un ordinateur quantique universel et sans erreur en vue de le commercialiser. Une telle machine, si un jour elle est créée, pourrait avoir des applications dans de nombreux domaines tels que la chimie quantique ou la cryptographie, pour n'en citer que deux. Toutefois, alors que la théorie a été élaborée dès les années 90, l'absence à ce jour d'un ordinateur quantique universel témoigne de l'extrême complexité inhérente à leur conception. Plusieurs approches semblent prometteuses et de multiples acteurs sont en compétition. Ce travail se concentre sur la technologie choisie par Alice&Bob pour tenter de relever ce défi.

Commençons par donner une définition : Un système dynamique peut être décrit par l'interaction de forces extérieures, de non-linéarités et d'effets dissipatifs [Guckenheimer and Holmes 1983]. L'étude de ces systèmes est une branche extrêmement riche des mathématiques. Parmi eux, les systèmes dynamiques bistables, tels que par exemple le renversement du champ magnétique terrestre, sont particulièrement intéressants. À une toute autre échelle, il a été démontré qu'en opérant des oscillateurs non linéaires à très petit nombre de photons, il était possible de produire des systèmes bistables manifestant de très grands temps de vie. De tels systèmes sont des candidats idéaux pour le traitement logique classique à ultra-faible puissance [Mabuchi 2011].

Dans le monde de l'information classique, l'unité de base est le bit, à savoir une entité qui peut se trouver dans deux états, que l'on nommera 0 et 1. Un très grand nombre de systèmes classiques de calcul numérique on recourt à des systèmes dynamiques bistables pour stocker l'information. On peut par exemple citer les mémoires flip-flops, obtenues en agençant entre elles plusieurs portes logiques formées de transistors. De tels systèmes sont extrêmement résilients aux erreurs et ont un fort potentiel de mise à l'échelle. L'explosion de la puissance de calcul des ordinateurs modernes

est une conséquence de ce fait.

Il peut être tentant de tirer parti de cette stabilité pour coder robustement l'information quantique, où la sensibilité au bruit reste le facteur limitant ["Suppressing quantum errors by scaling a surface code logical qubit" 2023]. L'unité d'information de base d'un ordinateur quantique est le qubit, à savoir un système qui peut être mis dans n'importe quelle superposition cohérente de deux états logiques $|0\rangle$ et $|1\rangle$. Les qubits sont soumis à deux types d'erreurs : d'une part les "bit-flips" liés au basculement aléatoire entre les états logiques et d'autre part les "phase-flips" qui résultent de la dérive aléatoire de la phase de superposition quantique [Nielsen and Chuang 2010]. Dans cette thèse, nous travaillons sur une technologie nommée qubit de « chat de Schrödinger ». Ce qubit est codé dans la variété des états métastables d'un système dynamique naturellement bistable. Ainsi, l'idée est qu'un tel qubit serait protégé contre les bit-flips au niveau matériel, héritant de la robustesse du système dynamique parent. Pour de tels qubits, le défi est alors de mesurer et contrôler l'information stockée sans briser sa protection. Si ce défi est relevé, la seule erreur restante, les phase-flips, devrait donc pouvoir être gérée. En intégrant plusieurs qubits de chat dans des architectures de correction d'erreurs, les phase-flips peuvent être corrigés avec un surcoût matériel considérablement réduit [Aliferis and Preskill 2008; Tuckett et al. 2018; Puri et al. 2020; Darmawan et al. 2021; Chamberland et al. 2022b] par rapport à celui qui est nécessaire pour corriger les bit-flips et les phase-flips en même temps [Fowler et al. 2012; "Suppressing quantum errors by scaling a surface code logical qubit" 2023], ce qui est le cas pour la plupart des technologies envisagées pour la conception d'un ordinateur quantique.

Cependant, à priori, il semble extrêmement compliqué d'utiliser un système dynamique bistable pour implémenter un qubit. En effet, ces systèmes doivent leur stabilité à des effets dissipatifs qui stabilisent leur position autour d'états métastables. Or la friction provient généralement de l'interaction avec un ensemble de degrés de liberté qui vivent dans un bain thermique. Cela a pour effet de faire fuir de l'information du système, détruisant ainsi les superpositions quantiques qui décohèrent en mélanges statistiques classiques [Zurek 2003]. De façon contre-intuitive, il existe un type de dissipation, connu sous le nom de dissipation à deux photons [Wolinsky and Carmichael 1988; Mirrahimi et al. 2014; Leghtas et al. 2015], qui assure la stabilité sans induire de décohérence. Les échanges de deux photons entre un oscillateur et un environnement froid ont pour effet de stabiliser deux états cohérents, sans pour autant divulguer celui dans lequel se trouve le système. C'est pourquoi de tels systèmes devraient en principe permettre d'atteindre des temps de bit-flip macroscopiques, tout en permettant la pré-

paration et la manipulation de leurs superpositions quantiques [Mirrahimi et al. 2014].

Pour illustrer ce phénomène, j'ai construit un démonstrateur mécanique avec l'aide d'une imprimante 3D. Cet objet purement classique, initialement proposé dans [Lescanne 2020a], est destiné à des fins pédagogiques. Ce modèle se compose d'un pendule suspendu à un ressort, disposé parallèlement à un amortisseur. En excitant le ressort à la moitié de la fréquence propre du pendule, il est possible de le mettre en mouvement. On peut alors démontrer deux effets. Premièrement, pour chaque oscillation du ressort, le pendule se déplace de gauche à droite ou de droite à gauche. Ainsi, en définissant la référence de temps sur le mouvement du ressort, on constate que le pendule admet deux trajectoires asymptotiques qui ont des phases opposées. Dans cette analogie, les deux trajectoires représentent les états logiques 0 et 1. En appliquant manuellement à ce système de légères impulsions, on peut facilement confirmer la grande stabilité de ces deux états. De plus, en enregistrant plusieurs trajectoires, on vérifie que le mouvement de l'amortisseur, qui symbolise le bain thermique, est indépendant de la trajectoire asymptotique du pendule. Ainsi, la dissipation n'entraîne pas la fuite vers l'environnement d'information relative à l'état du système. Théoriquement, en exploitant un système analogue dans des conditions adéquates, il serait donc possible de stabiliser des superpositions quantiques. L'une des grandes fiertés de mon travail de thèse est d'avoir pu utiliser ce démonstrateur mécanique dans le cadre d'un évènement de vulgarisation scientifique sur Twitch devant une très grande audience. Au moment où j'écris ces lignes, la vidéo Youtube associée a accumulé plus de 700 000 vues.

Bien sûr, il serait extrêmement difficile de mettre en œuvre un dispositif quantique utilisant des ressorts, des articulations et des pivots. Au lieu de cela, les résultats expérimentaux que nous présentons dans cette thèse appartiennent à un domaine connu sous le nom d'Electro-Dynamique Quantique en Cavité (CQED). À des températures suffisamment basses, de nombreux métaux conventionnels, comme l'aluminium, deviennent supraconducteurs. L'idée est de fabriquer un circuit électrique avec un tel matériau et de l'exploiter dans ce régime. Dans un supraconducteur, les électrons se déplacent sans dissipation par paires de spins opposés et forment des doublets appelés paires de Cooper. En utilisant une technique connue sous le nom de refroidissement à dilution, la température peut être descendue à quelques dizaines de milli-Kelvins, niveau auquel l'observation d'effets quantiques devient possible.

Ainsi, en pratique, la dissipation à deux photons est mise en œuvre entre deux modes résonants d'un circuit supraconducteur. Le premier, ana-

logue au pendule dans notre exemple, est un mode à haut facteur de qualité que nous nommerons la mémoire. Celle-ci est couplée, via un élément Josephson non linéaire, à un mode tampon dissipatif, que nous appellerons le buffer. Bien que des superpositions quantiques de deux états métastables aient déjà été observées, les temps de bit-flip rapportés saturent dans la gamme de la milliseconde [Lescanne et al. 2020b]. Les implémentations de qubits de chat basées sur l'effet Kerr ont atteint des échelles de temps similaires [Grimm et al. 2020; Frattini et al. 2022].

Dans de tels systèmes, une approche tout à fait naturelle pour tenter d'augmenter le temps de bit-flip est de maximiser le temps de vie du mode de mémoire. Ainsi, le premier projet expérimental présenté dans cette thèse consiste à coupler une cavité supraconductrice 3D à haut facteur de qualité avec un circuit destiné à mettre en œuvre la dissipation à deux photons. Ce projet, que nous nommons le « chat 3D », était le premier lancé par la startup Alice&Bob. Sur le papier, il semblait très prometteur, mais en pratique, la conception de cette expérience s'est avérée extrêmement complexe d'un point de vue technique.

En effet, coupler une cavité 3D à un circuit intégrant un mode fortement dissipatif sans compromettre son facteur de qualité représente un défi significatif en termes d'ingénierie mécanique et de micro-ondes. Pour tenter d'atteindre cet objectif, nous avons développé un porte-échantillon spécifiquement conçu pour cette expérience, accompagné d'un système de fixation mécanique élaboré pour l'occasion. Par ailleurs, un effort considérable a été consacré au développement de filtres coupe-bande efficaces, essentiels à la réalisation du projet. En effet, l'élément Josephson que nous utilisons pour implémenter la dissipation à deux photons possède deux boucles de flux et chacune requiert un biais de flux continu. Réussir à acheminer du courant continu pour chacune des deux boucles à proximité de la cavité tout en l'opérant dans de bonnes conditions représentait un challenge technique, que nous avons partiellement réussi à surmonter. Pendant près de deux ans, nous avons itéré sur la conception de cette expérience, ce qui nous a permis d'obtenir plusieurs résultats intermédiaires.

Tout d'abord, nous avons réussi à générer un état de chat de Schrödinger cohérent uniquement grâce à la dissipation à deux photons, démontrant ainsi que la dissipation, lorsqu'elle est bien maîtrisée, peut servir de ressource de cohérence quantique. Pour effectuer cette mesure, le système était couplé à un dispositif de tomographie auxiliaire constitué d'un qubit transmon. Dans une expérience récente [Berdou et al. 2023], ce dispositif de tomographie a été complètement retiré et des temps de bit-flip supérieurs à cent secondes ont été observés. Cependant, en l'absence du transmon, il était extrêmement

complexe de démontrer que le système fonctionne dans un régime quantique. Dans cette expérience, il n'est donc pas possible de parler de qubit.

Profitant de l'aspect modulaire du projet chat 3D et s'inspirant de [Berdou et al. 2023], nous avons également retiré le transmon, ce qui nous a permis d'atteindre cette fois-ci des temps de bit-flip de l'ordre de l'heure. Plus tard, nous avons même réussi à montrer in situ l'effet délétère du couplage au transmon sur le temps de bit-flip du qubit de chat obtenu. Le projet chat 3D était proche de produire un qubit de chat de Schrödinger, mais le manque de répétabilité du facteur de qualité de la mémoire, ainsi que celui des couplages, lié à notre méthode de fixation mécanique, ne nous a pas permis de produire un échantillon viable.

En dépouillant les données du projet chat 3D, nous en sommes venus à la conclusion qu'il fallait repartir sur une expérience au design plus simple. Nous avons alors conçu un circuit basé sur une géométrie 2D, où le transmon, incriminé par nos mesures, serait absent. Ainsi, pour réaliser la tomographie de l'état quantique de la mémoire, nous avons mis au point un nouveau protocole de mesure qui exploite exclusivement l'élément Josephson utilisé à la base pour implémenter la dissipation à deux photons. Nous démontrons qu'il peut servir à mapper les observables quantiques de la mémoire sur l'état du buffer, qu'il est ensuite possible de mesurer.

Dans cette dernière expérience, nous mettons en œuvre un qubit chat avec des temps de bit-flip dépassant dix secondes, ce qui représente une amélioration de quatre ordres de grandeur par rapport aux implémentations précédentes de qubits chat, et de six ordres de grandeur par rapport à la durée de vie des photons composant le qubit. Nous observons des temps de phase-flip supérieurs à 490 ns, principalement limités par le facteur de qualité de la mémoire. D'autre part, nous montrons que nous pouvaons contrôler la phase des superpositions cohérentes en effectuant une rotation autour de l'axe Z de la sphère de Bloch de notre qubit de chat par effet Zeno [Touzard et al. 2018], réalisant une rotation de π autour de l'axe Z en 235 ns. Nous vérifions également que cette manipulation ne réduit que marginalement le temps de bit-flip, le maintenant au-dessus de dix secondes. Cette expérience démontre la tomographie et le contrôle d'un qubit chat sans compromettre la protection contre les bit-flips. Cependant, des améliorations supplémentaires en matière de préparation d'états, de fidélités de mesure et de perte de photons uniques seront nécessaires avant de passer à un qubit logique pleinement protégé et efficace en termes de matériel [Guillaud and Mirrahimi 2019; Puri et al. 2020; Darmawan et al. 2021; Chamberland et al. 2022b]. Une approche possible pour augmenter le temps de cohérence des superpositions quantiques pourrait être de revenir à une géométrie 3D, mais

dans un contexte mieux maîtrisé.

Cette thèse est l'occasion de mettre en lumière les efforts considérables investis dans ces projets expérimentaux. Après une introduction contextuelle et une revue théorique approfondie, l'accent est placé sur le déroulement des expériences. La description est faite de manière chronologique afin d'immerger le lecteur dans le processus décisionnel qui a guidé nos avancées. Chaque étape est détaillée, reflétant les choix stratégiques et les ajustements techniques effectués en réponse aux défis rencontrés, illustrant ainsi la complexité de la démarche scientifique.

CONTENTS

LIST OF FIGURES

1
INTRODUCTION

1.1 The quantum computing era

I'm on a train, I don't know where it goes, yet I'm on board. Looking through the window, the landscape playing before my eyes appeases my clouded mind. While the stunning progress of artificial intelligence and the alarming degradation of the Earth's climate poses more existential threats than ever, I try to find peace in the wonders of physics. The future will tell if seeking the creation of a universal quantum computer is a vain enterprise, a threat to society, or a formidable new trip through unchartered physics land. Nevertheless, I humbly think it is an honest attempt to have a new and original discussion with reality. The destination is uncertain, but I'm happy I bought my ticket.

The hypothetical rise of the quantum computing era might have several outcomes. Historically, after a period of relative stagnation, the paper that injected fresh momentum into the field detailed how to use a quantum machine to factorise the product of large prime numbers [Shor 1994], breaching RSA, the most deployed asymmetrical cryptography protocol. Shor's algorithm quickly turned quantum computing into a strategic priority for several developed nations. I believe this is still a driving force behind the substantial investment in quantum research today. Classical quantum proof cryptography exists [Bernstein and Lange 2017] and efforts are made to make it efficient and scalable, though current hardware implementations are still lacking [Chen et al. 2016]. In my view, economic incentive should be made higher, as in modern society cryptography is the cornerstone of privacy.

Therefore, if not done correctly, the deployment of quantum computers might be followed by some dark days. If so, my hope is that they will not last for long. In my biased perspective, a quantum computer is first and foremost a formidable playground that might allow us to deepen our understanding of physics. The applications are still unclear, but here is my humble attempt to motivate this optimism.

It is of common knowledge that mathematics are unreasonably effective

in science, as the most complex phenomena can be summarized in short equations. However writting them is only half of the work, since their resolution can prove extremely challenging. For instance, the solution found by Joseph Fourier for the heat equation implied a new mathematical theory, at the heart of today's signal processing. In a lot of cases, tricks and technics are known, and the resolution can be done with pen and paper.

However, when by hand analysis is not possible, numerical simulation is the only alternative. In this regard, fluid dynamics is a fascinating case. In general, there is no known analytical solution for the 3D Navier-Stokes equation [Fefferman 2000]. The boom of digital computing allowed to simulate complex turbulent flow at an unprecedented level of precision. To take only one example, this development played a pivotal role in the industrialisation of modern aviation.

When it comes to simulating complex quantum systems, a paradigm shift is required. A simple calculation reveals that there are not enough atoms in the observable universe to simulate the interactions of 300 spins using classical devices. In contrast, it is reasonable to expect that a parametrizable quantum device of moderate size would be able to efficiently perform such simulations. For physicists, this is where the excitement lies! Such a machine is considered universal if it can simulate arbitrary interactions, and it is the long-term objective pursued by the company Alice&Bob.

Chemistry and materials science would be among the first fields to benefit from quantum computing. The electronic structure of complex molecules could become determinable [Cao et al. 2019]. Understanding the Hamiltonian of type II superconductors could be greatly enhanced. In biology, the protein folding problem could become tractable [Robert et al. 2021]. These advancements are made possible thanks to the quantum phase estimation algorithm [Kitaev 1995], which is a generalization of Shor's original proposal.

In the progression of this chapter, we transition gradually from classical computing to quantum computing. To make the discussion more engaging and tangible, we provide historical examples along the way. Our objective is to introduce the concept of biased noise qubits, with the cat qubit serving as an embodiment of this concept.

1.2 Classical computing

Let us start by giving a crude definition. A computer is an apparatus that allows to process information in a reliable way. It takes signals as inputs,

performs operations with them, and outputs the result. Among many requirements, a high degree of repeatability is expected, therefore the encoding of information on its physical support must have some robustness.

1.2.1 Analog computing

One preliminary approach involves constructing a system that encodes quantities using continuous physical variables, such as the positions of gears within a mechanical apparatus. Operations can then be executed by carefully configuring the interconnection of distinct moving components, thereby giving rise to an analog calculator. Analog computers, for the most part, are designed for specific tasks; their behavior mirrors that of phenomena under investigation, thus justifying their nomenclature. In practice, the term "analog" is retained whenever referring to a machine that encodes information on continuous variables.

The Antikythera mechanism, dated around 200 BC, stands as one of the very first instances of a computing device in history. This analog machine was probably invented by Greek astronomers and was able to predict with great accuracy the position of the celestial bodies known at the time. On close inspection, the structure of its internal gears reveals a peculiar mechanism.

Ancient Greeks possessed an excellent understanding of celestial dynamics, recognizing that the speed of the Moon around the Earth was not constant Fig.1.1b. Today, we understand this variability arises from the slightly eccentric orbit of the Moon Fig.1.1a. To address this phenomenon, astronomers of the time devised a clever mechanical coupling. A gear rotating at a constant speed was linked to a second gear through a spur sliding in a slot. The two gears were deliberately misaligned, causing the spur to oscillate within the slot as the first gear completed a full rotation Fig.1.1c. Through precise adjustments of the spur's position and misalignment, the Antikythera mechanism could accurately track the motion of the Moon Fig.1.1d.

The Antikythera mechanism is one of several analog devices that have existed throughout history. Despite their success, analog computers are generally limited in their applications. The program they are executing is directly encoded in the structure of the computer itself. A new calculation requires building a new one. Additionally, the continuous representation of signals inherently introduces a certain degree of noise, leading to some degree of unpredictability in their behavior. When reliable execution of complex algorithms is required, this design flaw renders analog computers un-

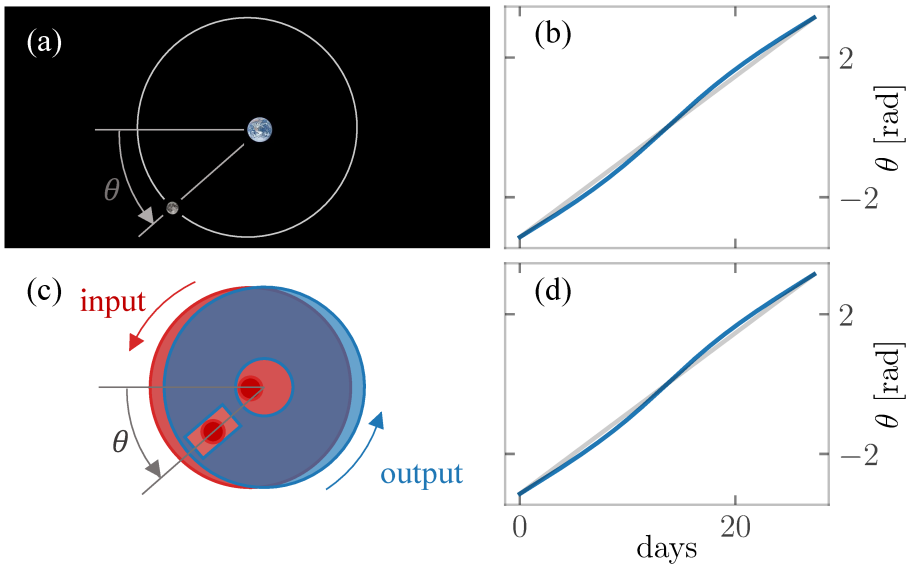

Figure 1.1 Antikythera mechanism. (a) The Moon orbiting the Earth (not to scale). The elliptical trajectory of the Moon is represented in white, and we denote θ as the angle it makes with its apoapsis. (b) The evolution of θ (blue solid line) plotted against time (x-axis). The non-circular orbit of the Moon causes a slight oscillation in its speed around its mean value (gray solid line). (c) Abstract representation of the mechanism that reproduces the motion of the Moon in the Antikythera mechanism. The red input gear engages the motion of the blue gear thanks to a spur sliding in a slot. By tuning the misalignment of the gears, the motion of the Moon can be reproduced with a great degree of accuracy (d).

suitable. Hence, in modern times, their digital counterparts are generally preferred for many applications.

1.2.2 Digital computing

The last decades have been characterized by the boom of digital computing. In the discrete domain, information is stored on a collection of bits—small abstract memories—whose states can be either 0 or 1. Physically, these memories take various forms: the local magnetization of a metal, the voltage in a specifically designed circuit, or the intensity of a radio wave signal at a certain frequency.

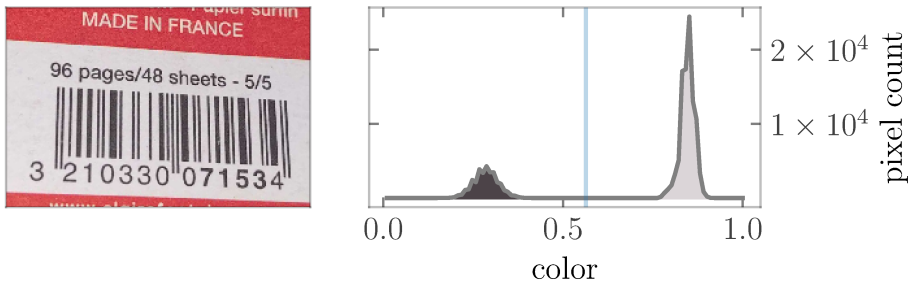

Figure 1.2 Thresholding operation on a barcode. (a) Photograph of a barcode. (b) Color (x-axis) histogram (y-axis) of the barcode. Two distributions akin to Gaussians are visible. They represent the colors of the black and the white bars. Since the standard deviation of these distributions is much smaller than their separation, the barcode can be accurately decoded from the picture.

To extract the value of a bit stored on physical support, a thresholding operation is performed. To illustrate this point, let us say I have the photograph of a barcode Fig. 1.2a, where white encodes for 1 and black for 0. In practice, to read the value of an individual bar, every pixel above a certain shade of gray is considered a 1, while those below are considered 0. Assuming Gaussian noise, it's easy to see that decoding errors are quickly suppressed as the signal-to-noise ratio (SNR) increases, since the overlap of Gaussian distributions is exponentially suppressed with their separation Fig. 1.2b. This feature, among others, makes digital information encoding extremely reliable for transmission and data storage.

To execute an algorithm, however, a computer needs a memory — a collection of bits whose states can be dynamically updated. In many cases, the price to pay for this controllability is a higher error rate. This is why, over time, bits may randomly flip between 0 and 1, events we call bit-flips. For instance, the physical quantity that encodes a bit can slowly drift. Nevertheless, as long as the state of a bit can be set and read reliably, it remains usable for computational tasks.

Magnetic-core memory, though naturally extremely stable, offers an excellent example of a system that continuously rewrites the states of its physical bits during operation. This early form of electronic digital storage was operational from 1950 to 1970 and notably utilized during the Apollo missions. In the Saturn V, the rocket that led humans to the Moon, the memory of the onboard computer responsible for navigation and guidance was com-

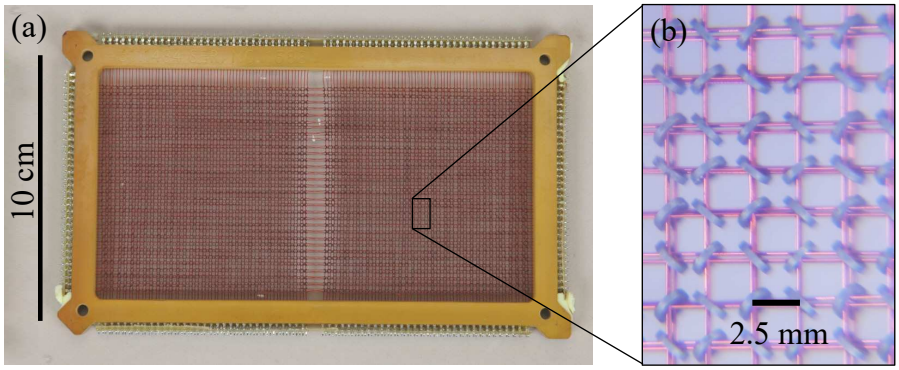

Figure 1.3 Magnetic core memory module. (a) Image taken from [*Ken Shirrif's blog* n.d.]. One core plane for the LVDC's memory, holding 8192 bits. The Saturn V had 8 memory modules, each possessing 14 core planes. (b) Image taken from [Sandlin 2020]. Inset zoom on the core plane. Horizontal and vertical conductive wires run through tiny ferrite cores. Each core represents a single physical bit and was threaded by hand.

prised of small ferrite rings, manually threaded onto a matrix of conductive cables (Fig. 1.3a). This technology was extremely reliable, which played a crucial role in the success of the Americans in the space race.

Here, bits are encoded by the magnetization of individual cores. To establish its directions, a core must undergo exposure to a magnetic field whose intensity surpasses a specific threshold. Let's denote I_0 the current required to impose this field. Consequently, the value of a particular bit can be set by passing a current of $I_0/2$ through the singular pair of horizontal and vertical wires that traverse this individual ring (Fig. 1.4a).

On the other hand, to determine the value of a bit, the computer attempts to switch its state to 1 (Fig. 1.4c). Throughout this process, the magnetic field response of the core is monitored via the current it induces through an additional line called sense. If the core's state was already 1, the response is minimal. However, when the core transitions from 0 to 1, a current spike occurs. The reading process is destructive, leaving the bit in the 1 state. Hence, after each read, its state must be promptly rewritten.

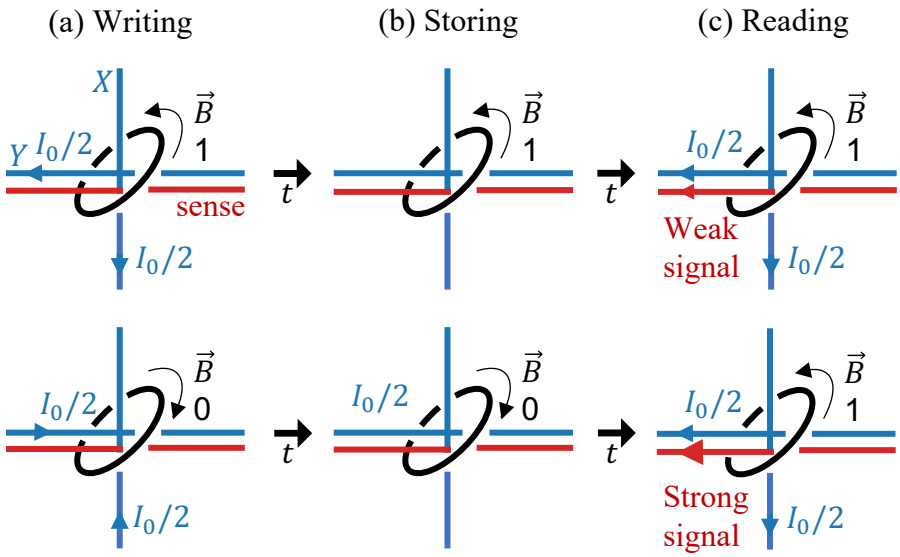

Figure 1.4 Magnetic core memory writing and reading cycles. (a) To set the state of a core (black circle), half the current necessary to flip its magnetisation is applied through the unique X and Y lines (blue solid lines) that passe through it. The direction of the magnetisation encodes either a 0 (bottom diagram) or a 1 (top diagram). (b) Once the magnetisation is set, the information is stored. (c) The system tries to flip the state of the core to 1 in order to extract its state. If the core is already in 1, no current is induced in the sense line (red solid line). Otherwise, the toggleling provokes a current spike that is interpreted by the system.

1.2.3 Classical error detection

Since the 1970s, computing capabilities have rapidly evolved. In a very small volume, silicon-based circuits allow for a density of bits that the mind struggles to conceive. For instance, among many technologies, today's computers can make use of Synchronous Dynamic Random Access Memory (SDRAM) to store ongoing calculations. A typical chipset [*Micron* 2015] can accommodate around 8 gigabits of data within a mere 1 square centimeter. To put things in perspective, this is roughly a tenth of the number of stars in the Milky Way galaxy.

In the realm of classical computing, the race for dense and energy efficient information storage medium is a very active and rich field in itself,

with very rich physics yet to be explored. For instance, magnetic tunnel junctions, operated through spin polarized current, allow to build nanoscale memories that can be operated at very reduced energy [Locatelli et al. 2014; Torrejon et al. 2017].

With such an incredible information density, errors can occur. Despite their high performance, individual bits are prone to failure due to various factors such as cosmic rays, electromagnetic interference, and transistor failures. The field of error correction is borne from the need to robustly encode information on multiple physical bits.

To illustrate our discussion in this section and the following one, based on [Hu et al. 2019], we will stick to the classic scenario where Alice wants to transmit information to Bob. Now, I'd love to shake things up a bit, but alas, given the name of the company I work for, I'm sworn to uphold tradition.

Alice sends a string of n bits to Bob. We denote p the probability for an error to occur during the transmission. Before diving into codes that can correct errors, let's start simple and consider the case where Bob can detect bit-flips but can't correct them.

To make errors detectable, Alice sacrifices a bit of data in favor of redundancy. The first $n-1$ bits encode the message, while the last bit serves as a parity check. Its value is chosen such that the total number of 1s in the string is even.

The message is transmitted with no errors when each individual bit remains unchanged. Assuming bit-flips are independent, the probability of a correct transmission is therefore given by:

$$P_{\text{correct}} = (1-p)^n \underset{p\to 0}{\approx} 1 - np \tag{1.1}$$

For Bob to detect an error, an odd number of bits has to flip. Using binomial coefficients, we can derive an expression for this probability to the leading order for small p:

$$P_{\text{detected}} = \sum_{1+2k\leqslant n} \binom{n}{1+2k} p^{1+2k}(1-p)^{n-1-2k} \underset{p\to 0}{\approx} np \tag{1.2}$$

Conversely, when an even number of errors occurs, they slip under the radar, leading to:

$$P_{\text{undetected}} = \sum_{0<2k\leqslant n} \binom{n}{2k} p^{2k}(1-p)^{n-2k} \underset{p\to 0}{\approx} \frac{n(n-1)}{2} p^2 \tag{1.3}$$

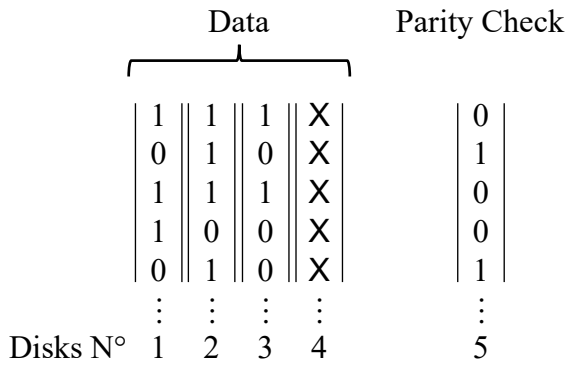

Figure 1.5 Parity check for data storage. An array of 5 hard drives is shown. The first 4 store the useful data, while the 5th is dedicated to redundancy. If one of the drives fails or is damaged, no data is lost. In this instance, the 4th drive is destroyed, yet its contents can be reconstructed as $1, 0, 1, 1, 0, \ldots$.

Therefore, assuming $np/2 \ll 1$, Bob can correctly detect errors during transmission. This technique is extremely cheap to implement and forms the backbone of many telecommunication protocols. However, I think it is also crucial to emphasize its significance in the context of data storage.

Let's consider a scenario where Alice has n hard drives to store her personal photos Fig. 1.5. For the sake of this example, let's say each drive has a capacity of 100 GB $= 8 \times 10^{11}$ bits. Fast forward a decade, and Alice's son, Bob Junior, seized by a nostalgic impulse, wants to revisit her photos. Now, let p represent the probability of a hard drive failure during this period. This time, the error we're concerned with is data loss rather than data corruption.

However, to prevent this from happening, Alice can once again employ the parity check strategy. The first $n-1$ drives encode the data, while the last one only stores parity checks. To clarify, what I mean is that for $m \in [1, 8 \times 10^{11}]$, the m^{th} bit of the n^{th} drive is the parity check of the m^{th} bits for the drives 1 to $n-1$.

Notice how, in this scenario, if a single drive is lost, no information is destroyed, as its state can always be reconstructed from the remaining others. By performing parity checks of parity checks, an arbitrary fraction of the n drives can be allocated to redundancy. For instance, RAID 6, one of the most commonly deployed data storage virtualization technologies, uses two disks for parity checks.

As a practical side note, the geek inside me wants to mention that having a single layer of redundancy for data storage is often not sufficient. In general, a disk failure is not an isolated event. While performing the replacement, you are already stressed, and you certainly don't want to feel the pressure of potentially losing all your data during the operation. This is why, when setting up the server bay for storing your scientific results, opting for at least RAID 6 is advisable.

1.2.4 Classical error correction

We now return to the scenario where Alice sends a message to Bob. For him to have any hope of correcting an error, he must at least know its location in the bit string. The value of the parity check bit does not encode this information, making error correction impossible with this simple technique. Perhaps the simplest scheme for error correction is the repetition code, of which the most basic instance uses 3 physical bits to encode the state of a unique logical bit.

$$\begin{aligned} 0_L &= 0.0.0 \\ 1_L &= 1.1.1 \end{aligned} \tag{1.4}$$

To send her message, Alice divides the bit string into $n/3$ registers. Therefore, Alice sacrifices $2/3$ of the storage capacity for redundancy, which is highly inefficient, but bit-flips can now be easily located. If Bob receives the string $0.1.0, 1.1.1, ...$, obviously, an error occurred in the first register. To recover its original value, assuming a single bit-flip happened, Bob can perform a majority vote to recover its original value, in this case $0.0.0, 1.1.1, ...$

We say that $0.0.0$ and $1.1.1$ are logical codewords, which should be seen as points in the register state space $\{0, 1\}^3$ as illustrated in Fig. 1.6. A bit-flip moves the state of the register within this space. As long as it stays close enough to the value of its original codeword, a correction can be achieved. The natural metric in this space is the Hamming distance, which simply counts the number of positions where two register states differ.

As long as the Hamming distance between the current state of the register and its original value stays below 1, errors are correctable. We say that we have a distance-1 code. Denoting $P_n, n \in 0, 1, 2, 3$ as the probability of n bit-flips occurring, we can compare the error rate of the code with that of the bare physical bit:

$$P_2 + P_3 < p \tag{1.5}$$

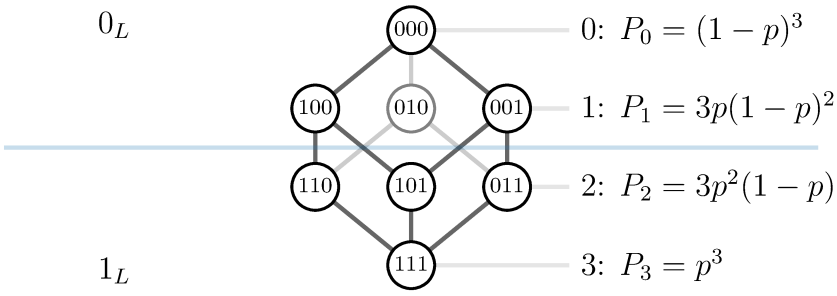

Figure 1.6 State space of a register for a repetition code with 3 physical bits. The points of $\{0,1\}^3$ are displayed as the vertices of a cube. Starting from the logical zero state $0_L = 0.0.0$, one can move from one state to another through the edges of the cube by flipping bits one by one. The Hamming distances from the starting point are displayed on the right column, along with their associated probabilities. The blue horizontal line indicates the point beyond which errors are not correctable.

Replacing with expressions of Fig. 1.6, this inequality signifies that the code is effective when $p < 1/2$, a value referred to as the error threshold of the code.

1.2.5 Bistable systems

The importance of dissipation Error mitigation strategies come in various forms. Thus far, we've explored protocols that operate at the software level, requiring continuous error correction operations from the computer. Alternatively, a hardware-level approach can be employed. Instead of encoding the state of a bit on a physical quantity that can continuously drift over time, it is intriguing to investigate systems with multiple stable points. Specifically, systems which have two hold immense value for information processing.

Nothing is more boring than observing a well-crafted bistable system, since its main feature is to stay in place. Yet, they quietly surround us, making their presence known only when they toggle - like when your neighbor on the train annoyingly clicks their ballpoint pen. Upon closer examination, in this modern age, they are everywhere: from the satisfying closure mechanism of your earbuds case to the switches controlling your oven.

In fact, the reliability of bistable systems makes them highly valued for

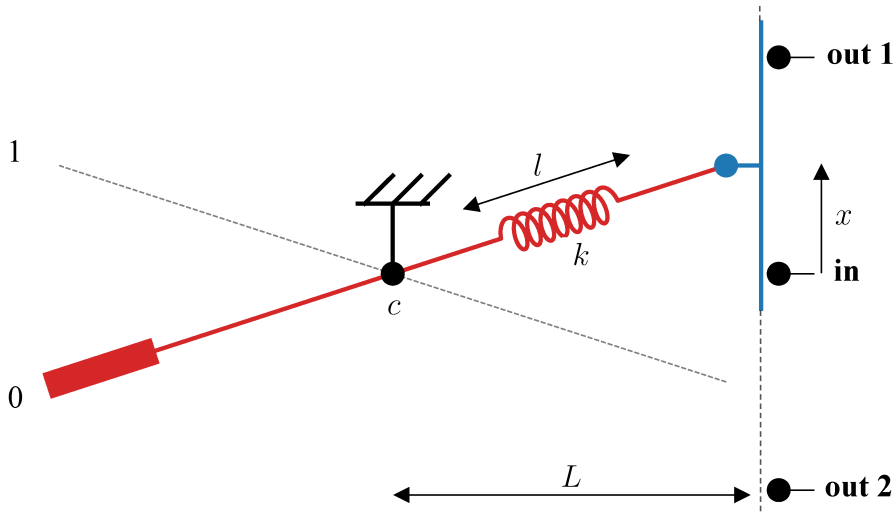

Figure 1.7 Representation of a double toggle switch. The red lever exhibits two stable positions thanks to a spring working in compression with a resting length of l and a spring constant of k. Positioned at a distance L from the lever pivot, a conductive broom establishes electrical contact between **in** and either **out1** or **out2** based on the switch's orientation. The system's stability is ensured by the friction coefficient c of the pivot.

the control interface of digital computers. The double-toggle switch mechanism Fig. 1.7 is an ubiquitous and robust design [Deb and Sen 2014]. It can be thought of as a memory that is operated by hand, and whose state can be accessed by the computer. From trains to airplanes, the daily lives of thousands of people rely on this design, as it covers their control panels. It basically never fails.

A state is stable when the system returns to it after a disturbance. Stability often arises from the friction. Without friction, energy cannot dissipate, resulting in endless oscillations around the energy minimum.

Therefore, a primary approach to constructing a bistable system is to introduce a double-well potential combined with significant dissipation. For example, we can analyze the energy potential $U(x)$ of the double-toggle switch mechanism to reveal this structure.

$$U(x) = \frac{1}{2}k(\sqrt{L^2 + x^2} - l)^2 \tag{1.6}$$

In this scenario, dissipation arises from the friction in the switch's joints,

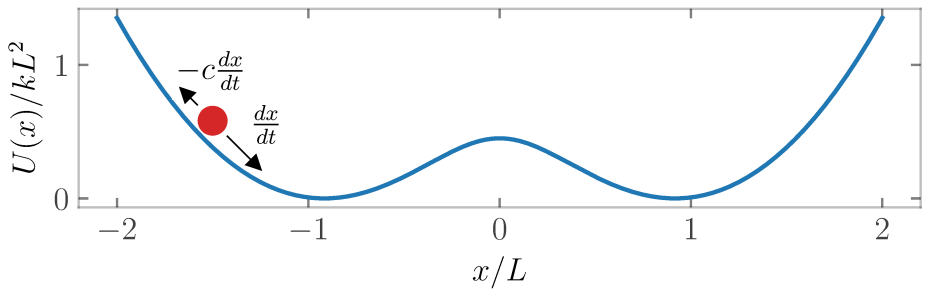

Figure 1.8 Double well potential. The potential energy of the double toggle switch is depicted (y-axis) relative to the position of the conductive broom (x-axis) responsible for establishing electrical contact. In this representation, the lever can be envisioned as a particle subject to gravity within the potential well, experiencing friction as it moves.

which must be precisely calibrated. Excessive friction renders the switch inoperable, while insufficient friction causes it to toggle with any perturbation. In an airplane, for instance, it's critical that the landing gear lever remains engaged during turbulent landings. Denoting c as the friction coefficient, we can express the equation of motion as:

$$\frac{d^2x}{dt^2} = \frac{dU}{dx} - c\frac{dx}{dt} \tag{1.7}$$

When toggling the switch, the energy barrier of the potential must be overcome Fig. 1.8. As the state transitions from one well to the other, this energy is rapidly dissipated, resulting in a distinctive clicking sound.

Bistability from parametric pumping Another insightful example of bistability can be found in everyday life, albeit seemingly unrelated to information processing. Among children's activities, the swing holds a special place. Through experimentation, one may have stumbled upon a neat trick as a child.

Here's how it works: you start by standing on the swing, then you crouch and wait until it comes to a standstill. Next, by alternating between standing and crouching at the right frequency, you can set the swing in motion. To achieve this, the frequency at which you stand and crouch should be twice the natural frequency of the swing (Fig. 1.9). If done correctly, you'll amplify the swing's motion until it reaches its maximum amplitude. It's worth

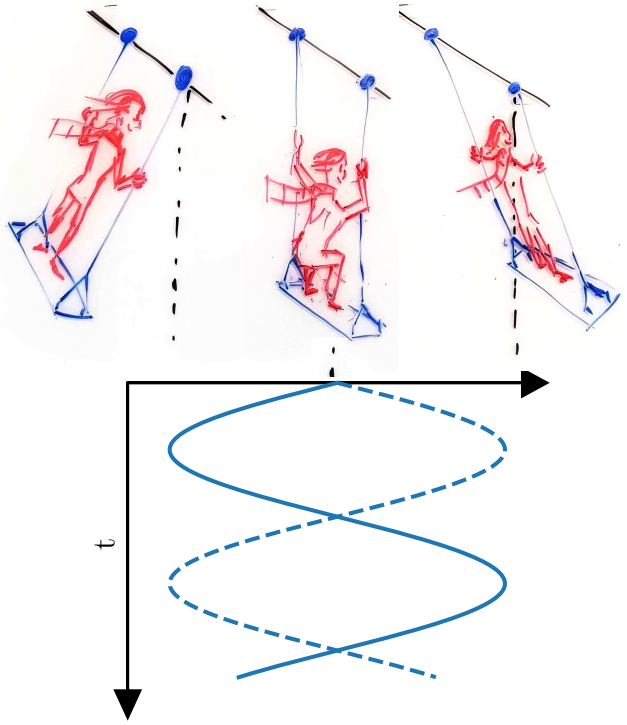

Figure 1.9 Parametric pumping of a swing. The swing's position (x-axis) is plotted against time (y-axis). The child stands upright when the swing reaches maximum amplitude and crouches when it is at the midpoint. Beginning from the middle position, the swing exhibits two stable trajectories (solid and dashed blue lines).

noting that this method differs from the traditional way of swinging, which involves sitting and moving your legs at the same frequency as the swing.

Here, as the swing moves from right to left or left to right, you'll alternate between standing, crouching when the swing is at its lowest point, and standing again when it reaches the other side. Starting from the middle position, the swing appears to have a "choice" for its initial direction. Is it going to move left or right first? Initially, the situation is symmetric, but once motion is imparted to the swing, symmetry is broken. If you record multiple videos of this phenomenon, synchronizing them with the motion of the child, you'll observe that the swing has two stable trajectories with opposite phases. These trajectories can be viewed as system states, except this time they are dynamic.

This phenomenon is an instance of parametric pumping. Essentially, what we're doing is periodically moving the center of mass of the pendulum, which is the swing. Later, we'll explore this principle further in the context of cavity QED to stabilize cat states.

1.3 Quantum computing

A quantum computer is a device that leverages quantum physics to accelerate computations that would be impractical with classical systems. Analogous to how a mechanical model might simulate the moon's orbit, one could envisage an "analog" quantum computer designed to investigate specific phenomena. Such devices, known as quantum simulators, are already within our technological grasp. Rydberg atoms arranged in optical tweezers have already been utilized to model the antiferromagnetic properties of many-body systems [Scholl et al. 2021].

Now, to explore the potential of a universal quantum computing, we must introduce the fundamental components of quantum information and examine their properties.

1.3.1 The quantum bit

So far, we discussed how to encode information on supports that are classical in nature. However, from a conceptual standpoint, there's no barrier preventing us from employing objects governed by quantum effects. The spin of an electron, its energy levels within an atom, or even the polarization of a photon are examples of quantum systems that can be operated and measured. This curiosity among physicists in the 1980s prompted speculation about whether quantum physics could inaugurate a new era in computing science.

While a classical bit is confined to the values 0 and 1, its quantum counterpart can exist in any superposition of its orthogonal basis states $|0\rangle$ and $|1\rangle$, also referred to as the computational basis. Generally, the state of a qubit is written $|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle$, where α and β are complex amplitudes satisfying $|\alpha|^2 + |\beta|^2 = 1$. As in quantum physics, the global phase can always be ignored, α can be taken as real, leading to an explicit parameterization:

$$|\psi\rangle = \cos(\theta/2)\,|0\rangle + e^{i\varphi}\sin(\theta/2)\,|1\rangle \tag{1.8}$$

Here, $\theta \in [0, \pi]$ and $\varphi \in [0, 2\pi[$ are real numbers analogous to the latitude and longitude on the surface of a sphere Fig. 1.10. Given this picture,

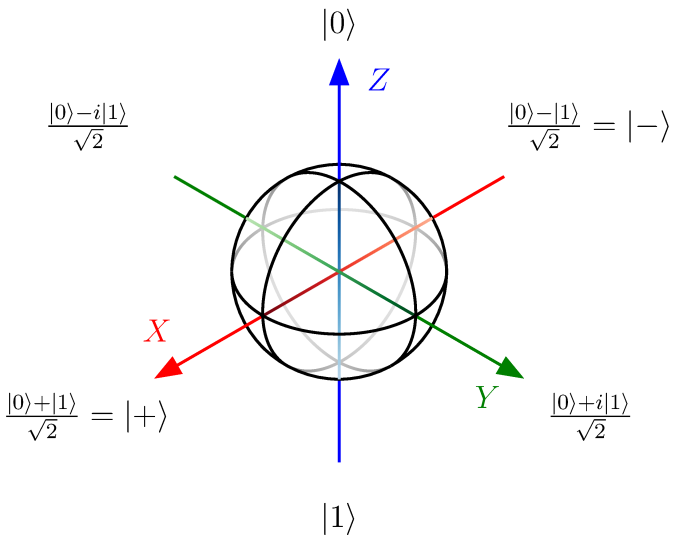

Figure 1.10 Representation of the Bloch sphere: The state $|\psi\rangle$ of a qubit can be envisioned as a point on the surface of the Bloch sphere. The computational basis states $|0\rangle$ and $|1\rangle$ correspond to the north and south poles, respectively. States along the equator are superpositions of computational states in equal proportions. Notably, states along the X axis are labeled $|+\rangle$ and $|-\rangle$. In this picture, θ from Eq. 1.8 is the latitude and φ is the longitude.

it becomes evident that a qubit has a lot of degrees of freedom. Its state belongs to a continuum and can either drift from one pole to the other or parallel to the equator. If these events are uncontrolled, they yield errors. The first kind can be identified with a classical bit-flip, while the latter is specific to quantum physics. Since it affects the relative phase of the complex numbers α and β, it is referred to as a phase-flip.

Despite the continuous nature of the state space, upon measurement a qubit seems to be completely classical. In the Z-basis, The output is either 0 with probability $|\alpha|^2$ or 1 with probability $|\beta|^2$. Subsequently, the system collapses into either $|0\rangle$ or $|1\rangle$ respectively. The projective nature of the measurement destroys quantum superpositions. Therefore, fully characterizing a state requires preparing it many times and accumulating statistics from multiple measurements.

The appeal of quantum computing becomes evident when one considers the state of an n-qubit register. In this case, the Bloch sphere representation breaks down, as a complex amplitude is associated with each classical state

of the total register. For instance, with 2 qubits, we have:

$$|\psi\rangle = \alpha\,|00\rangle + \beta\,|01\rangle + \gamma\,|10\rangle + \delta\,|11\rangle\,, \quad |\alpha|^2 + |\beta|^2 + |\gamma|^2 + |\delta|^2 = 1 \tag{1.9}$$

In general, 2^n complex numbers are needed to write the state of an n-qubit register. Accounting for the normalisation and the global phase, the dimension of state space manifold equals $4^n - 2$.

A well-designed quantum algorithm leverages this vast computational space by entangling multiple qubits. A highly entangled state is first prepared. Subsequently, through meticulously designed interference effects, the state undergoes a controlled evolution, gradually collapsing into a superposition of a select few states from the computational basis. Finally, these states are measured.

Unfortunately, these algorithms are highly susceptible to noise. Even a slight drift of individual qubits can completely disrupt their output. Similarly to the classical case, error correction is necessary. However, the continuous nature of the state space makes the situation seem utterly hopeless.

The very counterintuitive quantum error discretization theorem [Bennett et al. 1996, Shor 1997, Ekert and Macchiavello 1996] states that if one can correct for randomly occurring π rotations around the X and Z axes of the Bloch sphere, then any errors are correctable. Despite the state being continuous, correcting it requires only a finite set of operations.

The reason behind this lies in the fact that at any given moment, the information leaking out of a quantum system is classical in nature. This is why in reality, the interaction of the environment with the system is highly constrained, limiting the type of errors that can occur and the resources needed to mitigate them.

1.3.2 Quantum error correction

The main idea of quantum error correction is to use multiple qubits to encode one. The original quantum error correction code proposed by Shor [Shor 1995] is analogous to the repetition code. The complex amplitudes α and β of the logical qubit are encoded into two code words, which are simply obtained by placing kets around Eq. 1.4:

$$\begin{aligned} |0_L\rangle &= |000\rangle \\ |1_L\rangle &= |111\rangle \end{aligned} \tag{1.10}$$

Therefore, the state of this qubit reads $|\psi\rangle = \alpha\,|000\rangle + \beta\,|111\rangle$. Individual qubits are prone to errors, but the idea is that the 8 computational states of the register give us enough room to recover from a single bit-flip.

Contrary to the classical case, however, error detection should be performed indirectly. If at any point individual qubits are measured, the state gets projected, and the values of the complex amplitudes are lost. The solution is to measure the joint parity operators $\mathbf{Z}_1\mathbf{Z}_2$ and $\mathbf{Z}_2\mathbf{Z}_3$ Fig. 1.11a. As long as the system remains in the code space, the measurement of this error syndrome yields 0 and 0 independently of the qubit state. No information is leaked, and the quantum superposition is preserved.

In the event of a single bit-flip, measuring the error syndrome allows us to identify the faulty qubit without revealing its state Fig. 1.11bc. Applying an $\mathbf{X}$-gate on the identified qubit brings us back into the code space Fig. 1.11d, effectively correcting the bit-flip error. Similar to classical error correction, increasing redundancy enhances robustness. However, we don't have a logical qubit yet, as robustness against bit-flips is not equivalent to robustness against phase-flips, which require another level of protection.

A direct approach to achieve this is by layering repetition codes on top of another. The first layer is the one we just described, while the second slightly differ. The error syndrome is measured along the X quadrature of the Bloch sphere, specifically $\boldsymbol{X}_1\boldsymbol{X}_2$ and $\boldsymbol{X}_2\boldsymbol{X}_3$. What we get is a qubit protected against $\mathbf{Z}$ and $\boldsymbol{X}$ errors, satisfying the quantum error correction theorem. Ultimately, Shor's original repetition code for quantum error correction consisted of 9 qubits [Shor 1995]. Subsequent research has delved deeper into error-correcting schemes, resulting in numerous proposals.

We have just provided a concrete example of quantum error correction. Now, let's take a step back and reevaluate the requirements of a quantum computer from a more abstract, information-based perspective.

Imagine that Alice is not very careful and inadvertently sends an energetic particle to Bob. The particle transfers a fraction of its energy to Bob's skin, causing him harm. The resulting burn on Bob's arm serves as a record of the particle's trajectory—a form of measurement. In the realm of quantum computing, when bits are scrambled by external disturbances, the scenario is quite similar. A bit-flip error represents an instance of the computing device unintentionally measuring its environment in an uncontrolled manner.

Now, Bob sends Alice a spin, and she measures it, resulting in an outcome of either up or down. Crucially, all information about the quantum phase is lost after the measurement. When the state of a quantum computer is measured, all quantum superpositions collapse. Consequently, to conduct

Logical state (a)	Errors (b)	Measurement (c)
$\|\psi\rangle$ = $\alpha\|000\rangle + \beta\|111\rangle$	$\alpha\|000\rangle + \beta\|111\rangle$	00
	$\alpha\|001\rangle + \beta\|110\rangle$	01
	$\alpha\|010\rangle + \beta\|101\rangle$	10
	$\alpha\|100\rangle + \beta\|011\rangle$	11

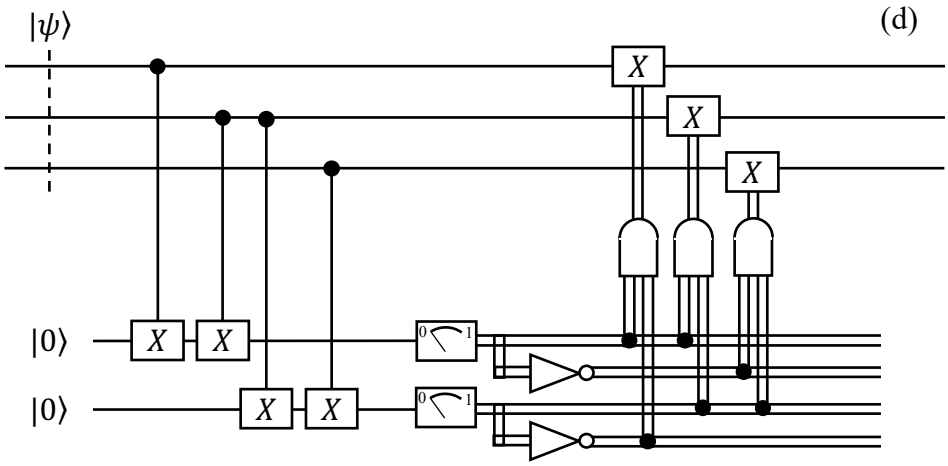

Figure 1.11 Repetition code circuit for bit-flip protection. (a) The state of the logical qubit $|\psi\rangle$ is redundantly encoded in 3 qubits. (b) Assuming at most a single qubit error occurs, there are 4 possible states the qubit can be in. (c) The measurement of the error syndromes allows identification of the faulty qubit. (d) A cycle of the repetition code is depicted. Solid lines represent quantum states, while double lines represent classical signals. Initially, $|\psi\rangle$ is prepared on 3 data qubits. With the help of two ancilla qubits, the error syndromes $\mathbf{Z}_1\mathbf{Z}_2$ and $\mathbf{Z}_2\mathbf{Z}_3$ are measured after the application of 2 controls X-gates on each. Depending on the outcome, individual X-gates are applied to the data qubits.

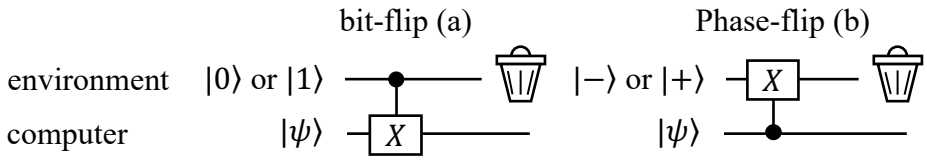

Figure 1.12 Representation of bit-flip and phase-flip. (a) The quantum circuit illustrates a bit-flip event. The environment, in an unknown state, may induce an error on a data qubit $|\psi\rangle$ within the computer via an X-gate operation. This error corresponds to an uncontrolled π-rotation around the X axis when viewed from the Bloch sphere of the data qubit, hence representing a bit-flip event. (b) Depicts the dual circuit for a phase-flip event. In this scenario, the control X-gate originates from the computer and acts on the environment. From the perspective of the data qubit, this results in an uncontrolled Z-gate operation.

proper quantum calculations, the computer must be refrained from measuring the environment, and conversely, the environment should not measure the computer's internal state Fig. 1.12. In essence, for effective operation, a quantum computer requires a bidirectional causal disconnection between itself and its environment. Moreover, this disconnection should be adjustable, allowing the operator to retrieve the computation's outcome at the conclusion of the calculation.

Measuring the error syndrome of the repetition code neither induces errors in the quantum computer's state nor leaks information about it. By activating and deactivating this correction scheme, a tunable disconnection is effectively implemented.

1.3.3 Quantum error correction in practice

On paper, quantum error correction offers a succinct solution to what initially appears as an intractable problem. Its theoretical framework is rich and fertile, and there is a certain mathematical beauty to it. Given this framework, let us derive the smallest number of physical qubits needed to implement a correction code.

If n physical qubits are used to encode quantum information, the error

syndrome measurement should output no more than $n-1$ bits. Otherwise, there is insufficient "room" to conceal the qubit's state from the environment. For example, in the simple 3-qubit repetition code we discussed earlier, the syndrome measurement leaks only 2 bits (Fig. 1.11c). Now, for correction to be effective, the syndrome must indicate whether an error occurred and, if so, which of the n physical qubits is faulty. With 3 types of errors ($\boldsymbol{X}$, $\mathbf{Y}$, and $\mathbf{Z}$), this gives rise to $3n+1$ possibilities. Thus, we have the following inequality:

$$3n + 1 < 2^{n-1} \tag{1.11}$$

This equation is satisfied for $5 \leqslant n$, and indeed it is possible to give an explicit construction of the 5-qubit code [Gottesman 2009]. Since it requires few physical qubits, we say it introduces a small hardware overhead. However, far from its abstract elegance, the practical implementation of quantum error correction is met with a lot of prosaic challenges, leading to a dramatic increase of this overhead.

Executing operations with physical qubits takes time, leaving room for errors to occur during gates, a factor we previously overlooked. Additionally, measuring an error syndrome, such as the joint parity operator $\mathbf{Z}_i\mathbf{Z}_j$, necessitates additional qubits, as illustrated in Fig. 1.11d. There are two classes of physical qubits now. The ones designated for storing information are data qubits, while those employed for syndrome measurement are termed ancilla. And of course, both are susceptible to errors.

Another limitation comes with the connectivity one can achieve for a given platform. In reality, qubits are arranged in 1D, 2D, or even 3D structures, and implementing long-range interactions may be extremely challenging. For instance, with superconducting qubits, gates are typically only feasible between the nearest neighbors of a 2D grid. The necessity to perform error correction under this stringent constraint drove the development of an approach known as the surface code [Dennis et al. 2002], which is now very popular.

In this architecture, data and ancilla qubits are arranged in an interlocking grid pattern (Fig. 1.13). Ancilla qubits measure the joint $\mathbf{Z}$ or $\boldsymbol{X}$ error syndromes of their neighboring data qubits: 4 if we are inside the grid, and 2 on the borders. The logical states are highly entangled and completely delocalized on the grid. Interestingly, the measure along the Z axis of the logical qubit can be achieved by measuring the joint parity operator of data qubits along any horizontal lines, while the X measurement is analogous but with vertical lines.

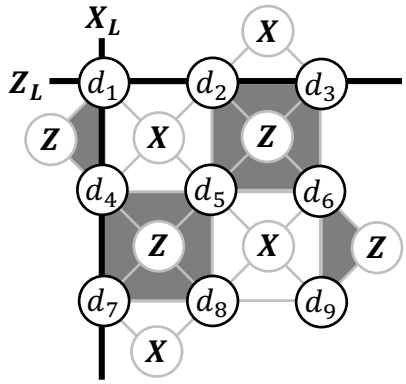

Figure 1.13 Surface code. Depiction of a 3×3 surface code. Data qubits are represented by black circles, while ancilla qubits are represented by gray circles. Ancilla qubits measure the joint $\boldsymbol{X}$ or $\mathbf{Z}$ error syndromes of their neighboring data qubits. The logical operators $\mathbf{Z}_L$ and $\boldsymbol{X}_L$ are obtained by taking a line and a column of data qubits, respectively.

When a physical qubit is compromised by an error, it is flagged by the neighboring ancilla qubits. The location of the error is therefore known. If the flag is raised by $\mathbf{Z}$, $\boldsymbol{X}$, or both types of ancilla, it indicates an $\boldsymbol{X}$, $\mathbf{Z}$, or $\mathbf{Y}$ error, respectively. For a logical error to occur, an entire row or column of physical bits must flip simultaneously, rendering the code highly robust.

The structure of surface code makes it inherently scalable; however, similar to the two-layer repetition code proposed by Shor, its scaling is quadratic with the error rate of the physical qubits, resulting in significant hardware overhead with current technology. It is estimated that implementing a single logical qubit requires approximately 10^3 physical qubits [Ruiz et al. 2024a]. Moreover, in order to run Shor factorization algorithm, it is projected that approximately $10^{3\sim4}$ logical qubits will be needed [Gidney and Ekerå 1905; Beverland et al. 2022], translating to the requirement of millions of physical qubits. To put things in perspective, the latest implementation by Google of a surface code is comprised of 25 data qubits ["Suppressing quantum errors by scaling a surface code logical qubit" 2023].

The surface code has many symmetries. One can notice that it is left unchanged after a 90 degrees rotation followed by an inversion of $\boldsymbol{X}$ and $\mathbf{Z}$ ancilla qubits. This is not a coincidence, since it treats bit-flips and phase-flips symmetrically. Apart from elegance, the motivation behind this choice is the fact that for most qubit technologies, both types of error are equally

likely. However, this is not a necessity, and it is extremely insightful to consider what happens when a channel has an error rate much lower than the other.

1.3.4 Biased noise qubit

Let us say we have a physical qubit with a bit-flip rate so low we can effectively assume phase-flip errors are the only to occur. At first glance, this scenario seems highly promising, as a simple repetition code tailored for phase-flip errors could suffice to create a logical qubit. Since this would require much fewer qubits than a complete surface code correcting both types of errors, it is reasonable to expect a significant reduction in hardware overhead.

For instance, research has shown that a single electron spin confined within a nitrogen-vacancy (NV) center in a diamond at temperatures of a few Kelvins exhibits a natural bit-flip time exceeding an hour, while its phase-flip time is on the order of seconds [Abobeih et al. 2018]. By applying optical pulses to the system, it is feasible to execute gates and measure the state of the physical qubit. Consequently, in theory, it should be achievable to conduct quantum computations by integrating these spins with an optical quantum network.

With error channels now exhibiting significant discrepancies, we describe the qubit as having a noise bias. The challenge here lies in maintaining this bias during operations. If at any time a gate maps a phase-flip error to a bit-flip error, the bias is not preserved, and the advantage it offers is gone. Consequently, the range of permissible gates becomes highly constrained, which raises two pivotal questions: What is the set of gates that preserves the bias? And is it large enough for us to perform universal quantum computations? A detailed answer to these questions is beyond the scope of this document, but let us give some main ideas.

On one hand, let's assume that we encode quantum information at the physical level using a collection of TLS, such as the spins in NV centers. For a single qubit, it's straightforward to visualize the rotations that preserve the bias. On the Bloch sphere, they should not mix the meridians and parallels. Therefore, the only admissible rotations are those with an arbitrary angle θ around the Z axis. It may appear that π rotations around the X or Y axis also preserve the bias. However, they are not feasible in practice. A gate is the result of a continuous process, and if, at any given point in time, the state of a qubit is exclusively supported on $|0\rangle$ and $|1\rangle$, then to move continuously

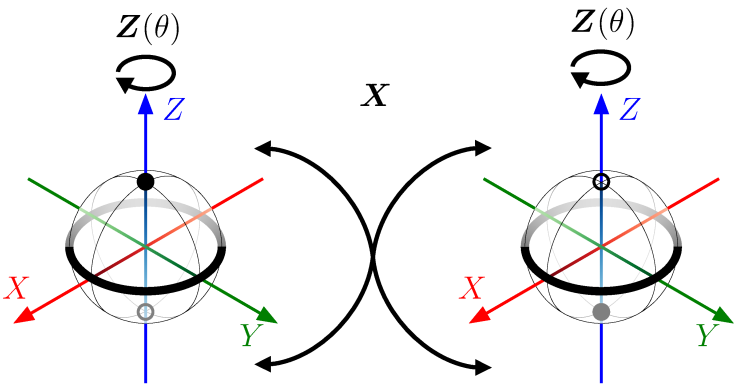

Figure 1.14 Bias-preserving single-qubit gates. The only bias-preserving single-qubit gates are rotations around the Z axis or the π rotations that flip the Bloch sphere upside down, such as the $\boldsymbol{X}$ gate. If one has the capability to prepare the qubit in the states $|0\rangle$, $|1\rangle$, and $|+\rangle$, then the attainable states are limited to the poles (solid and empty black circles) and the equator (solid black line).

from one pole of the Bloch sphere to the other, we must traverse the equator, at which point the noise bias is lost [Guillaud and Mirrahimi 2019, p. 19].

As explained in [Puri et al. 2020], this limitation is significant. The CNOT gate, which is a 2-qubit gate that performs an $\boldsymbol{X}$ rotation of a target qubit conditioned on the state of a control qubit, is therefore not achievable with a TLS. Unfortunately, the CNOT gate is commonly used to extract the value of error syndromes in many protocols, particularly those tailored for biased noise qubits [Tuckett et al. 2018]. It is still possible to extract error syndromes using CZ gates with a teleportation scheme [Aliferis and Preskill 2008], but this adds significant overhead. At first glance, this complexity makes biased noise qubits somewhat unappealing.

However, realizing a qubit embedded in a physical system with a large Hilbert space opens up the possibility of utilizing additional dimensions to execute rotations without mapping the Bloch sphere onto itself during intermediate steps (see Fig. 1.14). As we will see, this makes bosonic systems excellent candidates for biased noise qubits. With this enhanced capability, it is now possible to implement low-overhead error correction codes. Given the ability to perform the $\mathbf{Z}(\theta)$ and the CNOT gates, as well as the preparation of a state $|+\rangle$ and the measurement of the X observable, then universal

quantum computations can be performed [Guillaud and Mirrahimi 2019].

Operating a system which Hilbert space is large undoubtedly introduces complexity, but the potential benefits may well outweigh the challenges. For instance, the first practical applications of the Shor algorithm, aimed at breaking currently deployed cryptographic protocols, are projected to require around 10^5 physical qubits with reasonable noise bias, operated via a repetition codes [Gouzien et al. 2023]. This already represents an order of magnitude reduction compared to the surface codes. Furthermore, if the objective is to perform physical simulations that are infeasible on classical computers, the outlook is even more promising. With approximately 2×10^2 logical qubits, each with a per-gate-cycle error probability of less than 10^{-8}, it becomes feasible, for example, to simulate the dynamics of a 2D Hubbard model with a 10×10 lattice and around 100 particles [Daley et al. 2022]. By employing more sophisticated coding techniques than the simple repetition code, all these logical qubits can be encoded using only about 10^3 physical noise-biased qubits [Ruiz et al. 2024b], bringing us even closer to experimental feasibility.

Taking a step back, it is important to note that the definition we provided for a biased noise qubit is somewhat flawed. The classical bits inside your smartphone have exceedingly low bit-flip rates, while their phase-flip rates are so high they effectively behave classically, and are therefore not qubits. What's missing here? Well, for a system to qualify as a qubit, the gate time must be shorter than the typical time it takes for an error to occur on both channels. Classical silicon-based computing fails to meet this criterion, preventing its utilization for quantum computing.

Yet, this raises a question: Is it possible to create a quantum version of a classical bistable device ? After all, why can't we benefit from the inherent robustness these systems display in the classical realm ? At first glance, however, this notion seems entirely implausible. The dissipation that renders classical bistable systems stable also leaks energy, thereby divulging information about the state, which induces phase-flip errors. In its classical design Fig 1.7, what stabilises a switch also prevents it to exist in quantum superpositions. However, we demonstrate in the next chapters that dissipation, like many things, can be engineered: It is possible to modify the classical design of the switch so that the energy it leaks does not "knows" about its state, making it compatible with quantum information processing.

2
THEORY

To delve deeper into the world of quantum bistable systems, it is imperative to agree on a suitable formalism. Throughout our discourse, we will presuppose a certain level of familiarity with various methodologies: The eigenmodes of dynamical systems will be expressed in their second quantized form. To characterize their states, we will employ the density matrix formalism, or equivalently, the Wigner function representation. And given that dissipation will be a central theme in our discussion, it is also essential to have a straightforward approach to address open quantum systems. Consequently, we will employ the Lindblad master equation to elucidate the time evolution. As a warmup exercise, let us exemplify how we employ these methodologies on the simple harmonic oscillator. The forthcoming paragraphs are not intended as a reiteration of introductory material on quantum physics, as this topic has been extensively covered elsewhere (for instance in [Haroche and Raimond 2006a]). Instead, we try build on the reader existing knowledge to cultivate a geometric understanding of the dynamics in phase space.

A harmonic oscillator models a mass attached to a spring (Fig. 2.1a). Its classical equation of motion is obtained from Newton's law, with acceleration on the left hand side, and forces on the right.

$$m\frac{d^2x}{dt^2}(t) = -kx(t) \tag{2.1}$$

An excellent way to depict the state of a harmonic oscillator is certainly the phase space representation (Fig. 2.1b). In this 2D picture, the horizontal axis represents the position x, while the vertical axis denotes the momentum $p = m\frac{dx}{dt}$.

Due to the time independence of the dynamics, each point in this space can be associated with a vector representing the infinitesimal motion of the oscillator. This vector field, referred to as a flow (Fig. 2.2), is orthoradial in the case of the harmonic oscillator. As a result, trajectories follow circular paths of constant radius.

$$x(t) = A_0\cos(\omega t + \theta_0), \quad \omega = \sqrt{k/m} \tag{2.2}$$

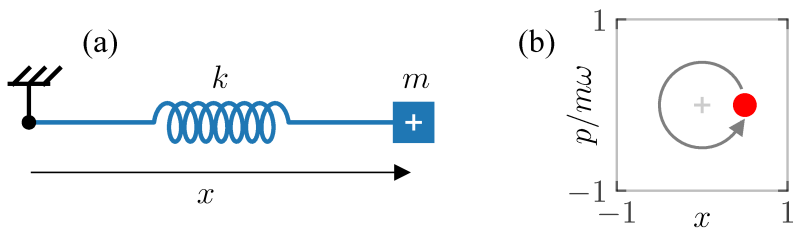

Figure 2.1 Harmonic oscillator. (a) In its simplest mechanical form, the harmonic oscillator comprises a spring k attached to an oscillating mass m. The position of the mass is denoted by x. (b) Phase space representation of the oscillator. The horizontal axis represents the position x, and the vertical axis represents momentum p. The state of the system is depicted as a point (red circle) which follows a circular trajectory in the absence of dissipation (solid black arrow).

Since the equation of motion is linear, the time it takes to perform a full revolution $2\pi/\omega$ is independant from the amplitude A. Adding a friction coefficient c to the model, a term proportional to the velocity appears.

$$m\frac{d^2x}{dt^2}(t) = -kx(t) - c\frac{dx}{dt}(t) \tag{2.3}$$

The circles transform into spirals (Fig. 2.2), and their radius decays exponentially at half the rate at which the system dissipates energy into its environment, denoted by κ_a.

Now, In the quantum realm, the Heisenberg uncertainty principle $\hbar/2 \leqslant \Delta x \Delta p$ implies that the state of the system cannot be represented as a point anymore, as it would require both momentum and position to be perfectly defined simultaneously. The quantum states that most resemble points, as they saturate the Heisenberg inequality, are called coherent states. In the Fock state basis, they are expressed as:

$$|\alpha\rangle = e^{-|\alpha|^2/2} \sum_{n=0}^{\infty} \frac{\alpha^n}{\sqrt{n!}} |n\rangle \tag{2.4}$$

Here the phase space is identified with the complex plan, $\alpha \in \mathbb{C}$ is a dimensionless number which real part encodes the position $\Re(\alpha) \propto \langle \boldsymbol{x} \rangle$, and its imaginary part the momentum $\Im(\alpha) \propto \langle \boldsymbol{p} \rangle$. The state is now blurred, and classical variables become the average values of operators (represented with bold characters). Conceptualizing and representing a quantum state can

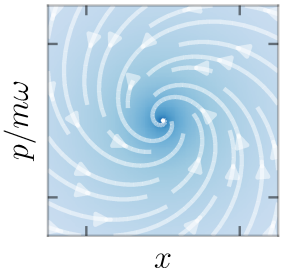

Figure 2.2 Harmonic oscillator flow. The horizontal axis represents the position x, and the vertical axis represents momentum p. The shades of blue in the picture are proportional to the magnitude of $|\dot{x} + i\dot{p}/m\omega|$, representing the amplitude of the flow, while white arrows depict typical trajectories. In the presence of dissipation, and in the laboratory reference frame, these trajectories take the shape of logarithmic spirals.

challenge intuition, especially considering our daily experiences are deeply rooted in classical phenomena. A technique that closely mirrors the phase space representation is known as the Wigner function $W(\gamma)$ (Fig. 2.3).

Its argument γ is the complex number coding for the phase space coordinate, and the Wigner function associates a real number to each of them. Under this representation, coherent states appear as Gaussian distributions evenly spreading across both quadratures. This image provides a sense of how the system's presence in specific points is now uncertain. However, it's crucial to emphasize it is not a probability amplitude. To access the distributions of $\boldsymbol{x}$ and $\boldsymbol{p}$, one must integrate this function along one quadrature or the other, essentially examining the shadow of the Wigner along a given direction (Fig. 2.3).

Now, we require a method to discuss the system's evolution. To this end, we describe its energy conservative part with the Hamiltonian written in the second quantized form. To achieve this, we introduce the usual annihilation and creation operators:

$$\mathbf{a} = \sqrt{\frac{m\omega}{2\hbar}}\left(\boldsymbol{x} + i\frac{1}{m\omega}\boldsymbol{p}\right), \quad \mathbf{a}^\dagger = \sqrt{\frac{m\omega}{2\hbar}}\left(\boldsymbol{x} - i\frac{1}{m\omega}\boldsymbol{p}\right) \tag{2.5}$$

Here, $\mathbf{a}$ and $\mathbf{a}^\dagger$ should be interpreted as operators that either extract (respectively inject) energy into the system one quantum at a time. From an

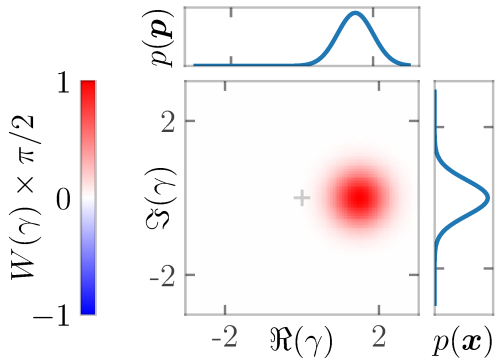

Figure 2.3 Wigner function of a coherent state. In the quantum realm, the state of a harmonic oscillator manifests as a coherent state. In phase space (x and y axes), its Wigner function adopts the shape of a 2D Gaussian (color). The probability distribution of both quadratures is obtained by integrating the Wigner function along the horizontal or vertical axes.

operational standpoint, it's important to recall two standard formulas: they follow to the usual commutation relation for bosonic modes, $[\mathbf{a}, \mathbf{a}^\dagger] = 1$, and the amplitude of a coherent state α is given by the expectation value of $\mathbf{a}$ via $\langle\alpha|\,\mathbf{a}\,|\alpha\rangle = \alpha$. With these operators, the Hamiltonian can be expressed in a compact form:

$$\mathbf{H}/\hbar = \omega \mathbf{a}^\dagger \mathbf{a} \tag{2.6}$$

And the Shrodinger equation allows to compute the time evolution:

$$\frac{d\boldsymbol{\rho}}{dt} = -\frac{i}{\hbar}[\mathbf{H}, \boldsymbol{\rho}] \tag{2.7}$$

In this context, $\boldsymbol{\rho}$ represents the density matrix operator, an alternate method of describing the state of a quantum system. Specifically, the expectation value of any operator $\boldsymbol{o}$ can be derived from $\boldsymbol{\rho}$ using $\langle \boldsymbol{o} \rangle = \mathrm{Tr}(\boldsymbol{\rho o})$. However, generally, we prefer working in the Heisenberg picture, where $\boldsymbol{\rho}$ remains constant, and instead, operators are now time-dependent.

$$\frac{d\boldsymbol{o}}{dt} = +\frac{i}{\hbar}[\mathbf{H}, \boldsymbol{o}] \tag{2.8}$$

Let us examine the evolution of the creation operator $\mathbf{a}$ in the Heisenberg picture:

$$\frac{d\mathbf{a}}{dt} = i\omega[\mathbf{a}^\dagger\mathbf{a}, \mathbf{a}] = -i\omega\mathbf{a} \tag{2.9}$$

This calculation reveals that, starting from a coherent state $|\alpha_0\rangle$, its complex amplitude rotates around the center of the phase space at a rate ω, resembling the classical scenario:

$$\alpha(t) = \alpha_0 e^{-i\omega t} \tag{2.10}$$

Therefore, it's customary to align ourselves with a frame that rotates at the same rate, simplifying calculations. In this frame, everything appears stationary, and the Hamiltonian reduces to a trivial expression $\mathbf{H}/\hbar = 0$.

Moving on to the treatment of dissipation, while a detailed exploration of the Lindblad master equation is outside our current scope, we will outline its expression, and comment on its structure:

$$\frac{d\boldsymbol{\rho}}{dt} = -\frac{i}{\hbar}[\mathbf{H}, \boldsymbol{\rho}] + \sum_{i=1}^{N} \boldsymbol{L}_i \boldsymbol{\rho} \boldsymbol{L}_i^\dagger - \frac{1}{2}\left\{\boldsymbol{L}_i^\dagger \boldsymbol{L}_i, \boldsymbol{\rho}\right\} \tag{2.11}$$

On the left hand side, we recognize the energy conservative part of the dynamics, while the right hand side allows to model the interaction with external baths. $\boldsymbol{L}_i$, for $i \in [1, N]$, are known as jump operators. Each represents a channel through which information about the state leaks into the environment. Again, it is possible to envision the evolution imposed by this equation as acting on operators rather than the density matrix, leading to the adjoint Lindblad master equation [Breuer and Petruccione 2002, p. 125]:

$$\frac{d\boldsymbol{o}}{dt} = +\frac{i}{\hbar}[\mathbf{H}, \boldsymbol{o}] + \sum_{i=1}^{N} \boldsymbol{L}_i^\dagger \boldsymbol{o} \boldsymbol{L}_i - \frac{1}{2}\left\{\boldsymbol{L}_i^\dagger \boldsymbol{L}_i, \boldsymbol{o}\right\} \tag{2.12}$$

Notice that compared to the equation obtained from the input-output theory [Gardiner and Collett 1985], we are missing the input noise, which accounts for the quantum fluctuations added by the thermal bath. For example, assuming $\boldsymbol{L}_i = \sqrt{\kappa_i}\mathbf{a}$, in the input-output theory we have the added terms $\sqrt{\kappa_i^c}\mathbf{a}_{\text{in}}^i$. This omission arises from the fact that in the adjoint master equations, operators are promoted to statistical mixtures, resembling density matrices. Consequently, if we were to plot the entropy $S = \text{Tr}(\boldsymbol{o}\log(\boldsymbol{o}))$ of a given observable $\boldsymbol{o}$, we would observe it increases over time.

If there is a simple energy leak induced by friction, it can be modeled with a single jump operator $\boldsymbol{L} = \sqrt{\kappa_a}\mathbf{a}$, where the annihilation operator $\mathbf{a}$

removes photons one by one from the oscillator at an average rate κ_a. Let us pause here to indicate that in the literature, this rate can be also denoted with κ or κ_a. Since later we will have to deal with multiple modes at the same time, we index the one-photon loss rate with them. From there, we express the master equation in the Heisenberg picture for $\mathbf{a}$ in the rotating frame:

$$\frac{d\mathbf{a}}{dt} = \kappa_a \left(\mathbf{a}\mathbf{a}\mathbf{a}^\dagger - \frac{1}{2}\mathbf{a}\mathbf{a}^\dagger\mathbf{a} - \frac{1}{2}\mathbf{a}^\dagger\mathbf{a}\mathbf{a} \right) = -\frac{1}{2}\kappa_a \mathbf{a} \tag{2.13}$$

And we find that it exhibits an exponential decay of the state amplitude.

$$\alpha(t) = \alpha_0 e^{-\kappa_a t/2} \tag{2.14}$$

Again, this is in total agreement with the classical case. We now possess tools that extend the behavior of the harmonic oscillator to the quantum realm. And on simple examples, they behave as one should expect. Therefore, throughout this document, to characterize the dynamics of a given quantum system, we simply provide the Hamiltonian $\mathbf{H}$ along with a collection of jump operators $\{\boldsymbol{L}_1, \ldots, \boldsymbol{L}_N\}$.

As a final introductive example, let us explore the effects of applying a resonant drive to the oscillator. In the rotating frame, this scenario can be modeled by adding a Hamiltonian term of the form $i\varepsilon(\mathbf{a} - \mathbf{a}^\dagger)$. Thus, we write:

$$\begin{aligned} \mathbf{H}/\hbar &= i\varepsilon\mathbf{a} + \text{h.c.} \\ \boldsymbol{L} &= \sqrt{\kappa_a}\mathbf{a} \end{aligned} \tag{2.15}$$

Observing the equation of motion in the Schrödinger picture, it becomes evident that the system possesses a unique steady state, which is a coherent state of amplitude of $\alpha = \varepsilon/2\kappa_a$. This outcome should not be surprising, since the situation is really the one of a child sitting on a swing, oscillating his legs back and forth at the swing's natural frequency. Looking back on our previous discussion about bistable systems sparks two intriguing questions: The drive depicted in Fig. 1.9 manifest two stable trajectories. How dose this translates in our formalism ? And since these trajectories are stable thanks to dissipation, can we devise a jump operator that conceals the system's chosen asymptotic trajectory from the environment ? We try to answer these questions, producing a way to create a physical qubit out of a harmonic oscillator.

2.1 The dissipative cat qubit

The Hilbert space of a quantum harmonic oscillator is vast, and information can be stored in more than one way. A highly efficient scheme, robust against both bit-flips and phase-flips at the hardware level, is embodied by the GKP code [Grimsmo and Puri 2021]. In the domain of trapped ion qubits, the preparation and measurement of GKP states have already been demonstrated [Flühmann et al. 2019]. In this section, we focus on a different approach and explain how coherent cat states can be used to store quantum information, and how they can be stabilized in a quantum harmonic oscillator. We abstract away from the physical realization of this resonating mode. For now on, to us, it is simply defined by a set of creation and annihilation operators $\mathbf{a}^\dagger$ and $\mathbf{a}$. We refer to it as the memory, since its purpose is to host our physical qubit. Later, we will explore methods for its concrete realisation.

As the reader will notice, the term "cat" will appear very frequently, although we are not referring to any feline creature. In our context, a cat denotes the quantum superposition of two coherent states. These states can have arbitrary complex amplitudes, but in practice, we often choose them to have identical magnitudes and opposite signs. Assuming they share the same quantum phase, such a superposition takes the form:

$$|\psi\rangle = \frac{1}{\mathcal{N}}\left(|\alpha\rangle + |-\alpha\rangle\right), \quad \mathcal{N} = \sqrt{2}\sqrt{1+\langle -\alpha|\alpha\rangle} \tag{2.16}$$

In the Wigner representation, this state is depicted by two Gaussians, centered at α and $-\alpha$, each with half the amplitude of a single coherent state. At the center of the picture, an interference pattern emerges (Fig. 2.4a). In this region, the Wigner function assumes negative values, a smoking gun evidence that superposition is quantum in nature. Loosely speaking, the system exists in both $|\alpha\rangle$ and $|-\alpha\rangle$ simultaneously. This is distinct from a situation where the system is definitively in one of the two states, but the observer is unaware of which one. In the density matrix formalism, such a state reads as:

$$\boldsymbol{\rho} = \frac{1}{\mathcal{N}^2}\left(|\alpha\rangle\langle\alpha| + |-\alpha\rangle\langle-\alpha|\right) \tag{2.17}$$

Such a state is said to be a decoherent mixture. In the Wigner representation, we still observe the presence of two Gaussians, but there is no discernible interference pattern (Fig. 2.4b). I like to say that a healthy cat has whiskers! In this work, we aim to present them in their best light.

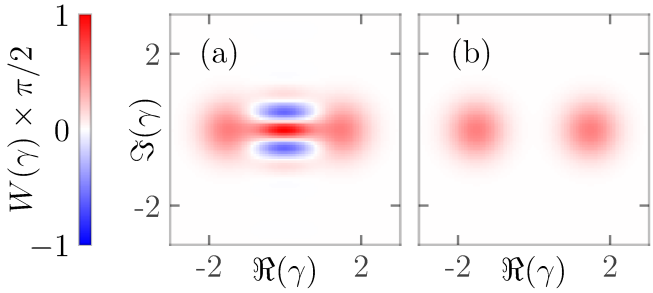

Figure 2.4 (a) Wigner representation of a coherent cat state of amplitude $\alpha = 1.5$. x and y axis are phase space coordinates, and color represents $W(\gamma)$. An interference pattern lies at the center of the image. The Wigner takes on negative values as indicated by the blue color. Since the central pixel is red, the parity of the state is 1. (b) Decoherent mixture of $|\alpha\rangle$ and $|-\alpha\rangle$. The absence of interference indicates the state is statistical.

2.1.1 Single mode model

The approach we describe, central to the Alice&Bob roadmap, relies on a combination of two-photon drive and two-photon dissipation. Consequently, we colloquially refer to our system as a dissipative cat qubit. An alternative approach, based on Hamiltonian confinement, also exists [Grimm et al. 2020; Frattini et al. 2022]. However, we choose not to delve into it to maintain conciseness in our discussion. The idealized dynamics of the dissipative cat qubit is given by combining a squeezing Hamiltonian with a two-photon loss channel.

$$\begin{aligned} \mathbf{H}/\hbar &= i\varepsilon_2 {\mathbf{a}^\dagger}^2 + \text{h.c.} \\ \boldsymbol{L}_2 &= \sqrt{\kappa_2}\mathbf{a}^2 \end{aligned} \tag{2.18}$$

Comparing Eq.2.18 to Eq.2.15, the reader can observe that the only modification made is placing a square above each operator. As we will see, this modification indeed results in the stabilization of two coherent states. As a side note, it is worth mentioning that replacing the exponent 2 with an arbitrary number indeed stabilizes an arbitrary number of coherent states.

Let us now try to give a concrete feeling of what is happening in terms of flow. To obtain its expression, we perform a semiclassical approximation. First, we write the dynamic in the Heisenberg picture:

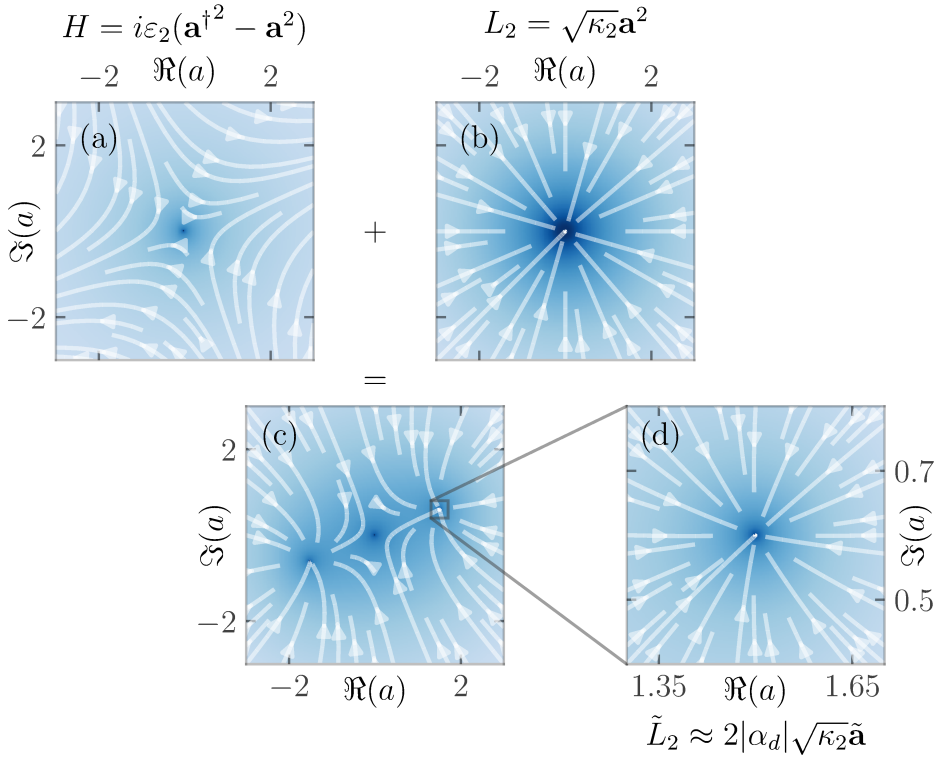

Figure 2.5 Flow imposed by the two-photon dissipative stabilization in the phase space. Color encodes for the magnitude of $\frac{da}{dt}$ in Eq. 2.20 as a function of the real part and imaginary part of a. White arrows display typical classical trajectories. Combining the squeezing Hamiltonian (a) with the two-photon dissipation (b) results in three equilibrium points (c). An unstable one lies at the center of the phase space. The other two are stable and have opposite phase. A zoom on the right stable point (d) is displayed. In the displaced frame, its effect resembles one-photon dissipation.

$$\frac{d\mathbf{a}}{dt} = 2\varepsilon_2\mathbf{a}^\dagger - \kappa_2\mathbf{a}^\dagger\mathbf{a}^2 \tag{2.19}$$

Then, replacing operators with their complex expected values $a = \langle\mathbf{a}\rangle$, and performing a mean field approximation $\langle\mathbf{a}^\dagger\mathbf{a}^2\rangle \approx a^*a^2$, one obtains the semiclassical equation of motion:

$$\frac{da}{dt} = 2\varepsilon_2 a^* - \kappa_2 a^* a^2 \tag{2.20}$$

Looking at the derivative of a in the phase space reveals a distinctive pattern. The memory is governed by two opposite forces. First, there's the squeezing Hamiltonian, which exerts forces in two opposite directions, orthogonal to the complex number ε_2, with an amplitude proportional to it (Fig. 2.5a). This part of the dynamic is totally analogous to the parametric pumping of the swing we described earlier, which "pushes" toward two trajectories with opposite phases.

The secret ingredient we didn't saw before is the two-photon jump operator $\boldsymbol{L}_2$, which removes pairs of photons at a rate κ_2. In a sense, this dissipation is stronger than the single-photon loss operator from Eq. 2.15, which wouldn't suffice to counteract the squeezing Hamiltonian — a necessary condition for the system to stabilize.

Two stable points emerge (Fig. 2.5c), with their separation increasing as the squeezing amplitude grows. In fact, it is clear from Eq. 2.19 that a coherent state of amplitude $\alpha = \pm\sqrt{2\varepsilon_2/\kappa_2}$ is an equilibrium point. To prove it is stable, we use the fact the Hamiltonian Eq. 2.18 can be factorized in the jump operator.

$$\begin{aligned} \mathbf{H}/\hbar &= 0 \\ \boldsymbol{L}_2 &= \sqrt{\kappa_2}(\mathbf{a}^2 - \alpha_d^2) \end{aligned} \tag{2.21}$$

Here, the two-photon drive $\alpha_d^2 = 2\varepsilon_2/\kappa_2$ is given in terms of photon number. Currently, it coincides with the number of photons in the steady state α^2, but later these quantities will differ. By going in the displaced frame centered around $\tilde{\mathbf{a}} = \mathbf{a} \mp \alpha_d$, the displaced two-photon jump operator $\tilde{\boldsymbol{L}}_2$ has the following expression:

$$\tilde{\boldsymbol{L}}_2 = \sqrt{\kappa_2}((\tilde{\mathbf{a}} \pm \alpha_d)^2 - \alpha_d^2) = \sqrt{\kappa_2}(\tilde{\mathbf{a}}^2 \pm 2\alpha_d\tilde{\mathbf{a}}) \tag{2.22}$$

Taking the perturbative limit where we consider a coherent state close to the one of the two stable states ($\langle\tilde{\mathbf{a}}\rangle \ll \alpha_d$), the term $\tilde{\mathbf{a}}^2$ can be neglected,

and we get that the two-photon jump operator confines the system around $\pm\alpha$ at a rate $\kappa_{\text{conf}} = 4|\alpha_d^2|\kappa_2$. This observation is stressed with the inset (d) of Fig. 2.5.

One of the fundamental aspects that makes the platform viable for quantum information storage is that, under pure two-photon stabilization, any coherent or decoherent combination of $|\alpha\rangle$ and $|-\alpha\rangle$ is stable. The system exhibits a manifold of steady states, of which, with some effort, it is possible to provide an explicit parameterization in terms of density matrix.

$$\boldsymbol{\rho} = \frac{1}{2 + 2X\langle\text{-}\alpha|\alpha\rangle}\begin{bmatrix}|\alpha\rangle & |\text{-}\alpha\rangle\end{bmatrix}\begin{bmatrix}1+Z & X-iY\\ X+iY & 1-Z\end{bmatrix}\begin{bmatrix}\langle\alpha|\\ \langle\text{-}\alpha|\end{bmatrix} \tag{2.23}$$

Here $(X, Y, Z) \in \mathbb{R}$ are the state coordinates. The positivity condition enforces that $X^2 + Y^2 + Z^2 \leqslant 1$, which effectively depicts the steady state manifold resembling a Bloch sphere. However, in this case, the density matrix formalism allows to write decoherent states, which reside within its interior. The amplitude of the two-photon drive defines what we call a code space: $|+\alpha\rangle$ plays the role of the logical zero $|0\rangle_\alpha$, while $|-\alpha\rangle$ is the logical one $|1\rangle_\alpha$. The cat plus and minus states are denoted as $|+\rangle_\alpha$ and $|-\rangle_\alpha$, respectively. The outer ring of Fig. 2.6 depicts the cat-qubit code space for a two-photon drive of amplitude $\alpha_d^2 = 2$.

For small value of ε_2, according to the convention of Eq. 2.23, the minus cat state is ill-defined. We can obtain its representation by taking $\alpha \to 0$, which is displayed in the inner ring of Fig. 2.6. Remarkably, it corresponds to the first Fock state $|-\rangle_0 = |1\rangle$, whereas the plus cat is represented by the vacuum state $|+\rangle_0 = |0\rangle$.

A key insight is to observe that this dynamics conserves the parity of the number of photons. Looking at Eq.2.19, they are injected or removed in pairs. An alternative way to comprehend the preservation of quantum superpositions is to recognize that the plus and minus cat states exclusively support the even and odd Fock states, respectively. Therefore, stabilizing a cat state requires maintaining the parity of the number of photons within the memory. In fact, one can prove the X coordinate on the Bloch sphere aligns precisely with the average value of the parity operator $\langle e^{i\pi\mathbf{a}^\dagger\mathbf{a}}\rangle$, which itself is equal to $\frac{\pi}{2}W(0)$. The X coordinate in the code space is equal to the color of central pixel of the Wigner.

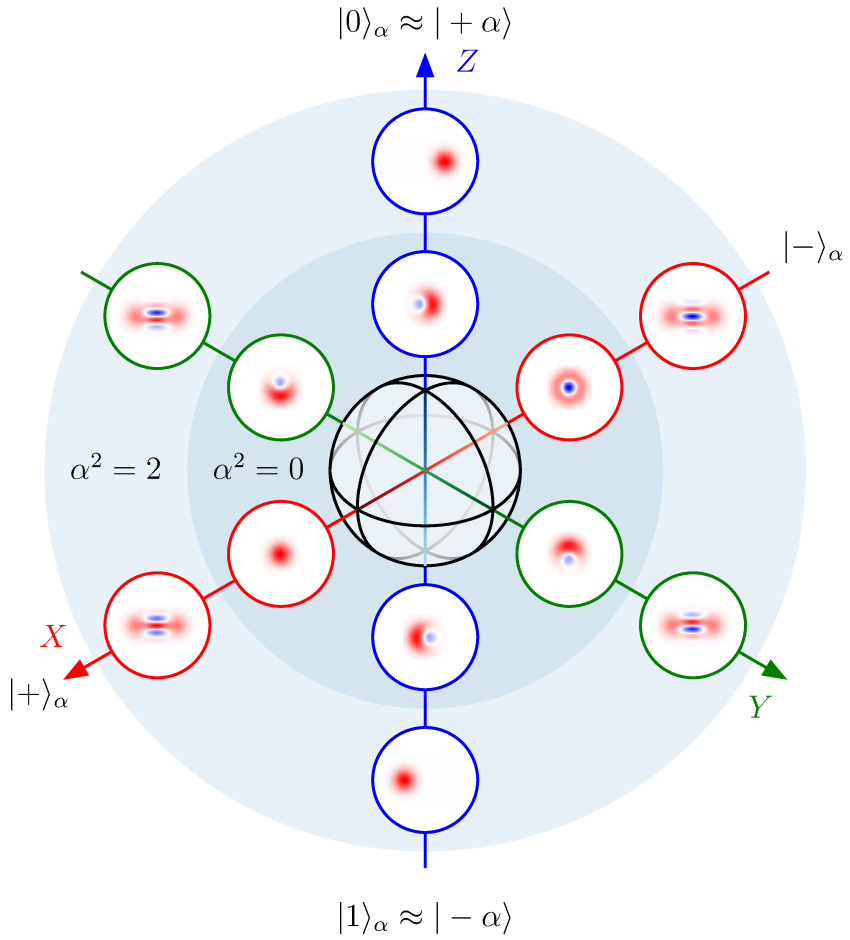

Figure 2.6 Bloch sphere representation of the cat qubit code space. The light blue crown illustrates the Wigner function of various logical states for a code space amplitude of $\alpha_d^2 = 2$. The inner dark blue halo depicts the logical states in the absence of a two-photon drive, representing a code space amplitude of zero.

2.1.2 Dissipative second order phase transition

In practice, the memory inevitably exhibits some imperfections. Therefore, it is crucial to investigate the behavior of two-photon stabilization in the presence of parasitic factors. Let us reintroduce a one-photon dissipation channel $\boldsymbol{L}_a$, featuring the damping rate κ_a.

$$\begin{aligned} \mathbf{H}/\hbar &= 0 \\ \boldsymbol{L}_2 &= \sqrt{\kappa_2}(\mathbf{a}^2 - \alpha_d^2) \\ \boldsymbol{L}_a &= \sqrt{\kappa_a}\mathbf{a} \end{aligned} \tag{2.24}$$

Before examining how the storage of quantum information is impacted by this adverse effect, it is insightful to investigate the behavior of simple and concrete observables. Specifically, the number of photons in the steady state is a quantity that is generally straightforward to access from an experimental perspective. In the subsequent section, we capitalize on the notion that a parallel can be drawn between our system and the thermodynamic limit of certain other models, an idea developed in [Minganti 2018, p. 69] to understand the behavior of the number of photons in the steady-state [Mylnikov et al. 2022].

In the presence of one-photon losses, the manifold of stable states collapses to a unique state. For a large two-photon drive, simulations reveal it looks like the decoherent mixture of two coherent states Fig. 2.7III. We can approximate its number of photons by accounting for dissipation in the semiclassical model Eq. 2.20. Setting the time derivative to zero yields the following system:

$$0 = \kappa_2 \alpha_d^2 a^* - \kappa_2 a^* a^2 - \frac{\kappa_a}{2} a \tag{2.25}$$

This model predicts that the vacuum state $a = 0$ is always a solution, although it becomes unstable when the two-photon drive is strong enough. Formalizing this observation, we obtain a formula for the number of photons in the stable steady state for the semiclassical model:

$$|a|^2 = \max\left(0, \alpha_d^2 - \frac{\kappa_a}{2\kappa_2}\right) \tag{2.26}$$

In a manner akin to the magnetization of the Ising model at finite temperature under an external magnetic field, the mean field approximation applied to our model results in a discontinuity in the derivative of $|a|^2$ with respect to α_d^2 at the critical point $\alpha_d^2 = \kappa_a/2\kappa_2$. Interestingly, the value of $\langle \mathbf{a}^\dagger \mathbf{a} \rangle$

increasingly coincides with $|a|^2$ as we either augment the two-photon drive α_d, the one-photon loss rate κ_a, or both.

To better expose this phenomenon, we renormalise α_d^2 and $\langle \mathbf{a}^\dagger \mathbf{a} \rangle$ by $\kappa_a/2\kappa_2$. As this ratio approaches infinity, we enter what we call the strong dissipative limit. The critical point happens at larger and larger two-photon drives. Looking at the renormalised curve Fig. 2.7, quantum fluctuations get washed out, and the curve of $\langle \mathbf{a}^\dagger \mathbf{a} \rangle$ approaches the solution of the semiclassical model, which slope exhibit an abrupt change. Formally, we have:

$$\lim_{\kappa_a/\kappa_2 \to \infty} \frac{2\kappa_2}{\kappa_a} \langle \mathbf{a}^\dagger \mathbf{a} \rangle |_{\alpha_d^2 = \frac{x\kappa_a}{2\kappa_2}} = \max(0, x-1) \tag{2.27}$$

In a sense, for this bosonic system, the strong dissipative regime is the same as considering a thermodynamic limit [Minganti 2018, p. 69]. We say there is a second-order dissipative phase transition at the critical point $\alpha_d^2 = \kappa_a/2\kappa_2$ [Beaulieu et al. 2023]. However, in practical terms, our intent is to operate this system with only a limited number of photons in the memory, specifically for small values of κ_a. In this case, the phase transition near the critical point is tempered by quantum fluctuations. To delineate the system behavior in its vicinity, we have to perform a higher order mean field approximation. We derive the equations of motion for $\mathbf{a}^\dagger \mathbf{a}$ and $\mathbf{a}^2$ in the Heisenberg picture:

$$\begin{aligned} \frac{d\mathbf{a}^2}{dt} &= \kappa_2 (2\mathbf{a}^\dagger \mathbf{a} + 1)(\alpha_d^2 - \mathbf{a}^2) - \kappa_a \mathbf{a}^2 \\ \frac{d\mathbf{a}^\dagger \mathbf{a}}{dt} &= \kappa_2 (\alpha_d^2 \mathbf{a}^{\dagger 2} + \alpha_d^{*2} \mathbf{a}^2 - 2\mathbf{a}^{\dagger 2} \mathbf{a}^2) - \kappa_a \mathbf{a}^\dagger \mathbf{a} \end{aligned} \tag{2.28}$$

Drawing inspiration from [Mylnikov et al. 2022], we introduce the second order moments $n = \langle \mathbf{a}^\dagger \mathbf{a} \rangle$ and $\psi = \langle \mathbf{a}^2 \rangle$, considered real for convenience. We replace higher-order operator expectations in Eq. 2.28 by the product of lower-order operator expectations:

$$\begin{aligned} \langle \mathbf{a}^\dagger \mathbf{a}^3 \rangle &\approx n\psi \\ \langle \mathbf{a}^{\dagger 2} \mathbf{a}^2 \rangle &\approx \psi^2 \end{aligned} \tag{2.29}$$

Substituting back this approximation into Eq. 2.28 and solving for the steady state yields a degree three polynomial system.

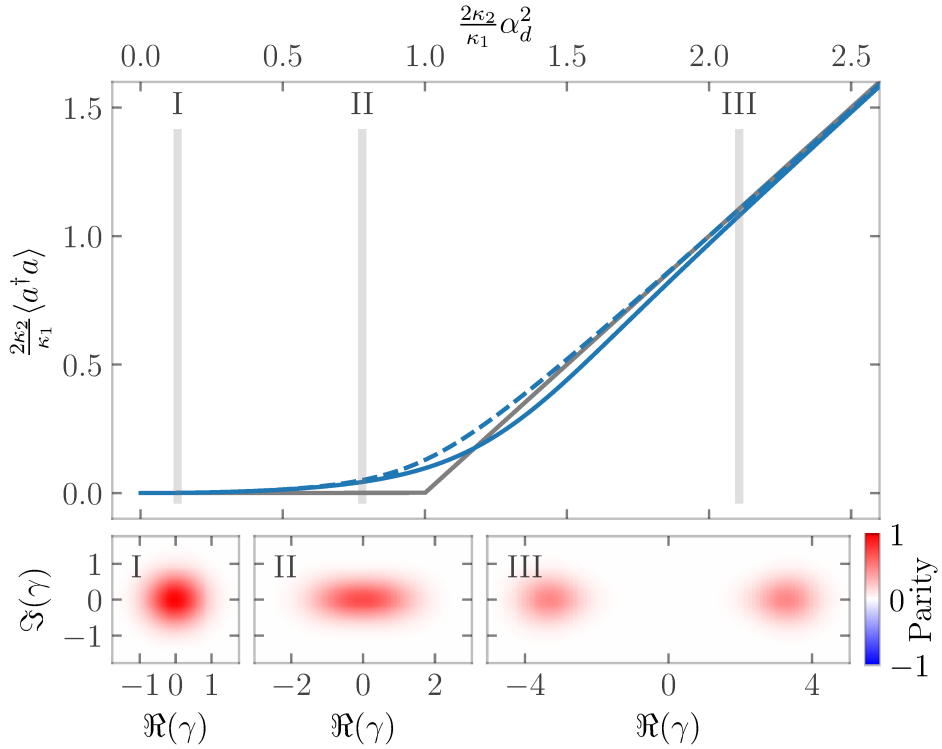

Figure 2.7 Second order dissipative phase transition. Normalized photon number (top) plotted against the normalized two-photon drive. A numerical simulation of the steady state (solid blue line) in the regime $\kappa_a/\kappa_2 = 20$ is compared to the thermodynamic limit (solid gray line) and the equation-based approximation from Eq. 2.30 (dashed blue line). At zero drive, vacuum is the only stable state (bottom-left). As the two-photon injection rate intensifies, squeezing is witnessed around the critical point (bottom-center). Going further to the right hand side of the diagram, a statistical mixture of two coherent states is evicted from the initial vacuum.

$$\begin{aligned}\psi &= \frac{\kappa_2}{\kappa_a}(2n+1)(\alpha_d^2-\psi)\\ n &= 2\frac{\kappa_2}{\kappa_a}(\alpha_d^2-\psi)\psi\end{aligned} \tag{2.30}$$

For a specific two-photon drive value, these equations admit a single stable solution displayed Fig. 2.7. The photon number now increases smoothly at the critical point. In appropriate units, this smoothness encodes quantum fluctuations' amplitude near the phase transition. To demonstrate this property, we temporarily assume the position of an experimentalist who just measured this curve on an actual system. He is not certain of his instruments, and don't know the units of the x and y-axis.

By fitting the data with a straight line at large two-photon drive, it is always possible to find the position of the critical point. In normalized units, the x-coordinate of this point is equal to 1 (Fig .2.7). Dividing the y-axis by the asymptotic slope, the two photon drive and the photon number are now in homogenous units. The key observation is to realize the polynomial curve of Eq. 2.28 we obtained is not scale invariant under the transformation:

$$\begin{aligned}&\frac{2\kappa_2}{\kappa_a}\langle \mathbf{a}^\dagger \mathbf{a}\rangle \to \lambda\frac{2\kappa_2}{\kappa_a}\langle \mathbf{a}^\dagger \mathbf{a}\rangle, \quad \lambda \in \mathbb{R}_+^*\\ &\frac{2\kappa_2}{\kappa_a}\alpha_d^2 \to \lambda\frac{2\kappa_2}{\kappa_a}\alpha_d^2\\ &\frac{\kappa_a}{2\kappa_2} \to \lambda\frac{\kappa_a}{2\kappa_2}\end{aligned} \tag{2.31}$$

Where λ can be seen as the scaling of the two-photon rate compared to the one-photon losses. Therefore, it is possible to find the unique ratio $\kappa_a/2\kappa_2$ that is compatible with the data. Combined with a prior knowledge of κ_a, both the photon number and the two-photon dissipation rate κ_2 can be calibrated. For lack of a better word, this provides us with a sense of scale. Notably, we employed this technique to calibrate the experiment detailed in [Berdou et al. 2023].

2.1.3 Scaling of the bit-flip and phase-flip rates

As we said, in the presence of parasitic one-photon dissipation, the manifold of steady states condenses into a single decoherent mixture. It indicates that, in the long run, quantum information eventually vanishes. The qubit

inherits finite lifetimes from its defects. However, we will observe that these lifetimes exhibit a hierarchical pattern for the cat qubit, and can differ by orders of magnitudes, thereby giving rise to a biased noise qubit.

Let's once again adopt the perspective of an experimentalist. In this scenario, the we aim to assess the quality of a given cat qubit. To achieve this, we prepare the state $|+\rangle_\alpha$. By definition, the speed at which the system transitions to a decoherent mixture, represented as $1/T_{\text{pf}}$, aligns with the phase-flip rate. Conversely, when starting from a logical zero state $|0\rangle_\alpha$, the typical time T_{bf} required to reach the steady state corresponds to the bit-flip time. By computing the time evolution in these two cases, we can compute T_{bf} and T_{pf}, and in particularly look at their theoretical scaling as α^2 grows.

On the one hand, the phase-flip rate is straightforward to compute from Eq. 2.24. It's determined by writing the Lindblad master equation for the parity operator in the Heisenberg picture, employing the Baker-Haudorff formula to commute terms.

$$\frac{de^{i\pi\mathbf{a}^\dagger\mathbf{a}}}{dt} = -2\kappa_a e^{i\pi\mathbf{a}^\dagger\mathbf{a}}\mathbf{a}^\dagger\mathbf{a} \tag{2.32}$$

Now, we can replace operators with their expectation values. When $\kappa_a \ll \kappa_2$, we have $\langle\mathbf{a}^\dagger\mathbf{a}\rangle \approx \alpha_d^2$. Also, as previously discussed, the value of the parity operator $\langle e^{i\pi\mathbf{a}^\dagger\mathbf{a}}\rangle$ is proportional to the X coordinate in the cat qubit code space. Altogether, this yields a first-order differential equation that results in exponential decay:

$$\begin{aligned}\frac{dX}{dt}(t) &\approx -2\kappa_a X(t)\alpha_d^2 \\ X(t) &= X(0)e^{-t/T_{\text{pf}}}\end{aligned} \tag{2.33}$$

A closed formula for the phase-flip time is therefore found:

$$T_{\text{pf}} \approx 1/2\kappa_a\alpha_d^2 \tag{2.34}$$

Importantly, it scales linearly with the two-photon drive amplitude, which almost equals the photon number in the steady state $\alpha_d^2 \approx \alpha^2$.

On the other hand, determining the scaling behavior of the bit-flip rate poses a greater challenge. We propose a method applicable under the condition where the one-photon loss rate is significantly smaller than the two-photon loss rate $\kappa_a \ll \kappa_2$.

We start by reconsidering the action of $\boldsymbol{L}_2$ the pure two-photon stabilisation in Eq. 2.24. Initializing the system with a coherent state $|\sigma\rangle$ of

amplitude σ, one should be able to compute the code space coordinates $X(\alpha_d, \sigma), Y(\alpha_d, \sigma), Z(\alpha_d, \sigma)$ towards which the system will eventually converge. In fact, this calculation is done in [Guillaud et al. 2023]. In our notations, we get:

$$\begin{aligned} X(\alpha_d, \sigma) &= e^{-2|\sigma|^2} \\ Z(\alpha_d, \sigma) - iY(\alpha_d, \sigma) &= \frac{2i\alpha_d \sigma^* e^{-i|\sigma|^2}}{\sqrt{2\sinh(2|\alpha_d|^2)}} I(\alpha_d, \sigma) \\ I(\alpha_d, \sigma) &= \int_0^{\phi} e^{-i\phi} I_0\left(|\alpha_d^2 - \sigma^2 e^{2i\phi}|\right) d\phi \end{aligned} \tag{2.35}$$

Formally, we say these 3 real numbers are left invariant by the dynamics imposed by Eq. 2.21. Now, taking inspiration from [Guillaud et al. 2023], let us give the hand waving interpretation of Eq. 2.24 that is the key to understand our argument: Loosely speaking, we can say the two-photon jump operator $\boldsymbol{L}_2$ projects the system into the cat-qubit code space at an average rate determined by κ_{conf}. In the interval between these successive projections, the system undergoes the decoherence induced by $\boldsymbol{L}_a$. Starting from the logical zero state $|0\rangle_\alpha \approx |\alpha\rangle$, it eventually settles into a coherent state of amplitude $\alpha_{\text{leak}} = e^{-\kappa_a/4\kappa_2\alpha_d^2}\alpha$ before being projected back into the code space. The projection on the code space conserves the Z coordinate, unlike one-photon dissipation. As this process gets repeated again and again, the Z coordinate exponentially converges to 0 at a rate that we can compute. Following this intuition, we obtain:

$$T_{\text{bf}} \approx \frac{1}{2\kappa_2\alpha_d^2} \frac{Z(\alpha_d, \alpha)}{Z(\alpha_d, \alpha) - Z(\alpha_d, \alpha_{\text{leak}})} \tag{2.36}$$

The values of both α and α_{leak} being in the vicinity of α_d, a perturbative expansion of the integral $I(\alpha_d, \alpha_d - \eta)$ around $\eta = 0$ enables us to derive a perturbative closed formula for $Z(\alpha_d, \alpha_d - \eta)$:

$$
\begin{aligned}
I(\alpha_d, \alpha_d - \eta) &\underset{\kappa_a/\kappa_2 \to 0}{=} I(\alpha_d, \alpha_d) + \eta \frac{dI(\alpha_d, \alpha_d - \eta)}{d\eta}|_{\eta=0} \\
&+ \frac{\eta^2}{2} \frac{d^2 I(\alpha_d, \alpha_d - \eta)}{d\eta^2}|_{\eta=0} + o(\eta^2)
\end{aligned}
$$

$$
\begin{aligned}
I(\alpha_d, \alpha_d) &= -2i \int_0^1 \frac{y}{\sqrt{1-y^2}} I_0(2\alpha_d^2 y) dy \\
&= -\frac{i \sinh(2\alpha_d^2)}{\alpha_d^2}
\end{aligned}
$$

$$
\begin{aligned}
\frac{dI(\alpha_d, \alpha_d - \eta)}{d\eta}|_{\eta=0} &= 4i\alpha_d \int_0^1 \frac{y^2}{\sqrt{1-y^2}} I_1(2\alpha_d^2 y) dy \\
&= i \frac{2\alpha_d^2 \cosh(2\alpha_d^2) - \sinh(2\alpha_d^2)}{\alpha_d^3}
\end{aligned}
$$

$$
\begin{aligned}
\frac{d^2 I(\alpha_d, \alpha_d - \eta)}{d\eta^2}|_{\eta=0} &= -8i\alpha_d^2 \int_0^1 \frac{y^3}{\sqrt{1-y^2}} I_2(2\alpha_d^2 y) dy \\
&+ 4i \int_0^1 \frac{y^2 - 2}{\sqrt{1-y^2}} I_1(2\alpha_d^2 y) dy \\
&= 4i \frac{\alpha_d^2 (1 + \cosh(2\alpha_d^2)) - (1 + \alpha_d^2) \sinh(2\alpha_d^2)}{\alpha_d^4}
\end{aligned}
\tag{2.37}
$$

Incorporating both Eq. 2.37 and Eq. 2.35 in Eq. 2.36 yields an asymptotic closed expression for the bit-flip rate in the limit $\kappa_a/\kappa_2 \to 0$. While the initial expression may seem intractable, employing formal calculus algorithms makes it possible to recover some known results. Particularly, we can compute the asymptotic scaling of the bit-flip rate.

$$
\lim_{\kappa_a/\kappa_2 \to 0} \left(\lim_{\alpha_d^2 \to \infty} \frac{\partial \log \left(T_{\text{bf}}(\alpha_d, \kappa_a, \kappa_2) \right)}{\partial \alpha_d^2} \right) = 4 \tag{2.38}
$$

We find that at large photon number, when the one-photon loss rate is small compared to the two-photon stabilisation, the bit-flip rate scales exponentially as $e^{4\alpha_d^2}$. This implies that for every photon introduced into the cat state, the bit-flip time increases by a factor of $e^4 \approx 54.6$. By considering typical numbers for superconducting qubit technology, one can see that

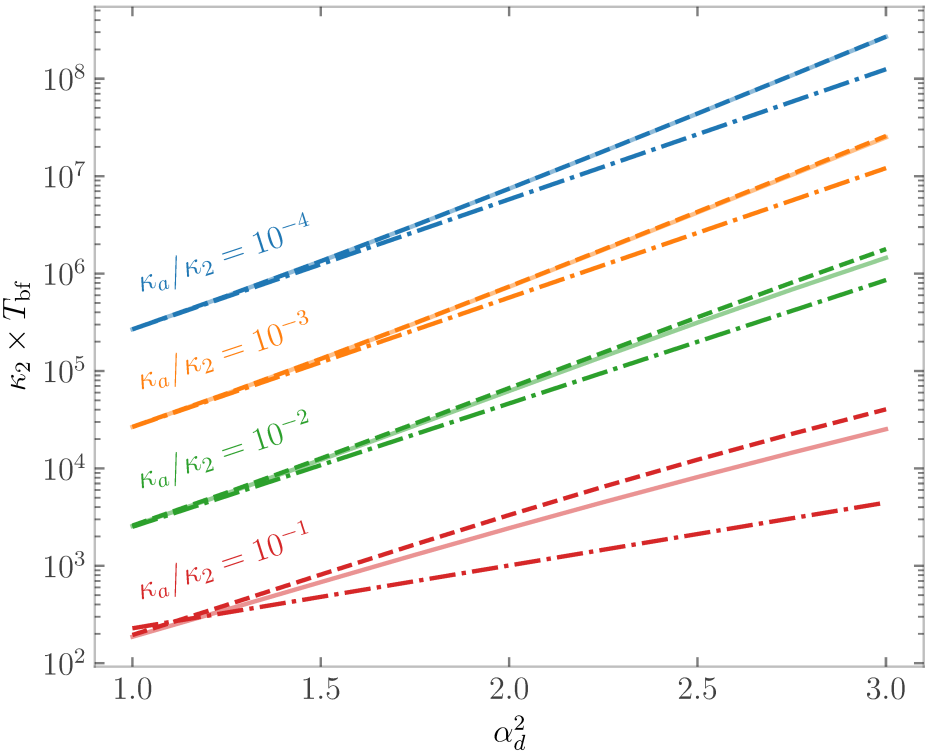

Figure 2.8 Bit-flip time scaling (y-axis) in the single-mode dissipative cat qubit model with respect to the two-photon drive amplitude. The numerical diagonalization of the Liouvillian (solid transparent lines) is contrasted with the projection from Eq.2.36 (dashed lines) and the series approximation of Eq.2.39 (dotted dash lines). Remarkably accurate predictions are achieved for $\kappa_2/\kappa_a \leqslant 10^{-1}$ using the analytical approach.

regimes where T_{bf} exceeds T_{pf} by more than 5 orders of magnitude are in reach, while still allowing for coherent quantum gate operations, at least on paper. This technique is therefore an excellent candidate for implementing a biased noise qubit.

As a side note, on closer examination of equation Eq.2.36 we find that the exponential scaling rate of the bit-flip time is not constant and varies with α_d. Since we intend to operate this system at a low photon count, we perform a series expansion around $\alpha_d = 1$ of T_{bf}, giving us an approximate formula for the bit-flip time.

$$
\begin{aligned}
T_{\text{bf}} \underset{\kappa_a/\kappa_2 \to 0}{=} & \left(a_0 \frac{1}{\kappa_a} + a_1 \frac{1}{\kappa_2} + a_2 \frac{\kappa_a}{\kappa_2^2} + o\left(\frac{\kappa_a}{\kappa_2}\right) \right) \times \\
& \exp\left((b_0 + b_1 \frac{\kappa_a}{\kappa_2})(\alpha_d^2 - 1) + o\left(\frac{\kappa_a}{\kappa_2}\right) \right)
\end{aligned}
$$

$$
\begin{aligned}
a_0 &= \frac{e^4 + 1}{2} \\
&\approx 26.8
\end{aligned}
$$

$$
\begin{aligned}
a_1 &= \frac{1 + 6e^2 - 9e^4 - 6e^6}{16} \\
&\approx -179.2
\end{aligned}
$$

$$
\begin{aligned}
a_2 &= \frac{(1 + e^2)(-237 + 33e^2 + 164e^4 + 136e^6)}{384} \\
&\approx 1394.4
\end{aligned}
$$

(2.39)

$$
\begin{aligned}
b_0 &= \frac{1 + 3e^4}{-1 + e^4} \\
&\approx 3.1
\end{aligned}
$$

$$
\begin{aligned}
b_1 &= \frac{-3 - 12e^2 + 67e^4 + 24e^6 + 3e^8 - 12e^{10} - 3e^{12}}{16(-1 + e^4)^2} \\
&\approx -15.9
\end{aligned}
$$

This formula demonstrates an exceptional correspondence with simulations, particularly noticeable for $\kappa_a < 0.1$ (Fig. 2.8). Discrepancies arise

when we enter the strongly dissipative regime since since the steady state cannot be approximated as a statistical mixture of coherent states. Within the experimental scope, ranging from one to ten photons, we find that the scaling of the bit-flip time varies between approximately $e^{3.1} \approx 22.2$ and $e^{3.8} \approx 44.7$. In the limit $\kappa_a \ll \kappa_2$ near one photon, the bit-flip time amounts to $(e^4+1)/2\kappa_a$, which can serve as a nice benchmark to validate exponential fit on experimental data.

2.1.4 Two photon dissipation with reservoir engineering

Pure two-photon exchange, involving the interaction of a mode with its environment, does not naturally occur. To implement this mechanism, we introduce an additional mode, referred to as the waste mode or buffer. This buffer is distinctly designed to exhibit a low quality factor, in contrast with the memory mode. Photons entering the buffer are rapidly released into the thermal bath at a rate denoted by κ_b. This leaky reservoir **b** is then intricately coupled to the memory mode **a** ensuring that photons in the buffer are exchanged for pairs of photons in the memory. The frequency of this exchange is defined by the two-photon exchange rate g_2. The situation is therefore the following:

$$\begin{aligned} \mathbf{H}/\hbar &= g_2 \mathbf{a}^{\dagger 2}\mathbf{b} - \varepsilon_d \mathbf{b}^\dagger + \text{h.c.} \\ \boldsymbol{L}_b &= \sqrt{\kappa_b}\mathbf{b} \\ \boldsymbol{L}_a &= \sqrt{\kappa_a}\mathbf{a} \end{aligned} \tag{2.40}$$

Since the operator responsible for the conversion between photons of the memory and those of the buffer acts on three photons at a time (two from the memory and one from the buffer), it is referred to as a three-wave mixer. Now, we will assume that the buffer degrees of freedom reach equilibrium much faster than the others. Under this condition, it becomes possible to exclude **b** from the motion equations by setting $\frac{d\mathbf{b}}{dt} = 0$:

$$\mathbf{b} = i\frac{\varepsilon_d}{2\kappa_b}\boldsymbol{I} \tag{2.41}$$

Notice that after this treatment, **b** is still an operator, indicated by the presence of the identity element $\boldsymbol{I}$. From there, we can write an effective theory solely for **a**. Among physicists, this technique is known as adiabatic elimination. If you have a background in automatic systems, you might see that it is in fact an application of Tikhonov's theorem [Petit and Rouchon 2009]:

$$\begin{aligned}\mathbf{H}/\hbar &= i\frac{2\varepsilon_d g_2}{\kappa_b}\mathbf{a}^{\dagger 2} + \text{h.c.}\\ \boldsymbol{L}_2 &= \sqrt{4|g_2|^2/\kappa_b}\mathbf{a}^2\\ \boldsymbol{L}_a &= \sqrt{\kappa_a}\mathbf{a}\end{aligned} \tag{2.42}$$

We get back the single mode model of the dissipative cat qubit, and we can identify variables:

$$\kappa_2 = \frac{4|g_2|^2}{\kappa_b}, \quad \kappa_a = \kappa_a, \quad \varepsilon_2 = \frac{2\varepsilon_d g_2}{\kappa_b}, \quad \alpha_d^2 = \frac{\varepsilon_d}{g_2} \tag{2.43}$$

In this simplified model, the buffer drive is effectively converted into squeezing, injecting pairs of photons into the memory, and the leakiness of the buffer results in effective two-photon dissipation. As before, these two processes reach equilibrium, leading to the stabilization of a cat state inside the memory mode. However, for this description to be accurate, the photon pairs escaping from the memory to the buffer must be leaked into the thermal bath before they have a chance to be converted back into the memory mode. Under this assumption, adiabatic elimination is valid. Intuitively, this implies that the two-photon exchange rate g_2 should be small enough compared to κ_b. Let us derive the precise condition these quantities must satisfy.

To further analyze and quantify our model, we derive the semi-classical equations of motion for $a = \langle \mathbf{a} \rangle$ and $b = \langle \mathbf{b} \rangle$ using a mean field approximation, ignoring memory losses:

$$\begin{aligned}\frac{da}{dt} &= -2ig_2 a^* b\\ \frac{db}{dt} &= -\frac{\kappa_b}{2}b - ig_2^* a^2 + i\varepsilon_d\end{aligned} \tag{2.44}$$

Again, the system displays an unstable solution where the memory is in the vacuum $a = 0$ and the buffer gets displaced $b = i\varepsilon_d/2\kappa_b$, and two stable solutions where the buffer $b = 0$ is empty and the amplitude of the field in the memory is given by $a = \pm\alpha_d$.

Let's delve into the stability analysis, and explicit the validity of the adiabatic elimination. To this end, we express the semi-classical dynamics within the displaced frame centered around the steady state. We introduce $a = \alpha_d + \delta a$ and $b = 0 + \delta b$ to examine the dynamics around this point, using a first-order approximation:

$$\begin{aligned} \frac{d\delta a}{dt} &= -2ig_2 a_d^* \delta b \\ \frac{d\delta b}{dt} &= -2ig_2^* \delta_{\mathrm{a}} - \frac{\kappa_b}{2}\delta b \end{aligned} \tag{2.45}$$

These differential equations can be casted in a more condensed form using matrix formalism. For simplicity sake, we assume g_2 to be real.

$$\begin{aligned} \frac{dX}{dt} &= \begin{bmatrix} A & 0_{2\times 2} \\ 0_{2\times 2} & A \end{bmatrix} X \\ A &= \begin{bmatrix} 0 & 2g_2\alpha_d \\ -2g_2\alpha_d & -\kappa_b/2 \end{bmatrix}, X = \begin{bmatrix} \Re(\delta a) \\ \Im(\delta b) \\ \Im(\delta a) \\ \Re(\delta b) \end{bmatrix} \end{aligned} \tag{2.46}$$

Presented in this manner, it becomes clear that the rate of convergence of the dynamics can be deduced from diagonalizing the matrix A. The roots of its characteristic polynomial have negative real parts, as we are looking at a stable solution.

$$\begin{aligned} \det(\lambda I - A) =& \lambda^2 + \frac{\kappa_b}{2}\lambda + 4|g_2|^2\alpha_d^2 \\ \lambda_\pm = -\frac{\kappa_b}{4} \pm \frac{1}{2}\sqrt{\Delta}, &\quad \Delta = \frac{\kappa_b^2}{4}\left(1 - \left(\frac{8g_2\alpha_d}{\kappa_b}\right)^2\right) \end{aligned} \tag{2.47}$$

The confinement rate κ_{conf} is determined by the root with the smallest real part in absolute value. In the above equation, Δ represents the discriminant of the system. We find the condition under which adiabatic elimination is valid: the discriminant is positive when the two-photon coupling g_2 scaled by the two-photon drive α_d is much smaller than $\kappa_b/8$. In this case, the convergence is rapid and uniform, and we find an expression compatible with our previous derivation.

$$g_2\alpha_d \ll \kappa_b/8 \quad \Rightarrow \quad \kappa_{\mathrm{conf}} = -2\lambda_- \approx \frac{16g_2^2\alpha_d^2}{\kappa_b} = 4\kappa_2\alpha_d^2 \tag{2.48}$$

When $\alpha_d g_2 > \kappa_b/8$, the roots have imaginary parts, and the system begins to oscillate. In the displaced frame around α_d, these oscillations can

be viewed as Rabi exchanges between the memory and the buffer. In the highly non-adiabatic limit, we observe the saturation of the confinement rate.

$$g_2\alpha_d \gg \kappa_b/8 \quad \Rightarrow \quad \kappa_{\text{conf}} = -2\Re\left(\lambda_-\right) \approx \frac{\kappa_b}{2} \tag{2.49}$$

In this limit, photons are constantly hopping between the buffer and the memory, being dissipated half of the time. It comes as no surprise that the effective confinement rate is thus half the buffer's leakiness. The scaling of the bit-flip rate we derived earlier did not account for this saturation, which imposes a limitation on the protection of the bit-flip degree of freedom practical applications. As we will see later, however, experimentally, we still recover an exponential scaling of the bit-flip rate in the non-adiabatic limit, but it is smaller than $e^{4\alpha^2}$.

2.1.5 Liouvillian formulation

In our previous discussions, we successfully derived analytical expressions for T_X, T_Z, and κ_{conf} in various limiting cases using semi-analytical methods. This approach provides valuable insights into our system's behavior but has limitations when dealing with more general scenarios. In this section, we introduce a numerical method that offers a unified approach to compute these quantities, known as the Liouvillian picture [Minganti and Huybrechts 2022].

The transition to the Liouvillian picture involves casting the Lindblad master equation into a linear formalism. Operators are elevated to superoperators, and the time evolution is represented as left multiplication by the Liouvillian $\boldsymbol{\mathcal{L}}$. For instance, we can rewrite the dynamics of the single-mode model (Eq. 2.24) in this formalism:

$$\begin{aligned}
\frac{d|\boldsymbol{\rho}\rangle\rangle}{dt} &= \boldsymbol{\mathcal{L}}|\boldsymbol{\rho}\rangle\rangle \\
\boldsymbol{\mathcal{L}} &= \boldsymbol{G}\otimes\mathbf{1} + \mathbf{1}\otimes\boldsymbol{G}^* + \sum_{i\in\{a,2\}} \boldsymbol{L}_i\otimes\boldsymbol{L}_i^* \\
\boldsymbol{G} &= \frac{-i}{\hbar}\mathbf{H} - \frac{1}{2}\sum_{i\in\{a,2\}} \boldsymbol{L}_i^\dagger\boldsymbol{L}_i
\end{aligned} \tag{2.50}$$

Where $|\boldsymbol{\rho}\rangle\rangle$ can be thought as the column vector obtained by flattening the matrix $\boldsymbol{\rho}$. Essentially, we have transformed the study of the dissipative

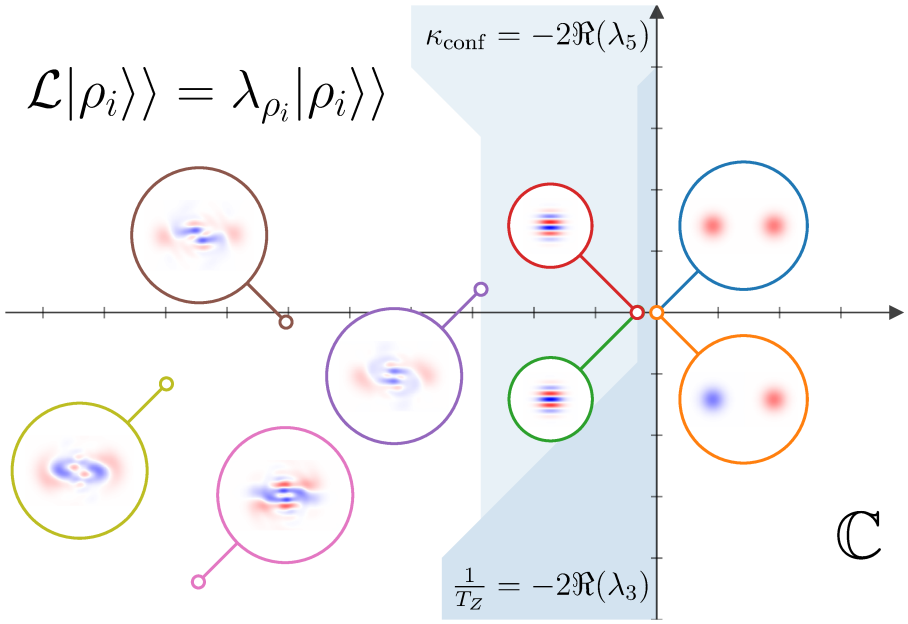

Figure 2.9 Liouvillian eigenvalues and eigenvectors in the complex plane. In the complex plane (x and y axis), we position each eigenvalue of the liouvillian from Eq.2.50 (small colored circles). The Wigner function of their associated normalized eigenvectors is displayed (large color circles). The Liouvillian is gapped between its 4th and 5th eigenvalues, defining the confinement rate κ_{conf}.

dynamics into the diagonalization of $\boldsymbol{\mathcal{L}}$. Let's denote $(\lambda_0, ..., \lambda_n, ...)$ as the eigenvalues of the Liouvillian, sorted in decreasing order of their real parts, and associate them with their respective eigenvectors $(|\tilde{\boldsymbol{\rho}}_0\rangle\rangle, ..., |\tilde{\boldsymbol{\rho}}_n\rangle\rangle, ...)$. Please note that the tildes are used to emphasize that the eigenvectors are not proper density matrices, as the diagonalization process doesn't guarantee either $\text{Tr}\,(\tilde{\boldsymbol{\rho}}) = 1$ or $\tilde{\boldsymbol{\rho}}^\dagger = \tilde{\boldsymbol{\rho}}$. We will address this technicality shortly.

After numerical diagonalization, we can plot the spectrum in the complex plane, giving us a complete view of the system's behavior Fig. 2.9. The first eigenvalue is strictly zero and corresponds to the steady state. Its eigenvector is the only one that represents a valid density matrix, and we will offer an intuitive interpretation of the remaining spectrum.

Examining the Wigner function of the second eigenvector, we observe that it appears to be associated with the differential occupation of $|0\rangle\,\alpha$ and $|1\rangle\,\alpha$, hinting at a connection with the bit-flip time. Similarly, the Wigner

functions of the third and fourth eigenvectors depict fringes, the lifetimes of which are associated with the phase-flip rate. In fact, we can rigorously define both these quantities based on the Liouvillian spectrum.

$$T_X = \frac{-1}{2\Re(\lambda_1)}, \quad T_Z = \frac{-1}{2\Re(\lambda_3)} \tag{2.51}$$

Noticeably, the first four eigenvectors of the Liouvillian provide a valid parameterization of the cat-qubit code space. Another crucial insight is the presence of a gap in the Liouvillian spectrum between its fourth and fifth eigenvalues. This gap gives a lower bound on the rate at which the system converges to the code space when initiated in an arbitrary state. This observation allows us to establish a precise definition of the confinement rate.

$$\kappa_{\text{conf}} = -2\Re(\lambda_5) \tag{2.52}$$

To make sense of the Liouvillian eigenvectors, we need to revisit the definition of a density matrix. The manifold P of admissible density matrices exhibits two key properties: it is affine ($\boldsymbol{\rho} \in P \Rightarrow \mathrm{Tr}(\boldsymbol{\rho}) = 1$) and convex ($\boldsymbol{\rho}', \boldsymbol{\rho}'' \in P^2 \Rightarrow \forall t \in [0,1], t\boldsymbol{\rho}' + (1-t)\boldsymbol{\rho}'' \in P$). Therefore, if the spectrum of the Liouvillian consisted of more than one admissible density matrix, the first property would imply that the dynamical state $\boldsymbol{\rho}(t) = \sum \exp(\lambda_n t)\tilde{\boldsymbol{\rho}}_n$ would not remain confined to P, thereby leading to a contradiction.

Instead, the steady state should be regarded as an anchor point within P, with the other eigenvectors representing directions along which this ensemble can be explored. The convexity ensures that we remain within P as long as we remain in a reasonable vicinity of the anchor point. With this mental framework in place, we define $(\boldsymbol{\rho}_0, ..., \boldsymbol{\rho}_n, ...)$ the normalized Liouvillian eigenvectors as follows:

$$\begin{aligned} \boldsymbol{\rho}_0 &= \frac{1}{2}(\tilde{\boldsymbol{\rho}}_0 + \tilde{\boldsymbol{\rho}}_0{}^\dagger) \\ \forall n > 0, \boldsymbol{\rho}_n &= S_n \frac{1}{2}(\tilde{\boldsymbol{\rho}}_n + \tilde{\boldsymbol{\rho}}_n{}^\dagger) \end{aligned} \tag{2.53}$$

Where S_n is the largest positive real value such that:

$$\min\left(\mathrm{Sp}\left(\boldsymbol{\rho}_0 + S_n \frac{1}{2}(\tilde{\boldsymbol{\rho}}_n + \tilde{\boldsymbol{\rho}}_n{}^\dagger)\right)\right) \geqslant 0 \tag{2.54}$$

Under this condition, it is guaranteed that $\boldsymbol{\rho}_0 \pm \boldsymbol{\rho}_n, n > 0$ lies on the boundary of the ensemble of admissible density matrices P. In the context

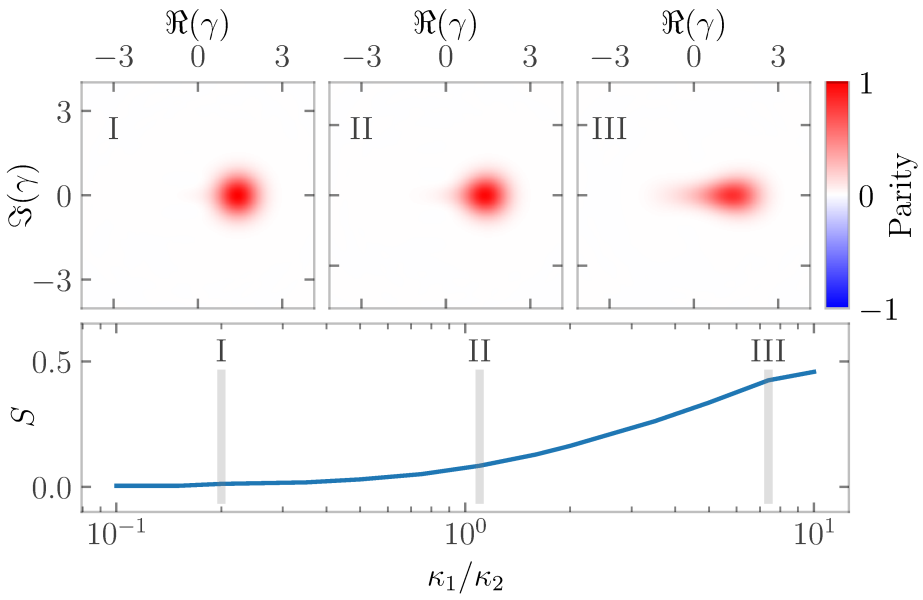

Figure 2.10 Deformation of the logical zero state in the presence of losses. The top panels show the Wigner representation of the logical zero state, computed through Liouvillian diagonalization for increasing values of κ_a/κ_2. From left to right, the coherent state becomes increasingly distorted towards the origin of the phase space. The bottom panel displays the Von Neumann entropy as this ratio increases. As soon as memory losses enter the picture, the logical state becomes impure.

of dissipative cat qubits, this allows us to give meaning to the notion of the code space, even when $\kappa_a > 0$. By solving for the code space with different values of the ratio κ_a/κ_2, one can observe the distortion of the coherent states encoding the logical zero and one of our qubit (Fig. 2.10). In particular, an important observation is realizing that these states are no longer pure: we are encoding quantum information in a coherent mixture of decoherent states, a nice little intriguing apparent paradox, which, upon closer examination, turns out not to be a paradox at all.

2.1.6 Bias preserving single qubit gates

Now that we have a clear understanding of the noise bias inherent to our quantum bistable system, we can explore the implementation of the bias-

preserving operations discussed in the previous chapter.

The $\mathbf{Z}(\boldsymbol{\theta})$ gate To begin, let's elucidate how we can perform rotations around the Z axis of the Bloch sphere. By examining the fringes of the states lying on its equator for the code space of amplitude $\alpha^2 = 2$ (as depicted in Fig. 2.6), we observe a distinctive gliding motion, either upwards or downwards, contingent upon whether the rotation is clockwise or anticlockwise around the sphere. This observation is particularly evident for the code space of amplitude $\alpha^2 = 0$, where the blue spot moves up or down through the red one, appearing and disappearing as it traverses.

Based on this empirical observation, one can intuitively grasp that applying a Hamiltonian whose flow is vertical, in conjunction with the two-photon stabilization, will result in the $\mathbf{Z}(\theta)$ gate. Remarkably, this intuition is entirely accurate, and such a flow is achieved by introducing a resonant drive with the appropriate phase:

$$\mathbf{H}_{\text{drive}} = i\varepsilon_z \mathbf{a} + \text{h.c.} \tag{2.55}$$

Essentially, as long as the amplitude of the flow drive ε_z remains sufficiently small compared to the confinement rate κ_{conf}, the system remains within the code space. Formally, if the system is initialized in the state $|+\rangle_\alpha$, the weak drive induces Rabi oscillations between $|+\rangle_\alpha$ and $|-\rangle_\alpha$ Fig. 2.11. Following the procedure outlined in [Guillaud et al. 2023], the rate of these oscillations can be computed by projecting the drive Hamiltonian Hdrive onto the code space:

$$\begin{gathered}(|+\rangle_\alpha\langle+|_\alpha + |-\rangle_\alpha\langle-|_\alpha)\mathbf{H}_{\text{drive}}(|+\rangle_\alpha\langle+|_\alpha + |-\rangle_\alpha\langle-|_\alpha) \\ = \\ 2\Im(\alpha\varepsilon_z)(|+\rangle_\alpha\langle-|_\alpha + |-\rangle_\alpha\langle+|_\alpha) + O(e^{-2|\alpha|^2})\end{gathered}$$

The concept that this gate results from the projection of a perturbative Hamiltonian onto the space toward which dissipation continuously projects the state is nothing short of a Zeno effect, as demonstrated in [Azouit et al. 2016]. Therefore, later on, we will colloquially refer to this scheme as a Rabi-Zeno gate. We indeed find that the gate speed is maximized when the drive ε_z is orthogonal to the cat state amplitude α. To rotate on the Bloch sphere by an angle θ, we find that the time T during which $\mathbf{H}_{\text{drive}}$ must be active is given by:

$$T = \frac{\theta}{4\alpha\varepsilon_z} \tag{2.56}$$

Assuming that $2\kappa_a\alpha^2 < \varepsilon_z \ll \kappa_{\text{conf}}$, the value of the resonant drive ε_z is such that the gate is faster than the decoherence induced by the one-photon losses, and the confinement rate is strong enough to maintain the state in the code space. However, if, on the other hand, the drive amplitude starts to be on par with the confinement rate, the leakage out of the code space causes the fringes of the cat to fade faster.

This phenomenon can be understood as drive-induced dephasing. Let us compute the steady state in the semi-classical theory of the two-mode model Eq.2.40 while adding the drive Hamiltonian Eq.2.55. We find that the state of the buffer now depends on the state of the memory. Denoting a and b as the mean fields of the memory and the buffer, respectively, we get:

$$b = \mp\frac{\varepsilon_z}{2a^* g_2} \tag{2.57}$$

Therefore, an information leak proportional to the drive amplitude is created through the field that is constantly radiated by the buffer at a rate $\kappa_b/2$. Following [Gambetta et al. 2006], the induced dephasing is related to the rate at which information is gained on the memory state. Replacing a by the amplitude of the cat state in the memory α, we get:

$$\Gamma_Z^{\varepsilon_z} = \frac{\kappa_b}{2}\left|\frac{\varepsilon_z}{2\alpha g_2}\right|^2 \tag{2.58}$$

Where $\Gamma_Z^{\varepsilon_z}$ is the drive induced dephasing, the rate at which the fringes fade is now given by $2\kappa_a\alpha^2 + \Gamma_Z^{\varepsilon_z}$.

The X gate As we stated previously, the challenge with the X gate is to preserve the noise bias while traversing from one pole of the Bloch sphere to the other. If we restrict our view to the subspace it spans, then this operation is obviously impossible, as at some point the equator must be crossed.

However, examining the Wigner function of a cat state reveals that a π rotation in phase space corresponds to a π rotation of the Bloch sphere around its X axis. Importantly, the fact that at any point during this rotation the supports of the information carried by the computational states $|\pm\alpha\rangle$ are disjoint means that phase-flips have no chance to propagate to bit-flips. During the intermediary steps of the rotation, the state of the system evades the steady-state manifold of the two-photon stabilization, allowing the noise bias to be preserved Fig. 2.11.

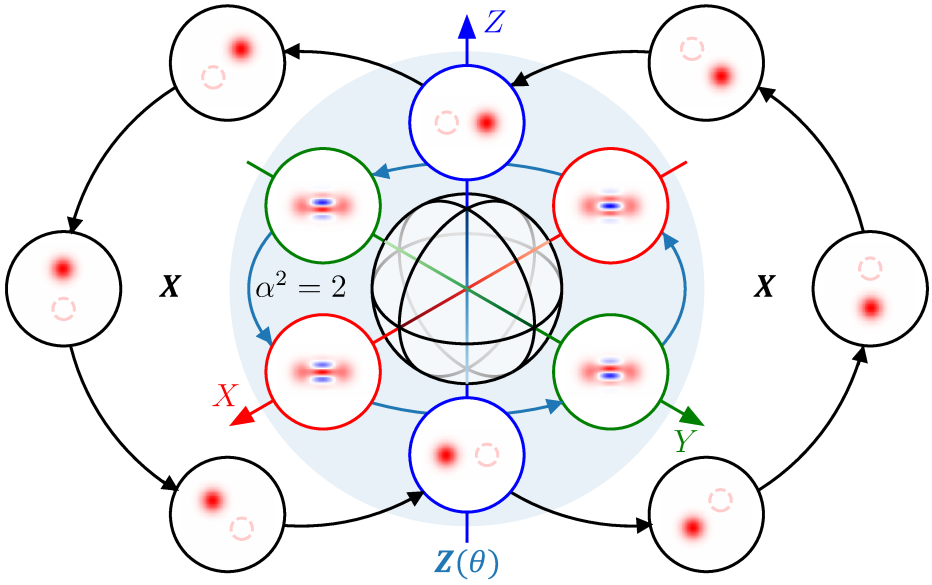

Figure 2.11 Bias preserving gates on a cat qubit. The code space for amplitude $\alpha^2 = 2$ is displayed. The blue arrows show a Rabi-Zeno scheme allows to perform a $\mathbf{Z}(\theta)$ gate. Following them, the fringes of the coherent cat state go up. The black arrows show how a rotation of the phase space allows to escape the Bloch sphere to perform a bias preserving $\boldsymbol{X}$ gate.

This gate can be implemented in several ways, an extremely easy scheme to do it experimentally is by applying a detuning Hamiltonian:

$$\mathbf{H}_{\text{detuning}} = \Delta_a \mathbf{a}^\dagger \mathbf{a} \tag{2.59}$$

Here, Δ_a represents the detuning of the memory mode during the gate operation. The duration required to perform the $\boldsymbol{X}$ gate is therefore:

$$T = \Delta_a T/2 \tag{2.60}$$

However, during gate operation, the qubit is not protected, and errors could be induced by parasitic dephasing or thermal population in the memory. However, as long as the gate time is much shorter than these timescales, we are relatively safe. Since it is straightforward to implement large detunings Δ_a in practice, this shortcoming should not be too worrying from an experimental perspective.

The CNOT gate The bias-preserving single qubit gates have been shown to be possible with a dissipative cat-qubit. However, to formally prove that universal quantum computing is achievable with cat qubits, we also need to demonstrate the feasibility of a CNOT gate between a control qubit $\mathbf{a}_c$ and a target qubit $\mathbf{a}_t$. We do not delve into this rabbit hole in this manuscript, as our focus lies solely on single-qubit physics from an experimental perspective. However, it is worth mentioning that a CNOT gate can be achieved by implementing a longitudinal Hamiltonian $\mathbf{H}_{\text{long}}$, also referred to as the optomechanical coupling in other contexts [Singh et al. 2014], between these two modes:

$$\mathbf{H}_{\text{CNOT}} = g_{\text{CNOT}}(\mathbf{a}_c + \mathbf{a}_c^\dagger)\mathbf{a}_t^\dagger \mathbf{a}_t \tag{2.61}$$

Intuitively, such a Hamiltonian implements a detuning of the target qubit conditional to the $\mathbf{Z}$ observable of the control qubit at a rate $2\alpha g_{\text{CNOT}}$. By disabling the stabilization of the target qubit during the gate time, it is possible to perform the CNOT gate.

However, it is important to note that in the literature, the initial proposal for implementing the CNOT gate involved a conditional rotation of the two-photon stabilization of the target qubit [Guillaud et al. 2023]. This scheme utilized a longitudinal Hamiltonian (referred to as the feedforward Hamiltonian in this context) as a means to enhance the fidelity of the operation by combining both techniques. As of now, from a practical perspective, it remains uncertain which approach is the most convenient for implementing the CNOT gate while ensuring decent fidelities.

For completeness, we must mention that our universal gate set also includes the ability to prepare a $|+\rangle_\alpha$ state, as well as the ability to measure the value of the $\boldsymbol{X}$ operator. The preparation of the $|+\rangle_\alpha$ state can be achieved by enabling the two-photon stabilization starting from the vacuum state. On the other hand, measuring the value of the parity operator $\langle e^{i\pi \mathbf{a}^\dagger \mathbf{a}} \rangle$ is feasible from an experimental perspective and is not a significant challenge.

At the end of the day, dissipative cat-qubits look extremely promising on paper as they should offer many advantages compared to other technologies. Let us demonstrate how they can be implemented experimentally.

3
MECHANICAL ANALOG

I believe that a physical phenomenon is best understood when it is embodied in a tangible object. One can then manipulate the object with ones hands, and observe the effects with ones eyes. For this reason, I built a mechanical oscillator endowed with two-photon dissipation. The idea is not new, as it was first proposed in [Lescanne 2020a]. The structure of the device is presented in Fig. 3.1.

Here, the memory is embodied by a pendulum composed of a mass suspended to a rod, which natural frequency id denoted ω_a. The pendulum's pivot is attached to the bottom end of a spring capable of vertical movement, effectively inducing a parametric pumping reminiscent of a child standing and crouching on a swing (see Fig. 1.9). In effect, the translational-to-rotational motion conversion ensured by this connection is responsible for a 3-wave mixing between the pendulum, and the oscillating mode associated with the spring, which takes on the role of the buffer.

However, unlike the swing in the previous chapter, we assume that the pendulum's pivot is frictionless. Instead, the stability of the system is ensured by a damper parallel to the spring, which role is to dissipate the energy of its oscillations. This arrangement results in effective two-photon dissipation. By agitating the top-end of the spring at twice the natural frequency of the memory $\omega_d = 2\omega_a$, it is possible to set the pendulum in motion. Intuitively, we enter a regime where, for every full oscillation of the spring, the pendulum completes half an oscillation, swinging from left to right or right to left. Thanks to the effective two-photon dissipation imposed by the damper, the oscillations of the pendulum while eventually stabilize around an asymptotic amplitude.

Before unveiling the mechanical device I've constructed, I'd like to emphasize the potency of the Lagrangian formalism using this concrete example. While many may find its machinery cumbersome, in reality it offers a remarkably efficient and systematic approach to deriving the equations of motion for even the most complex systems. Unlike the Newtonian formalism, which deals primarily with forces, Lagrangian reasoning operates at the level of energies.

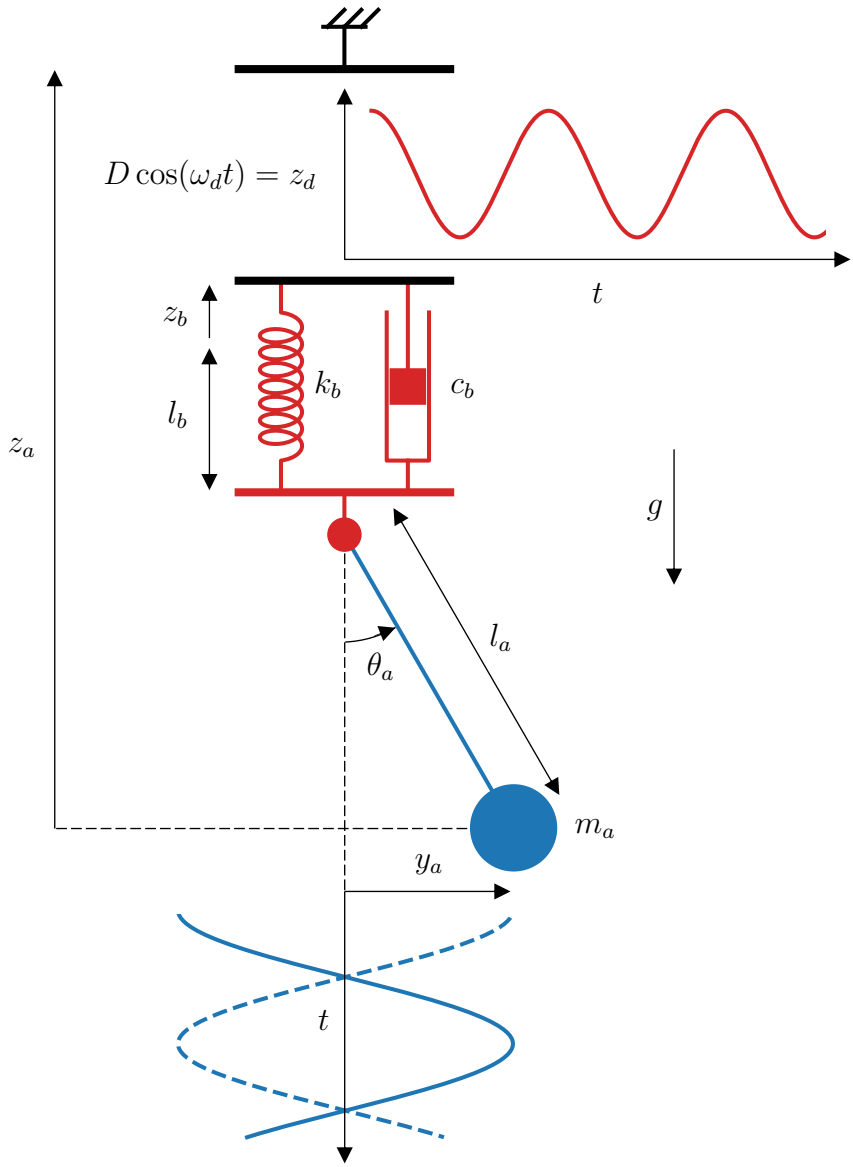

Figure 3.1 Mecanical analogy schematic. The memory mode (blue) is embodied by a pendulum, while the role of the buffer is played by a vertical oscillating spring (red). The damper parallel to the spring ensures that the buffer mode dissipates its energy at a high rate in the thermal bath. Applying a drive at half the memory natural frequency, the system converges to one of two steady states.

The first step involves writing the degrees of freedom of our system. In our case, looking at Fig. 3.1, there are two: the angle of the pendulum θ_a, and the position of the spring z_b. From there, the only tricky part consist in parameterizing the potential energy V and the kinetic energy T in terms of these variables. Essentially, it involves simple geometrical considerations based on vector addition and derivatives. We get:

$$\begin{aligned} T &= \frac{1}{2} m_a (\dot{z_a}^2 + \dot{y_a}^2) \\ &= \frac{1}{2} m_a (l_a^2 \dot{\theta}_a^2 + \dot{z}_b^2 + \dot{z}_d^2) - m_a \dot{z}_b \dot{z}_d - m_a l_a \sin(\theta_a) \dot{\theta}_a (\dot{z}_b + \dot{z}_d) \\ V &= \frac{1}{2} k_b z_b^2 - m_a g z_a \\ &= \frac{1}{2} k_b z_b^2 + \frac{1}{2} m_a g l_a \theta_a^2 - m_a g (z_d + l_a \cos_4(\theta_a)) + \text{Const.} \end{aligned} \tag{3.1}$$

By definition, the Lagrangian is the difference between the kinetic energy and the potential energy:

$$\mathcal{L} = T - V \tag{3.2}$$

Since dissipation plays a pivotal role in our system, we require a way to account for it. To this end, we derive the expression for the power dissipated R by the damper, a quantity also recognized as the Rayleigh dissipation function [Goldstein et al. 2002, p. 22], which can be seen as a quick and dirty way to add dissipation to Lagrangian formalism.

$$R = \frac{1}{2} c_b \dot{z_b}^2 \tag{3.3}$$

Following this, although the calculations may not yield aesthetically pleasing results, they can be executed almost mechanically. The equations of motion are derived by applying the Lagrange equations to θ_a and z_b, augmented by the terms introduced by the Rayleigh function.

$$\frac{d}{dt} \frac{\partial \mathcal{L}}{\partial \dot{x}} = \frac{\partial \mathcal{L}}{\partial x} - \frac{\partial R}{\partial \dot{x}}, \quad x \in \{\theta_a, z_b\} \tag{3.4}$$

Replacing terms, the left-hand side represents the derivative of momentum, while the right-hand side accounts for the sum of applied forces on the system, corresponding with the Newtonian picture.

$$\begin{aligned} m_a l_a (l_a \ddot{\theta}_a - (\dot{z}_b + \dot{z}_d)\dot{\theta}_a \cos(\theta_a)) &= m_a g l_a \sin(\theta_a) \\ m_a (\ddot{z}_b - \ddot{z}_d - l_a \sin(\theta_a)\dot{\theta}_a) &= k_b z_b - c_b \dot{z}_b \end{aligned} \tag{3.5}$$

Following this, employing a numerical solver, one can integrate the equations of motion to forecast the system's behavior. Essentially, the classical dynamics is solved thanks to the Lagrangian technique, and it requires no further elaboration.

Another advantage of this formalism becomes apparent when delving into quantum physics, as the transition to the Hamiltonian picture is straightforward. To illustrate this, as a small exercise, we explicitly undertake the second quantization of the system and map it to the two-mode model defined in Eq. 2.40. This entails defining the reduced variables Θ_a and Z_b by considering the linear portion of the Lagrangian, along with their conjugate variables P_a and P_b. By definition, these are determined by taking the partial derivative of the Lagrangian with respect to the derivatives of these variables.

$$\omega_a = \sqrt{\frac{g}{l_a}} \qquad \Theta_a = \sqrt{\frac{m_a \omega_a}{\hbar}} l_a \theta_a \qquad P_a = \frac{\partial \mathcal{L}}{\partial \dot{\Theta}_a} = \sqrt{\frac{m_a}{\hbar \omega_a}} l_a \dot{\theta}_a$$

$$\omega_b = \sqrt{\frac{k_b}{m_a}} \qquad Z_b = \sqrt{\frac{m_a \omega_b}{\hbar}} z_b \qquad P_b = \frac{\partial \mathcal{L}}{\partial \dot{Z}_b} = \sqrt{\frac{m_a}{\hbar \omega_b}} \dot{z}_b$$

After this effort, the Hamiltonian can be straightforwardly derived using its textbook definition:

$$\mathbf{H} = \dot{\Theta}_a P_a + \dot{Z}_b P_b - \mathcal{L} \tag{3.6}$$

Replacing terms and looking at its expression, one can remark it coincides with the sum of the kinetic energy and the potential energy. However, here, caution is needed, as it is not always the case! On this one, I will quote my PhD adviser, who once told me: "I shall never obtain the Hamiltonian from flipping the sign of the potential energy in the Lagrangian."

With the most challenging part behind us, variables are transformed into operators: $\Theta_a \rightarrow \boldsymbol{\Theta}_a$, $Z_b \rightarrow \boldsymbol{Z}_b$, $P_a \rightarrow \boldsymbol{P}_a$ and $P_b \rightarrow \boldsymbol{P}_b$, and the second quantization is achieved by expressing them in terms of creation and annihilation operators.

$$\mathbf{a} = \frac{1}{\sqrt{2}}(\boldsymbol{\Theta_a} + i\boldsymbol{P}_a) \qquad \mathbf{a}^\dagger = \frac{1}{\sqrt{2}}(\boldsymbol{\Theta}_a - i\boldsymbol{P}_a)$$
$$\mathbf{b} = \frac{1}{\sqrt{2}}(\boldsymbol{Z}_b + i\boldsymbol{P}_b) \qquad \mathbf{b}^\dagger = \frac{1}{\sqrt{2}}(\boldsymbol{Z}_b - i\boldsymbol{P}_b)$$

Within the framework of the master equation formalism, we can express the Hamiltonian and the jump operators to describe the dynamics in a second quantized manner:

$$\mathbf{H} \approx \hbar\omega_a \mathbf{a}^\dagger \mathbf{a} + \mathbf{H}_{\text{drive}} + \mathbf{H}_{\text{2ph}} \tag{3.7}$$

Where:

$$\begin{aligned} \mathbf{H}_{\text{drive}} &= i\sqrt{\frac{\hbar\omega_b m_a}{2}}(\mathbf{b}^\dagger - \mathbf{b})\dot{z}_d \\ \mathbf{H}_{\text{2ph}} &= -\frac{\hbar}{2}\sqrt{\omega_a\omega_b}\sin\left(\frac{1}{l_a}\sqrt{\frac{\hbar}{2m_a\omega_a}}(\mathbf{a}^\dagger + \mathbf{a})\right)(\mathbf{a}^\dagger - \mathbf{a})(\mathbf{b}^\dagger - \mathbf{b}) \end{aligned} \tag{3.8}$$

By examining the non-linear part of the Hamiltonian, the underlying principle of the design becomes more evident. $\mathbf{H}_{\text{2ph}}$ to the following interpretation: as the pendulum deviates from the vertical position, the torque experienced by the pivot becomes increasingly proportional to the motion of both the spring and the pendulum itself. When we replace the terms with **a** and **b**, this component is responsible for the three-wave mixing phenomenon. Conversely, $\mathbf{H}_{\text{drive}}$ is born from the vertical motion of the spring anchor, providing a means to inject energy into the buffer mode. Through a Taylor expansion and a rotating wave approximation, while imposing $\omega_d = 2\omega_a$, we can identify this model as the two-mode model we described earlier in Equation 2.40.

$$|g_2| = \frac{1}{2l_a}\sqrt{\frac{\hbar\omega_b}{2m_a}}, \quad |\varepsilon_d| = \frac{D\omega_d}{2}\sqrt{\frac{\hbar\omega_b m_a}{2}}, \quad \kappa_b = \frac{c_b}{m_a} \tag{3.9}$$

Let us now make the discussion concrete! As I was writing my PhD, a built a real-world implementation of this mechanism, as illustrated in Fig .3.2, using my 3D printer. Beginning from the idle state, a motorized crankshaft, connected to the spring, initiates motion in the system. Over time, the oscillations of the spring decrease in amplitude as the pendulum's

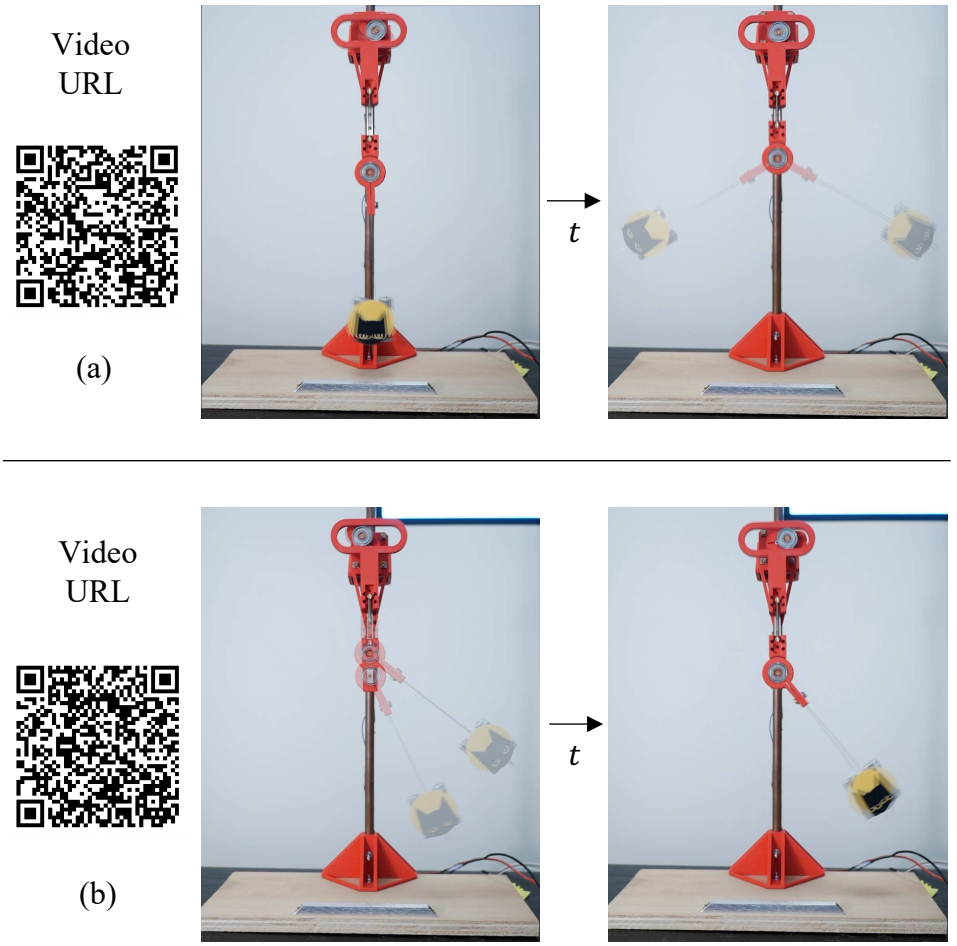

Figure 3.2 Mechanical analogy implementation. Synchronized video captures of the pendulum dynamics. Starting from the idle state (a), the pendulum state converges to one side of the phase space or the other. Conversely, starting from two initial conditions on the same side of the phase space (b), the pendulum converges to the same asymptotic trajectory. QR Codes leading to videos are displayed for both experiments.

swing widens. The system eventually reaches a state where the spring exhibits minimal motion, while the pendulum continues its constant oscillations. From there, it is possible to demonstrate two effects.

First, when several videos of the system's dynamics are overlaid, and the motion of the motor is synchronized, it becomes evident that the system converges to one of two steady states with opposite phases. In the rotating frame, these states are entirely analogous to $|+\alpha\rangle$ and $|-\alpha\rangle$ in the cat qubit code space. By manipulating this contraption, one can firsthand observe the system's stability, which is responsible for bit-flip protection Fig. 3.2a.

Secondly, while this mechanism operates in the classical realm, it can serve as an excellent tool to explain why two-photon stabilization preserves quantum superpositions. The key insight is to recognize that the motion of the damper is independent of the pendulum's asymptotic trajectory Fig. 3.2b. From the damper's perspective, the state of the memory remains concealed. The system's stability is achieved through energy dissipation into the thermal bath, yet the information it carries is unrelated to the pendulum's phase. The state remains separable, and coherences are preserved.

Something I'm extremely proud of is that I was able to use this mechanism in front of a very large audience (700 000 views on YouTube to this date), during the course of an event on Twitch, to popularize the principle underlying the cat-qubit technology.

4

CAVITY QED IMPLEMENTATION OF DISSIPATIVE STABILISATION

Of course, implementing a quantum device using springs, joints, and pivots would be exceedingly challenging. Today, multiple approaches to quantum computing are in competition. For instance, neutral atoms operated in arrays of optical traps are excellent candidates to store and operate quantum information [Henriet et al. 2020]. In a very recent paper, surface codes from distances $d = 3$ to $d = 7$ were demonstrated using this technic [Bluvstein et al. 2023]. Our work focuses on another promising field, known as Cavity Quantum ElectroDynamics (CQED).

As explained in [Girvin 2012], At sufficiently low temperatures, a lot of conventional metals, like aluminum, transition to a superconducting state. The idea is to craft an electrical circuit out of such a material, and operate it in this regime. Here, electrons with opposite spins form doublets called Cooper pairs, effectively moving through the metal without dissipation. Collectively, their behavior resembles a plasma, which plasma frequency falls in the ultraviolet region of the optical spectrum.

Since in CQED we operate well below this plasma frequency, we get an interesting phenomenon: the degrees of freedom associated with the creation and annihilation of plasmons effectively freeze. Despite typical circuits comprising billions to trillions of aluminum atoms, they exhibit only a few degrees of freedom. Consequently, an LC oscillator crafted from a centimeter-wide circuit can be accurately modeled by the quantized harmonic oscillator. Maybe surprisingly, the discretization of its single degree of freedom is for instance quite easy to demonstrate experimentally as we will see later Fig. 5.14.

Thanks to a technology called dilution cooling, it's feasible to reach temperatures of a few milli-Kelvins in what's called a cryostat. Operating these circuits at energies E_{op} significantly below the energy threshold for breaking Cooper pairs 2Δ, and ensuring that the energy of thermal excitations

k_BT remains well below the operational energy E_{op}, enables the observation of dissipationless quantum phenomena.

$$\begin{array}{ccccc} \approx 20\text{ mK} & & & & \approx 340\ \mu\text{eV} \\ | & & & & | \\ k_BT/h & \ll & E_{\text{op}}/h & \ll & 2\Delta/h \\ (\approx 0.4\text{ GHz}) & & (\approx 5\text{ GHz}) & & (\approx 160\text{ GHz}) \end{array} \tag{4.1}$$

The procedure to properly quantify a given circuit can be extremely complex depending on its topology. In this text, we restrict the discussion to the simple cases of the Harmonic oscillator and the Josephson junction, and show how two-photon stabilisation can be implemented. To this end, we heavily rely on the calculations done in [Girvin 2014] and [Devoret et al. 1995]. If the reader needs more detail, we encourage him to check out these texts. Here, we simply introduce the notations we need to treat our circuit.

4.1 The superconducting quantum LC-oscillator

The LC-oscillator is obtained by assembling a capacitor of capacitance C and an inductor of inductance L, a situation shown in Fig. 4.1. Following Devoret [Devoret et al. 1995] we define the node flux at the junction of the inductor and the capacitor as:

$$\phi(t) = \int_{-\infty}^{t} V(\tau)d\tau \tag{4.2}$$

In this expression, V is the node voltage. Since $V = \dot{\phi}$, the energy stored in the capacitor can be written in terms of the magnetic flux. As latter, we will work with Josephson junctions, it will be convenient for us to consider it as a kinetic energy T:

$$T = \frac{1}{2}C\dot{\phi}^2 \tag{4.3}$$

Similarly, using Faraday's law in Fig. 4.1, we get the relation between the current flowing through the inductor L and the magnetic flux:

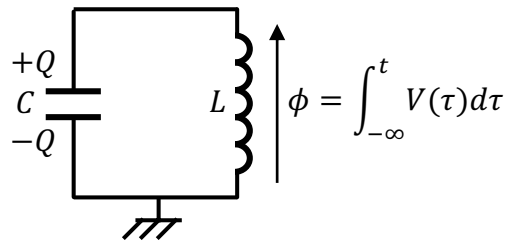

Figure 4.1 LC-oscillator diagram. A capacitor C and an inductor L are arranged in parallel. φ is defined as the integral of the voltage V between the ground and the point where the capacitor and the inductor are connected. The energy stored in the inductor is proportional to the square of φ, while the one in the capacitor is connected to the charge Q.

$$V = L\dot{I} = \dot{\phi} \tag{4.4}$$

ϕ is indeed the magnetic flux threading through the inductor loops, and we can write the energy it stores:

$$U = \frac{1}{2L}\phi^2 \tag{4.5}$$

From there, the expression of the Lagrangian is recovered in terms of the superconducting phase ϕ:

$$\mathcal{L} = T - U = \frac{1}{2}C\dot{\phi}^2 - \frac{1}{2L}\phi^2 \tag{4.6}$$

As soon as a Lagrangian is written, we enter autopilot mode. By applying its definition, we recover the expression Q of the conjugated variable of ϕ, which happens in this case to coincide with the charge of the capacitor:

$$Q = \frac{\partial H}{\partial \dot{\phi}} = C\dot{\phi} \tag{4.7}$$

And the Hamiltonian can be written in terms of Q and ϕ:

$$\mathbf{H} = n\dot{\phi} - \mathcal{L} = \frac{1}{2C}Q^2 + \frac{1}{2L}\phi^2 \tag{4.8}$$

In the usual way, variables are replaced with operators ($\phi \rightarrow \boldsymbol{\phi}$ and $Q \rightarrow \boldsymbol{Q}$) which follow the usual commutation relation:

$$[\boldsymbol{\phi}, \boldsymbol{Q}] = +i\hbar \tag{4.9}$$

And we can obtain the usual form for the second quantized formulation of the Hamiltonian:

$$\mathbf{H} = \hbar\omega\left(\mathbf{a}^\dagger\mathbf{a} + \frac{1}{2}\right) \tag{4.10}$$

We find the usual value for the natural pulsation $\omega = 1/\sqrt{LC}$ of the oscillator, and the creation and annihilation operators read:

$$\mathbf{a} = \frac{1}{\sqrt{2L\hbar\omega}}\boldsymbol{\phi} + i\frac{1}{\sqrt{2C\hbar\omega}}\boldsymbol{Q}, \quad \mathbf{a} = \frac{1}{\sqrt{2L\hbar\omega}}\boldsymbol{\phi} - i\frac{1}{\sqrt{2C\hbar\omega}}\boldsymbol{Q} \tag{4.11}$$

Writing this relation the other way around, it is also useful to introduce the zero point fluctuation of the phase ϕ_{zpf} and of the number of Cooper pairs Q_{zpf}:

$$\phi_{\text{zpf}} = \sqrt{\frac{\hbar Z}{2}} \qquad \boldsymbol{\phi} = \phi_{\text{zpf}}(\mathbf{a}^\dagger + \mathbf{a})$$
$$Q_{\text{zpf}} = \sqrt{\frac{\hbar}{2Z}} \qquad \boldsymbol{Q} = iQ_{\text{zpf}}(\mathbf{a}^\dagger - \mathbf{a})$$

We recover that the zero point fluctuation of the flux is proportional the square root of the characteristic impedance Z of the mode $1/C\omega = \sqrt{L/C} = Z$.

4.2 The Josephson junction

In a superconducting circuit, it is possible to instantiate numerous bosonic modes by arranging multiple LC-oscillators. However, the intriguing physics arises when these modes are nonlinearly coupled, a feat not achievable with capacitors or inductors. For instance, in the mechanical analogy of the previous section, the two-photon conversion is implemented thanks to the translating motion to rotating motion conversion Fig. 3.1.

In the quantum realm, we use Josephson junctions as a dissipationless nonlinearity resource. Thanks to this tool, a wide variety of interactions can be implemented. Loosely speaking, in a way similar to how modern electronics rely on transistors, cavity QED is built around the utilization of Josephson junctions.

Josephson junctions consist of two superconducting islands separated by an ultra-thin insulating layer, typically a few nanometers thick. Diagrammatically, a Josephson junction is represented by a cross (see Fig. 4.2a). Remarkably, Cooper pairs within the superconductors can coherently tunnel across this junction. Here, we define $\mathbf{n}$ as the operator that counts the number of Cooper pairs that traversed from one side to the other. In the charge basis, it is represented as:

$$\mathbf{n} = \sum_{n=-\infty}^{\infty} n \, |n\rangle\langle n| \tag{4.12}$$

We won't delve into the microscopic details of the physics underlying the Josephson effect here, but it can be understood phenomenologically.

$$\mathbf{H}_T = -\frac{1}{2} E_J \sum_{n=-\infty}^{\infty} |n\rangle\langle n+1| + |n+1\rangle\langle n| \tag{4.13}$$

We define E_J as the Josephson coupling energy, representing the energy associated with the traversal of a Cooper pair from one side of the junction to the other. In practice, this quantity is determined by the thickness and width of the oxide barrier. The operator $\mathbf{H}_T$ causes the number of pairs n to increase or decrease one by one, depending on whether a pair tunnels from right to left or left to right through the junction. It is now insightful to introduce the plane-wave states indexed by the wave vector φ. Loosely speaking, they are obtained by taking the Fourier transform of the n representation, ignoring temporarily that the energy of $|\varphi\rangle$ is undefined:

$$|\varphi\rangle = \sum_{n=-\infty}^{\infty} e^{in\varphi} \, |n\rangle \tag{4.14}$$

In this basis, it is straightforward to see that the states $|\varphi\rangle$ are eigenvectors of $\mathbf{H}_T$. Their eigenvalues are given by:

$$\mathbf{H}_T \, |\varphi\rangle = -E_J \cos(\varphi) \, |\varphi\rangle \tag{4.15}$$

Now, a fundamental fact that unfolds from the second Josephson relation is that φ is actually proportional to a magnetic flux:

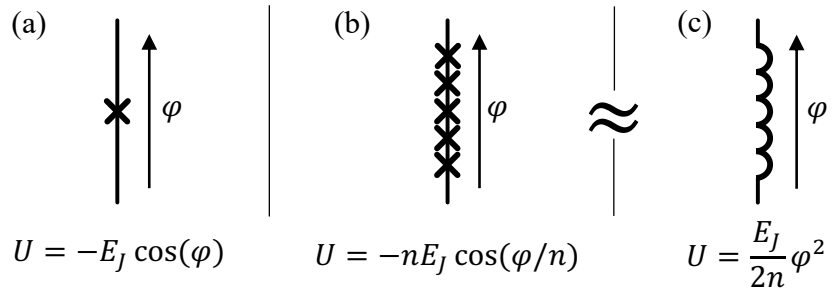

Figure 4.2 Josephson junction Diagram: (a) A single Josephson junction is depicted by a small cross, with an arrow indicating the superconducting phase difference φ. (b) By arranging multiple junctions with the same energy in an array, it becomes possible to approximate a kinetic inductance. (c) This arrangement enables the simulation of a large inductance in a compact space.

$$\varphi = \frac{2e}{\hbar}\phi \tag{4.16}$$

Since in a lot of cases, we will be dealing with Josephson junctions coupled to LC-oscillators, we will write the Hamiltonian in terms of a flux operator, which will therefore enter $\cos$ functions.As it will be extremely convenient for us to deal with dimensionless units, we express them in terms of the phase φ, which we refer to as the superconducting phase. As this variable is associated with a Harmonic oscillator, it is legitimate to cast it to an operator $\varphi \to \boldsymbol{\varphi}$, allowing us to write Hamiltonians of the form:

$$\mathbf{H}_T = -E_J \cos(\boldsymbol{\varphi}) \tag{4.17}$$

Writing a Taylor expansion yields:

$$\mathbf{H} = -E_J \cos(\boldsymbol{\varphi}) = \frac{1}{2}E_J\boldsymbol{\varphi}^2 + \cos_4(\boldsymbol{\varphi}^2) + \text{Const.} \tag{4.18}$$

In this expression, $\cos_4$ is shorthand for the Taylor expansion of the cosine starting at the fourth order. Ignoring the constant terms and focusing on the quadratic part, it is natural to interpret a junction as a nonlinear inductor. We can define the equivalent inductance $\boldsymbol{L}_J$ of the junction as inversely proportional to its energy:

$$\boldsymbol{L}_J = \frac{\hbar^2}{4e^2}\frac{1}{E_J} \tag{4.19}$$

One of the first practical applications of Josephson junctions arises from this observation. By placing junctions in series, the superconducting phase evenly spreads along the chain. Thanks to this effect, it is possible to approximate the behavior of a large inductor in a very small footprint (see Fig.4.2b). In practice, when an inductor is represented in a superconducting circuit diagram (see Fig.4.2c), it is almost always implemented with a chain of junctions.

However, the richest physics induced by the Josephson effect resides in the nonlinear part of its dynamics Eq. 4.18. By arranging Josephson junctions together, it is possible to shape arbitrarily complex interactions. This art is delicate and resides at the boundaries of our current physical understanding. Constrained by technology, simplicity is its beauty paradigm.

The physics of Josephson junctions holds immense potential, but practical considerations impose significant constraints on their utilization. Currently, achieving consistent thicknesses for the oxide layer remains challenging. Therefore, circuit designs that heavily rely on precise fabrication are not advisable. What distinguishes an engineer from a scientist is their awareness of tolerances. Hence, when designing a circuit, it's prudent to account for a margin of error of around ten percent on E_J.

Another wise practice is to routinely conduct room temperature measurements of the junctions within a circuit before proceeding with its full characterization. The Josephson energy is related to the junction's room temperature resistance R through the Ambegaokar-Baratoff formula [Ambegaokar and Baratoff 1963]. With knowledge of the superconducting gap Δ, an indirect measurement of E_J can be obtained:

$$E_J = \frac{\pi \Delta \hbar}{8e^2 R} \tag{4.20}$$

However, this procedure, while not yielding highly precise results, enables the identification and elimination of circuit samples that deviate significantly from specifications.

4.3 ATS induced two-photon exchange

Josephson junctions can be arranged in various configurations for different purposes. The design aspect of cavity QED is certainly among its most appealing as involves a significant degree of creativity. To mediate the two-photon coupling, we highlight the geometry known as the Asymmetrically Threaded SQUID [Lescanne et al. 2020b], abbreviated as ATS Fig. 4.3b.

$$U = -E_J \cos(\varphi) \cos\left(\frac{1}{2}\varphi^{\text{ext}}\right)$$

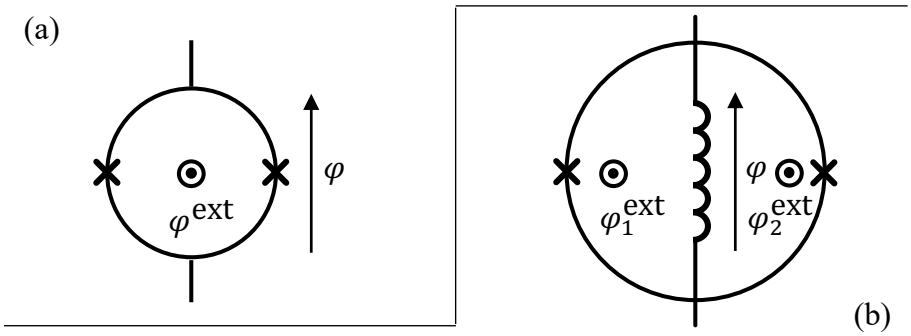

$$U = \frac{1}{2}E_L\varphi^2 - E_J \cos\left(\frac{\varphi_1^{\text{ext}} + \varphi_2^{\text{ext}}}{2}\right) \cos\left(\varphi + \frac{\varphi_1^{\text{ext}} - \varphi_2^{\text{ext}}}{2}\right)$$

Figure 4.3 ATS diagram. (a) When two junctions are placed in parallel, they form a loop where the magnetic flux φ_{ext} alters the relative phase of both of its arms, causing interference effects. A SQUID is made. When the external flux equals π, these interferences become entirely destructive, resulting in vanishing dynamics ($U(\varphi) = 0$). (b) The ATS emerges from combining a SQUID with an inductor, endowing the system with a finite equivalent inductance even when the total flux threading through the ATS is π.

This geometry possesses a unique character in more than one way: it is currently favored by Alice&Bob as it offers several advantages in the context of dissipative cat qubits. Additionally, it is the only instance known to the writer of a meta-acronym, as SQUID stands for Superconducting QUantum Interferometry Device Fig. 4.3a.

Denoting φ as the phase of the inductor in the classical picture and $\varphi_{1/2}^{\text{ext}}$ as the flux threaded through the two loops of the ATS, its the potential energy reads:

$$U_{\text{ATS}}(\varphi) = \frac{1}{2}E_L\varphi^2 - E_{J,1}\cos(\varphi + \varphi_1^{\text{ext}}) - E_{J,2}\cos(\varphi + \varphi_2^{\text{ext}}) \qquad (4.21)$$

To reflect the symmetry of the ATS in the expression of its energy, it is convenient to introduce the common flux $\varphi_\Sigma = (\varphi_1^{\text{ext}} + \varphi_2^{\text{ext}})/2$ and the differential flux $\varphi_\Delta = (\varphi_1^{\text{ext}} - \varphi_2^{\text{ext}})/2$. This parameterization is a gauge

choice, and justifying it would require taking into account the presence of parasitic capacitances [You et al. 2019]. In practice, this choice yields good agreement between theory and experiments.

$$U_{\text{ATS}}(\varphi) = \frac{1}{2}E_L\varphi^2 - 2E_J\cos(\varphi_\Sigma)\cos(\varphi+\varphi_\Delta) - 2\Delta E_J\sin(\varphi_\Sigma)\sin(\varphi+\varphi_\Delta) \tag{4.22}$$

In this expression, the asymmetry of the ATS $\Delta E_J = E_{J,1} - E_{J,2}$ is quantified by the energy difference between the two junctions. In the rest of the section, we will assume this term cancels perfectly.

To obtain the two-photon exchange mechanism, we will couple the ATS to two oscillators: strongly to the buffer mode **b**, and weakly to the memory mode **a**. This can be achieved in more than one way, and an example is presented in Fig. 4.4. After quantifying the circuit, the superconducting phase of the ATS central inductance can be expressed as the sum of contributions from the buffer and memory oscillators.

$$\varphi \rightarrow \boldsymbol{\varphi} = \varphi^a_{\text{zpf}}(\mathbf{a}^\dagger + \mathbf{a}) + \varphi^b_{\text{zpf}}(\mathbf{b}^\dagger + \mathbf{b}) \tag{4.23}$$

Where $\varphi^a_{\text{zpf}} \ll \varphi^b_{\text{zpf}}$ are the zero-point fluctuations of the memory and buffer modes across the inductance of the ATS, respectively.

Now, we must choose the flux that we are going to impose in both loops of the ATS. Depending on this flux point, we can obtain a wide variety of behaviors. To obtain the two-photon exchange process, we are going to use the ATS as a four-wave mixer, for which it is best to thread no flux through one loop and π flux through the other. Therefore, we write the common and differential fluxes as small perturbations $\delta\varphi_\Sigma$ and $\delta\varphi_\Delta$ relative to this flux point:

$$\varphi_\Sigma = \pi/2 + \delta\varphi_\Sigma, \quad \varphi_\Delta = \pi/2 + \delta\varphi_\Delta \tag{4.24}$$

In the Appendix, we show that this flux point is, in fact, a saddle point of the buffer frequency flux map (see Sec. 7.3). Injecting Eq.4.24 and Eq.4.23 into the expression of the ATS potential energy Eq. 4.22, we can isolate the nonlinear part and examine the Hamiltonian it gives rise to:

$$\mathbf{H}_{\text{ATS}} = -2E_J\sin(\delta\varphi_\Sigma)\sin(\varphi^a_{\text{zpf}}(\mathbf{a}^\dagger + \mathbf{a}) + \varphi^b_{\text{zpf}}(\mathbf{b}^\dagger + \mathbf{b}) + \delta\varphi_\Delta) \tag{4.25}$$

We then absorb the quadratic part of the ATS potential in Eq.4.22 into the definition of modes **a** and **b**. Consequently, the dynamics of the circuit shown in Fig.4.4 can be accurately described by the following Hamiltonian:

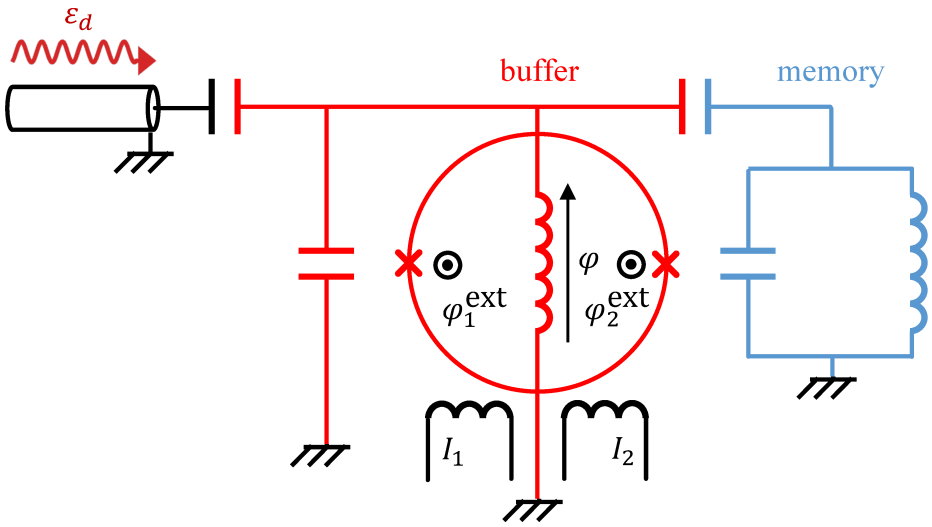

Figure 4.4 Cat-qubit circuit diagram. The inductance of the ATS is positioned at the center of an LC-oscillator, implementing the buffer mode (red). Through a coupling capacitance, the buffer is coupled to the memory mode (blue) which slightly hybridizes and therefore participates in the phase of the ATS inductance. Two current lines (I_1 and I_2) facilitate flux injection through the ATS loops. Additionally, the buffer mode is strongly capacitively coupled to a coaxial line, enabling mode driving and dissipation.

$$\begin{aligned} \mathbf{H} &= \mathbf{H}_{\text{sys}} + \mathbf{H}_{\text{ATS}} + \mathbf{H}_{\text{drive}} \\ \mathbf{H}_{\text{sys}}/\hbar &= \omega_a \mathbf{a}^\dagger \mathbf{a} + \omega_b \mathbf{b}^\dagger \mathbf{b} \end{aligned} \tag{4.26}$$

In this setup, $\mathbf{H}_{\text{sys}}$ models the resonances of $\mathbf{a}$ and $\mathbf{b}$, while $\mathbf{H}_{\text{drive}}$ represents a capacitive drive applied to the buffer mode:

$$\mathbf{H}_{\text{drive}} = 2\varepsilon_d \cos(\omega_d t) \tag{4.27}$$

Here, ε_d is the drive amplitude, and ω_d is its frequency. In practical terms, this drive is generated by sending a microwave tone through a coaxial line capacitively coupled to the buffer mode. It is this drive that is responsible for the two-photon injection. However, unlike the mechanical analogy, which corresponds to a 3-wave mixer, here we are utilizing the ATS as a 4-wave

mixer. Hence, an additional tone is required. For this purpose, we apply a parametric pump to the common flux φ_Σ.

$$\delta\varphi_\Sigma = \varepsilon_p \cos(\omega_p t), \quad \delta\varphi_\Delta = 0 \tag{4.28}$$

Where ε_p is the pump amplitude, and ω_p is its frequency. Typically, this is achieved by inductively coupling the flux loops to lines through which an oscillating current circulates. These lines also serve to impose the flux point. Substituting these terms into Eq. 4.25 and performing a Taylor expansion of the first $\sin$ term yields:

$$\begin{aligned}\mathbf{H}_{\text{ATS}} = &- E_J\varepsilon_p \cos(\omega_p t)\cos(\varphi^a_{\text{zpf}}(\mathbf{a}+\mathbf{a}^\dagger))\sin(\varphi^b_{\text{zpf}}(\mathbf{b}+\mathbf{b}^\dagger)) \\ &- E_J\varepsilon_p \cos(\omega_p t)\sin(\varphi^a_{\text{zpf}}(\mathbf{a}+\mathbf{a}^\dagger))\cos(\varphi^b_{\text{zpf}}(\mathbf{b}+\mathbf{b}^\dagger))\end{aligned} \tag{4.29}$$

Upon considering this expression, it becomes apparent that the coupling from $\cos$ to $\sin$ generates even powers of the operator $\mathbf{a}$ and odd powers of the operator $\mathbf{b}$ after Taylor expansion. It is therefore this term that we utilize to implement the two-photon exchange.

At this stage, we transition into the rotating frame for the buffer at frequency ω_d and for the memory at frequency $(\omega_p+\omega_d)/2$. Following a Taylor expansion and subsequent rotating wave approximation, we obtain the celebrated two-photon exchange Hamiltonian $g_2\mathbf{a}^2\mathbf{b}^\dagger + \text{h.c.}$. To revert to the pure two-mode model however Eq. 2.40, we must nullify $\mathbf{H}_{\text{sys}}$. This is achieved by selecting the drive and pump frequencies to satisfy the resonance matching condition:

$$\omega_d = \omega_b, \quad \frac{\omega_p + \omega_d}{2} = \omega_a \tag{4.30}$$

The parametric pump enables the two-photon exchange, while the drive injects pairs of photons in the memory. Identifying terms, the amplitude of the two-photon coupling can be deduced from the Hamiltonian parameters:

$$|g_2| = \frac{1}{2} E_J \varepsilon_p {\varphi^a_{\text{zpf}}}^2 \varphi^b_{\text{zpf}} \tag{4.31}$$

Now, armed with a solid theoretical understanding of dissipative cat physics and their implementation in circuit quantum electrodynamics (CQED), the next crucial step is to transition from theory to practice.

5

DISSIPATIVE CAT QUBIT IN A HIGH QUALITY 3D ALUMINIUM MICROWAVE CAVITY

In this chapter, we describe the initial experiment undertaken by Alice & Bob as a company. After two years of intensive work, involving numerous design iterations and hundreds of hours dedicated to data acquisition, the "3D cat experiment" did not yield the results we had anticipated. The future of the project is uncertain, and the current preferred direction for quantum computing with superconducting circuits is focusing on more scalable 2D architectures. Furthermore, the extensive iterative process this project underwent makes it challenging to publish, as the dataset covers a wide range of samples, with only a few demonstrating intriguing but partial results.

To a certain extent, as detailed in the unfolding chapters, we managed to meet some experimental challenges with success. However, issues arose due to the poor repeatability of the mechanical fixation of the chips within the 3-dimensional cavity, compounded by our insufficient comprehension of the system as a whole at that time, preventing us from producing a conclusive dataset. Instead, our experimental campaign yielded a series of insights prompting us to redesign a new sample from scratch, a decision that ultimately proved worthwhile.

This thesis provides an opportunity to showcase the tremendous effort invested by Alice & Bob teams in this project. The results it produced, though negative for the most part, are valuable clues in the jigsaw puzzle that is the design of a quantum computer. As is often the case in experimental science, the journey was the destination, and a special effort will be made to extract as much information as possible from the various design iterations. Furthermore, it serves as an excellent playground to illustrate the typical calibration procedures employed in circuit QED. Given this conducive environment, we will aim to explain them in the most pedagogical

manner.

5.1 Experiment conception

As the saying goes in engineering: your requirements are dumb! Certainly, an experiment which concept is to stabilize a dissipative cat qubit in a superconducting 3D cavity is no exception to this rule of thumb, as it almost stance for an oxymoron. Let us detail how we get to this insane idea.

The main figure of merit of a cat qubit is the ratio κ_a/κ_2. A natural way to improve on it is to push the quality factor of the memory resonator to the limit. When it comes to achieving a high-quality factor, several strategies are at our disposal. For instance, AWS outlines in its roadmap [Chamberland et al. 2022a] the use of mechanical resonators coupled to superconducting circuits. Indeed, mechanical resonators have demonstrated high quality factors [Singh et al. 2014, Huttel et al. 2009], but engineering and coupling them to superconducting circuits remains challenging. On the other hand, 3-dimensional superconducting cavities offer cutting-edge results as of the writing of this manuscript.

For dissipative cat qubits, however, this choice comes with drawbacks. The apparatus for 4-wave mixing is implemented using a 2-dimensional Josephson based circuit: establishing strong capacitive coupling between an on-chip antenna and a 3-dimensional mode remains notably challenging. This specific design decision imposes a stringent limitation on the achievable values for the two-photon coupling rate g_2. Furthermore, achieving strong coupling to the buffer mode without compromising the intrinsic 1-photon loss rate κ_1, which initially justified this architecture, becomes imperative. Excellent filtering is needed. The geometry we discuss is the result of these non trivial constraints.

5.1.1 Stub-geometry 3-dimensional cavity

In circuit QED, three main loss mechanisms have been identified. The first arises from surface resistance as predicted by BCS theory [Halbritter 1974] for superconductors. It depends on temperature and goes as $e^{-\Delta/k_bT}/T$, where Δ represents the superconducting gap. The second originates from dielectric losses that can emerge in chip substrates or the oxide layer covering the surface of superconductors. Lastly, if the resonator involves multiple metallic pieces, there will be residual resistances associated with each seam due to contact imperfections. Within a 3D cavity, the electromagnetic field

predominantly exists in the vacuum, thereby dielectric and surface resistance losses are minimal.

Throughout my PhD research, we explored various geometries. For this experiment, we sought a solution that balanced fabrication complexity, form factor, and single photon loss rate. A stub-geometry cavity [Kudra et al. 2020] offered a favorable compromise across these criteria.

This particular design features a $\lambda/4$-stub positioned within a rectangular waveguide drilled into a block of 4N aluminium. The stub is grounded at the base of the cavity, positioned away from its orifice. At the top of the stub, an open boundary condition is met, concentrating the electric field, while at the bottom, a short circuit favors the magnetic field, giving rise to a quarter-wavelength eigenmode. Field intensities exponentially diminish towards the lid, limiting losses linked to residual seam resistance.

Comparing this stub geometry to ellipsoidal cavities, which currently represent the state of the art, reveals distinct trade-offs. In the stub geometry, adjusting the waveguide length can minimize seam participation, yet the inability to polish surfaces results in the oxide layer causing significant dielectric losses. Conversely, recent literature demonstrates that employing an ellipsoidal cavity geometry achieves superior surface finishes, reducing surface and dielectric losses to a minimum. However, the seam resides at the cavity's center. Notably, ellipsoidal cavities have attained lifetimes of approximately 100 ms [Lei et al. 2023].

5.1.2 Cavity caracterisation

Before delving into stabilizing coherent cat states within this geometry, our initial step involved designing standalone cavities to characterize the achievable quality factors within our laboratory using our fabrication methods. These samples were equipped with two antenna ports, enabling measurements of the cavity in both reflection and transmission.

Writing the input-output theory [Gardiner and Collett 1985] for a single harmonic oscillator of frequency ω_0 coupled to two external lines at rates $\kappa_c^{1,2}$, along with internal losses κ_i, leads to the derivation of formulas for the reflection S_{11} and transmission S_{21} scattering parameters:

$$\begin{aligned} S_{11}(\omega) &= \frac{\kappa_i + \kappa_c^2 - \kappa_c^1 + 2i(\omega - \omega_0)}{\kappa_c^1 + \kappa_c^2 + \kappa_i + 2i(\omega - \omega_0)} \\ S_{21}(\omega) &= \frac{2\sqrt{\kappa_c^1 \kappa_c^2}}{\kappa_c^1 + \kappa_c^2 + \kappa_i + 2i(\omega - \omega_0)} \end{aligned} \tag{5.1}$$

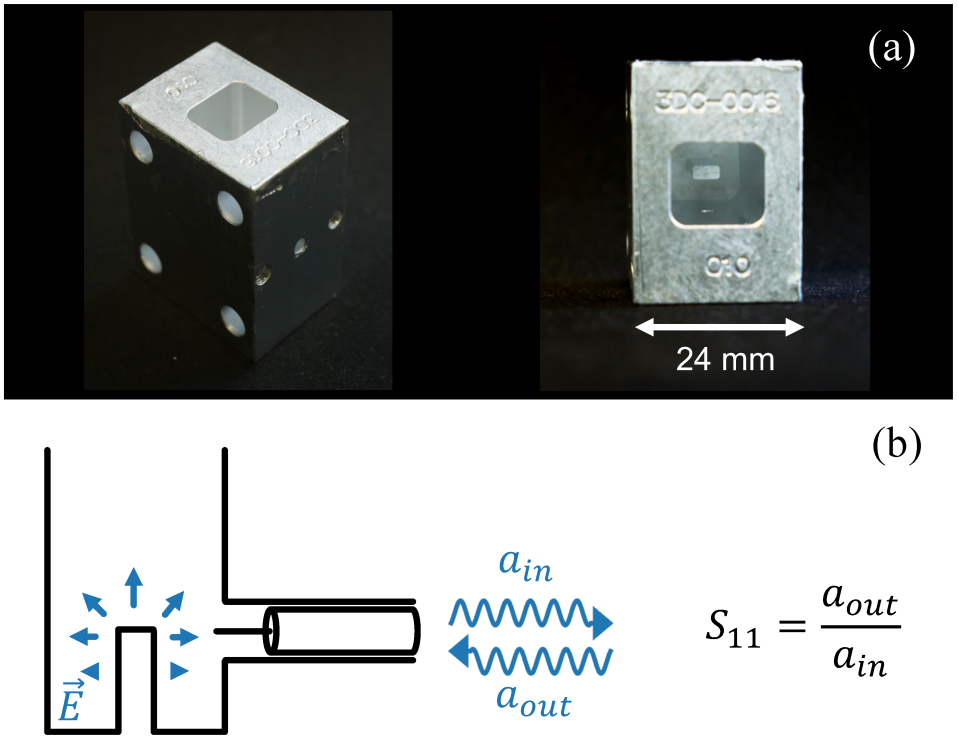

Figure 5.1 Stub-geometry 3-dimensional aluminum cavity used for characterization. (a) Two photographs of the cavity with different angles. (b) Implementation of the cavity probing apparatus: an antenna is introduced through a tunnel in the cavity, coupling to its electric field. By measuring the ratio of the incoming field a_{in} and the reflected field a_{out}, one can reconstruct the reflection scattering parameter S_{11}.

These formulas (redefined in Sec. 8.1) may not seem like much, yet they are the bread and butter of the circuit QED engineer. They describe the appearance of a resonance in a spectroscopy, which is the most fundamental feature an experimentalist can seek. The Vectorial Network Analyzer (VNA) serves as a biologist's binoculars to a circuit QED engineer. For an experimentalist to be effective, a profound understanding of these formulas, akin to having one's best friends, is crucial. Spending time with them unravels valuable insights.

Primarily, these scattering parameters represent first-degree rational functions of ω, resembling circles in the complex plane. The prevailing method used to fit these formulas to experimental data, known as the circle fit pro-

cedure, is flawed. It disregards the parametric curve's 'kinematics' as functions of ω and is susceptible to misinterpretations due to electrical delay. In this manuscript, we introduce a semi-analytical method, termed **abcd-rf-fit** (see Chap. 8), surpassing the circle fit procedure based on our benchmarks.

Assuming no prior knowledge of line attenuation, reflection, and transmission doesn't yield equivalent information. The former reveals the cavity frequency ω_0 and the total cavity loss rate $\kappa_{\text{tot}} = \kappa_c^1 + \kappa_c^2 + \kappa_i$ in a robust manner, resistant to noise. The latter provides richer information, allowing the extraction of ω_0, κ_c^1, and $\kappa_c^2 + \kappa_i$. However, the precision of the coupling rate reading quickly deteriorates with the Signal-to-Noise Ratio (SNR), reaching its maximum when approaching the critical coupling regime, where $\kappa_c^1 \approx \kappa_c^2 + \kappa_i$. Given the expected loss rate for our design at around 100 Hz at 10 mK, designing antennas at room-temperature that would allow for this reading at cryogenic temperatures becomes a challenging task.

To calibrate the cavity antennas correctly, we begin by coupling two antennas to the cavity near the critical coupling at room temperature. This setup allows us to measure the cavity in both reflection and transmission, extracting κ_c^1, κ_c^2, and κ_i simultaneously from observations of $S_{21}(\omega)$, $S_{11}(\omega)$, and $S_{22}(\omega)$. Subsequently, one of the antennas is iteratively made shorter while monitoring the amplitude of $S_{21}(\omega_0)$, which scales as $\sqrt{\kappa_c^1}$. This precise calibration method enables us to handle extremely low values of κ_c^1, even when the resonance becomes invisible in reflection.

Since we are characterizing the bare cavity, we lack an in-situ nonlinearity that would precisely calibrate the number of photons within the cavity in its steady state. To estimate the regime of photon numbers at which the cavity operates, we rely solely on our measurement of line attenuation. Denoting P as the power entering the cavity, and n the average photon number, we can use the following equation to link them:

$$n = \frac{4P\kappa_c}{\omega_0 \hbar \kappa_{\text{tot}}^2} \tag{5.2}$$

Consequently, we measure the cavity quality factor as a function of the number of photons and fit the results to a TLS model [Burnett et al. 2017, Macha et al. 2010]:

$$\kappa_i = \omega_0 F \delta_{\text{TLS}} \tanh\left(\frac{\hbar\omega_0}{2k_B T}\right)\left(1 + \frac{n}{n_c}\right)^{-\beta} + \kappa_{\text{res}} \tag{5.3}$$

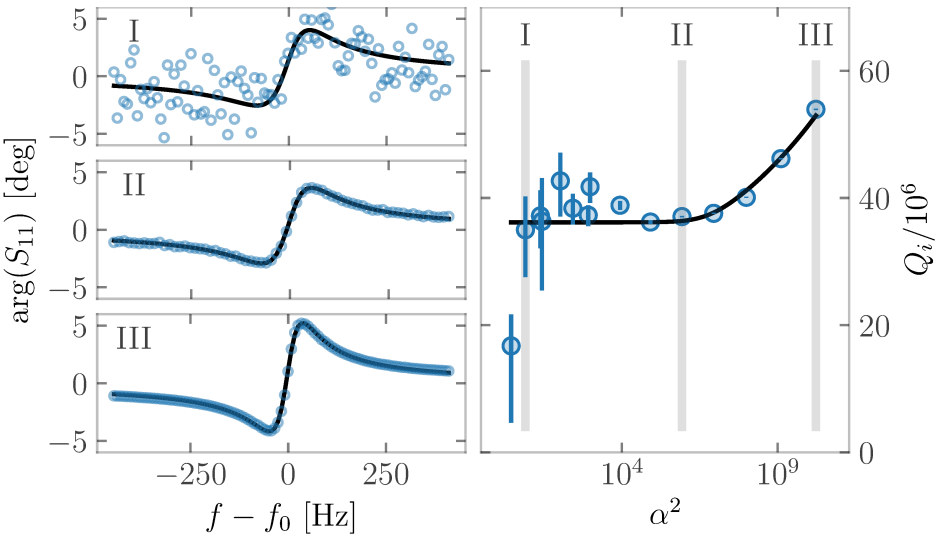

Figure 5.2 Spectroscopy of the bare stub cavity. Amplitude (top) and phase (bottom) of the reflection scattering parameter S_{11} as a function of frequency f. The data (blue open circles) closely follow a first-degree rational function (orange solid line), allowing us to extract the cavity resonance frequency f_0, the antenna coupling rate κ_c, and the intrinsic one-photon loss rate of the cavity κ_i.

Here, F represents the participation ratio of the electric field in the TLS, δ_{TLS} is the two-level system loss tangent, n_c is the typical number of photons at which the TLS saturates, and β is a fitting parameter describing the rate of saturation. Additionally, κ_{res} accounts for the residual cavity losses. Our data fitting yields $n_c = 5.6 \times 10^6$, $\beta = 0.05$, $\kappa_{\text{res}}/2\pi = 2$ Hz, consistent with values reported in the literature for 3D resonators [Kudra et al. 2020]. Importantly, we extract the cavity frequency and loss rate in the low photon number limit, obtaining $\omega_0/2\pi = 4.6$ GHz and $\kappa_i = 120$ Hz, respectively, aligning with the state-of-the-art values for this geometry. In forthcoming experiments, we will monitor how this figure of merit evolves as we couple the cavity to the rest of the experimental apparatus.

5.1.3 Filter design

Given the exceptional quality factor achieved for the 3D mode, it is crucial to avoid compromising it with subpar filter design. This discipline itself constitutes a dedicated research field with a multitude of strategies

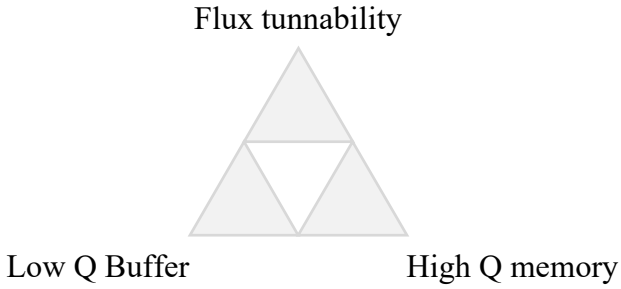

Figure 5.3 Experiment constraints diagram:

at one's disposal [Hong and Lancaster 2004, Abbas et al. 2007]. However, when dealing with an ATS, a specific requirement comes to the forefront: the ATS comprises two flux loops, both of which should be addressable with DC and AC flux. On-chip DC current is therefore needed, constraining us to implement a stop-band, as it serves three simultaneous purposes: it prevents the leakage of the memory's frequency ω_a into the thermal bath, while simultaneously allowing a substantial coupling rate κ_b for the buffer at frequency ω_b. Additionally, it allows the uninterrupted flow of DC current. These constraints were the cornerstones we had to navigate while building this experiment from the ground up, a situation we summarize by the mean of a triangular diagram Fig. 5.3.

Ensuring a good electrical contact between a 2D track and the aluminum bulk of the cavity requires the use of indium bumps [Lei et al. 2020], a technology that was not available in our lab at the time. This limitation compelled us to adopt a microstrip geometry instead of a CPW (coplanar waveguide) configuration, which would have required transporting the ground defined by the 3D cavity bulck to an on-chip plan.

The implemented stop-band needed to be a few MHz wide, as we estimated that cavity frequencies could shift by approximately this amount due to mechanical tolerances. To achieve this, we opted for a design consisting of three $\lambda/4$-stub resonators separated by a quarter wavelength [Matthaei et al. 1963, p. 758] Fig. 5.4. To save space, the stubs are integrated into the strip lines. At resonance, each stub functions as a second-order stop-band filter. By slightly offsetting their frequencies, a large stop-band filter can be engineered.

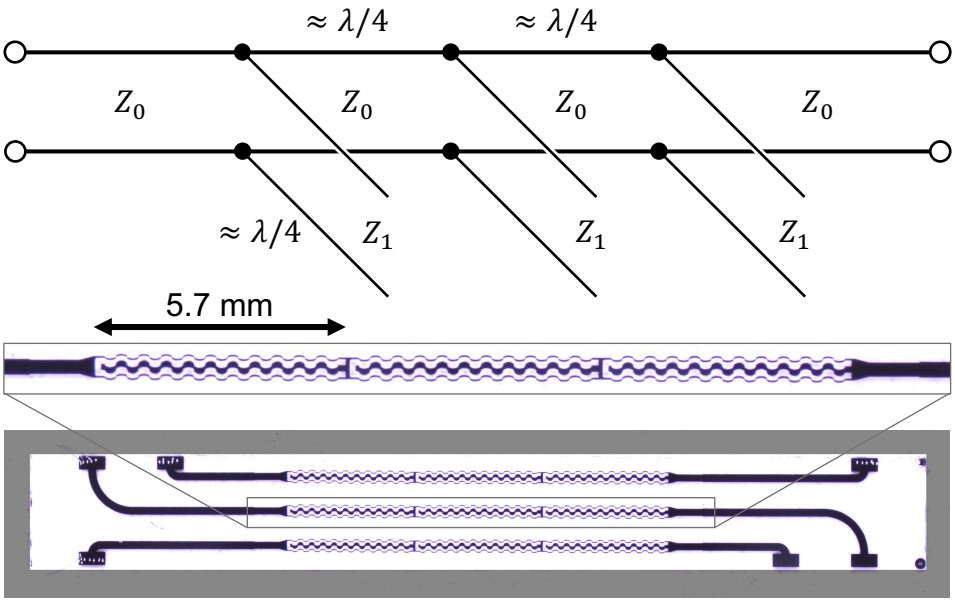

Figure 5.4 Filter sample: A photograph of the filter test sample is presented (bottom panel). Three niobium tracks run horizontally from left to right, each accommodating its own filter. A closer look at the middle track (middle inset panel) reveals the three inline stubs that make up the filter. In the distributed element representation (top panel), the length of these stubs is shown to be one-quarter of the wavelength to be filtered, and their spacing follows the same proportion.

5.1.4 Filter caracterisation

To evaluate this design, we developed a specialized sample holder taylored for this specific purpose Fig. 5.5. Its primary function was to enable the direct measurement of the filter's transmission parameter, which would not be accessible in the final architecture. Specifically, the aluminum tunnels through which the filters are inserted mimic the ones that will ultimately be used in the cat-qubit experiment.

The tunnels cross section (1.2 mm × 5.5 mm) was chosen to avoid any leakage of the memory resonator hosting the cat state through TEM (Transversal ElectroMagnetic) propagating modes [Pozar 2011, p. 96]. This choice resulted in abrupt aspect ratios, as the typical length of these tunnels is about 10 cm. Carving them in the aluminium was achieved by the mean of EDM (Electrical Discharge Machining) [Abbas et al. 2007], a fabrication

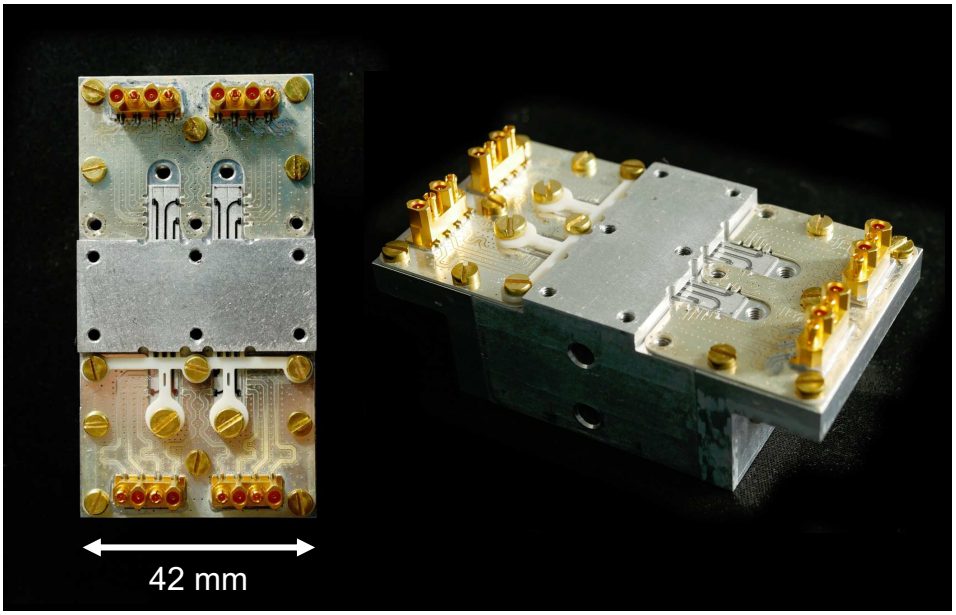

Figure 5.5 Filter sample holder. Photographs of the sample holder designed for conducting unitary tests of the filters are depicted. A top-down view (left panel) illustrates a sapphire sample passing vertically through an aluminum tunnel. The sample is connected to two PCBs using wire bonds, and held in place using a custom made Teflon piece. The side view (right panel) provides a closer look at the connectors employed to link the sample with external lines.

method that exercises almost no force on the piece during manufacturing, as it relies on electric sparks to dig into the material.

Regrettably, this technique leaves a lot of carbon residue. After extensive cleaning, inside the tunnels, it resulted in a rather rough surface state. Therefore flatly positioning the samples posed a minor mechanical challenge. To address this issue, specially crafted Teflon bridges Fig. 5.5 were created to secure in place the sapphire chips at both ends of the sample holder.

The filters themselves consist of tantalum tracks. As the final design necessitates three lines to operate the two flux loops of the ATS, the chip incorporates three tracks, each employing the same filter architecture Fig.5.4. To assess their performance, we measured the signal propagation from one end to the other at various frequencies, revealing the transmission parameter Fig.5.6.

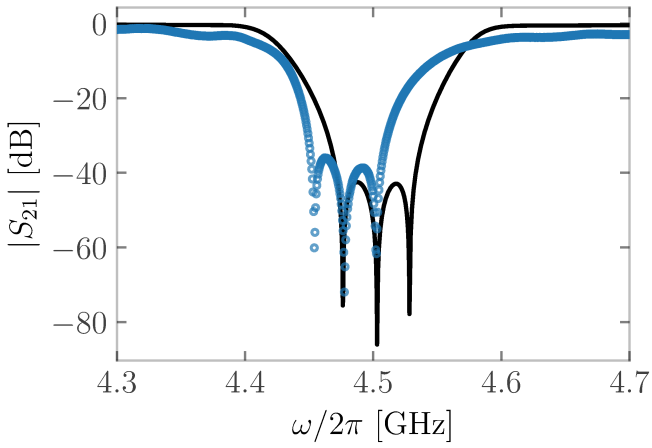

Figure 5.6 Spectroscopy of the spurline filters.

Simulations and measurements are presented Fig.5.6, showcasing a band of 50MHz where 40 dB of attenuation is achieved. However, we observe that the stop-band implementation is offset by approximately 25 MHz when compared to theoretical predictions. This deviation is likely attributed to mechanical imperfections, as multiple experimental runs yielded random offsets. The suboptimal surface finish of the tunnels may have contributed to this discrepancy, as simulations indicate a strong correlation between line impedance and the chip's relative position to the tunnel. This lack of repeatability was certainly a shortcoming of our design choices, born from unwise requirements at the conception stage. We payed the price of this unreliability during data acquisition as this randomness resulted in significant time overhead.

5.1.5 Sample holder and circuit design

We've arrived at a pivotal juncture in our journey. After painstaking theoretical studies and meticulous unitary tests, all the building blocks are in place. In a process akin to assembling a well-crafted LEGO set, we'll now delve into how these components come together. It's time to dive into the heart of the matter - an experiment we affectionately refer to as the "3D-cat".

The sample holder, approximately the size of a tennis ball, is carved from a block of high-purity 4N aluminum. Removing the covers that shield its internals reveals the experiment's structure. On one end, a deep cavity

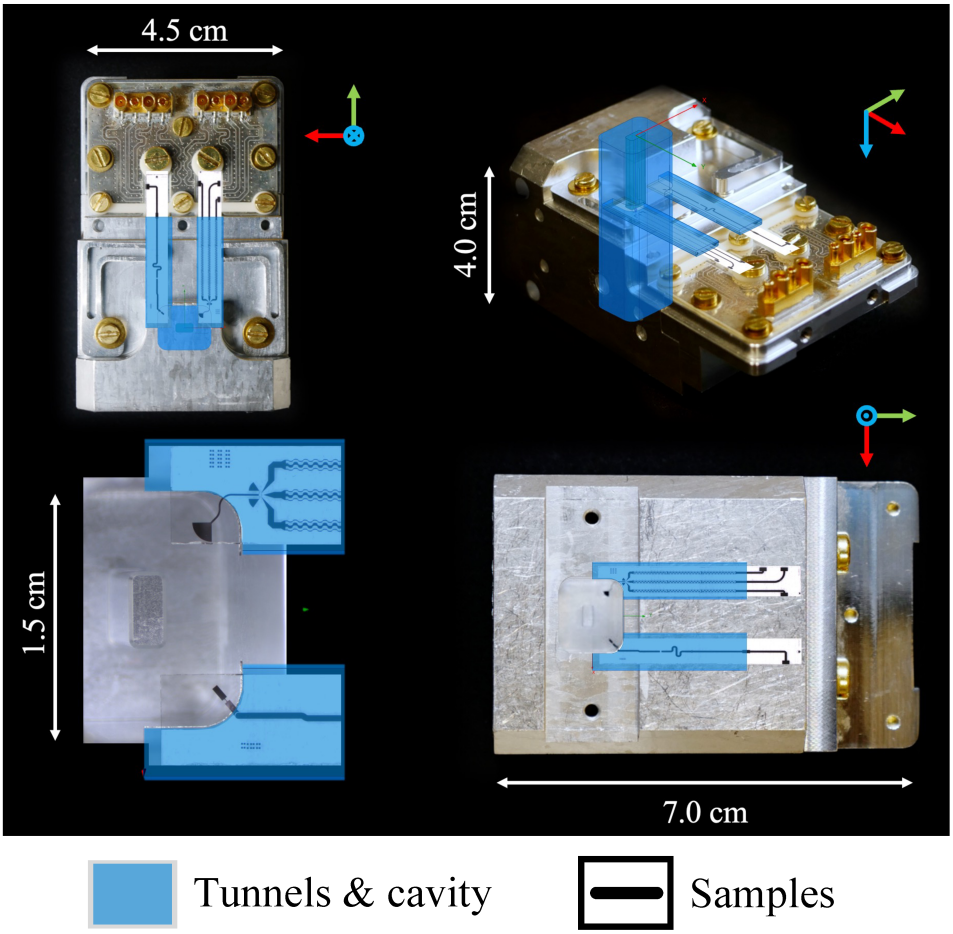

Figure 5.7 Photographs of the sample holder: An overhead view (top left) displays the PCB, which allows the connection of the device to external lines. The side view (top right) reveals two tunnels running from the PCB to the cavity, with a sample placed in both. The bottom view (bottom right) provides a direct view of the cavity, while the zoomed-in image of the cavity (bottom left) highlights two antennas extending from both samples into the cavity.

houses a high-quality-factor pole resonator, highlighted with the blue transparent overlay in Fig. 5.7. Running from the opposite side of the piece, two tunnels connect the cavity to a custom-made PCB. Nested within are two sapphire chips, securely held in position by a Teflon pin that exerts pressure to keep them firmly in place. Both samples have very different purposes.

The first sample plays a crucial role in implementing the two-photon sta-

bilization. It features a tantalum $\lambda/2$-resonator, indicated in red in Fig. 5.8. This resonator is used as the buffer and is designed like an antenna, with one end extending into the cavity to ensure efficient capacitive coupling, enhancing the memory's participation in the ATS φ^a_{zpf}. The other end of the resonator is split into two symmetrical sections, resembling 'Mickey ears'. This strategically positions the inductor near the current antinode of the buffer, ensuring strong participation φ^b_{zpf}, while allowing the ATS to be exposed. Three filtered flux lines run from there to the PCB through the tunnel, enabling the delivery of DC and AC on-chip current to both flux loops. The filters depicted in orange prevent any leakage from the memory to the thermal bath.

In its original design, the experiment features a second circuit which only purpose is enabling the reconstruction of the quantum state residing in the memory. To achieve this effect, a transmon qubit positioned within the cavity is capacitively coupled to the pole resonator, which is visible in green in Fig.5.8. Dispersively coupled to a readout $\lambda/2$-resonator, depicted in blue, the state of the transmon can be accessed. To prevent any memory leakage through this connection, the readout resonator is shielded by a $\lambda/2$-purcell filter, shown in purple. An abstract representation of the experimental principle is presented in Fig.5.9.

Throughout this chapter, we aim to present the results obtained from two years of data acquisition, spanning over a dozen designs. Each iteration of the experiment is the combination of a sample holder and one or two chips. We refer to them as assemblies. Our discussion focuses on the four that produced the most intriguing data. To maintain clarity and organization, an overview is provided in Fig. 5.10. Each assembly is given a little cartoon, and their main experimental parameters are summarized. We try to unfold the data in chronological order, and every plot displays the cartoon of the assembly from which the data has been acquired.

For instance, the first column in Fig. 5.10 presents the loss rate of the 3D cavity when samples are in place. The data shows that the quality factor of the memory can vary by approximately an order of magnitude across different assemblies. This variability stems from the inconsistent reproducibility of our mechanical fixation system for the samples, which leads to randomness in the alignment of the stop-band, as discussed in the previous section.

I should also note that we do not have precise measurements of the participations of the buffer and the memory in the central inductance of the ATS. Based on electromagnetic simulations, we can estimate that their magnitudes are approximately $\varphi^b_{\text{zpf}} \approx 0.2$ and $\varphi^a_{\text{zpf}} \approx 0.02$, respectively. However, these numbers should be taken with a grain of salt. It is safe to say

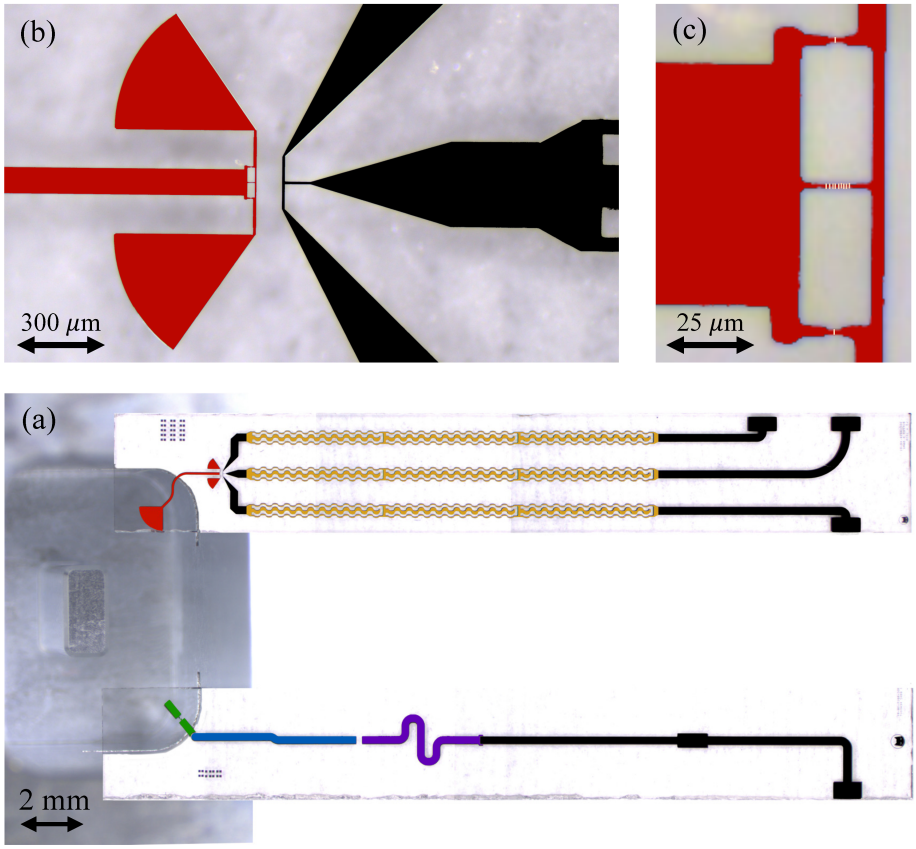

Figure 5.8 Optical micrographs of the sample. (a) A composite photograph shows the 3D cavity on the left with two sapphire chips inserted. The bottom sample features a transmon qubit (green) coupled to a readout resonator (blue), which is connected to an external line via a Purcell $\lambda/2$-filter (purple). The top sample includes the buffer mode (red), coupled to three lines using the filter design discussed in the previous section (yellow). (b) A close-up of the buffer mode reveals the inductive nature of the coupling between the ATS at the center of the buffer and the external lines. (c) A zoom-in on the ATS shows the central inductance, composed of 9 junctions in series (orange, false colors), and the two junctions on each side of the ATS.

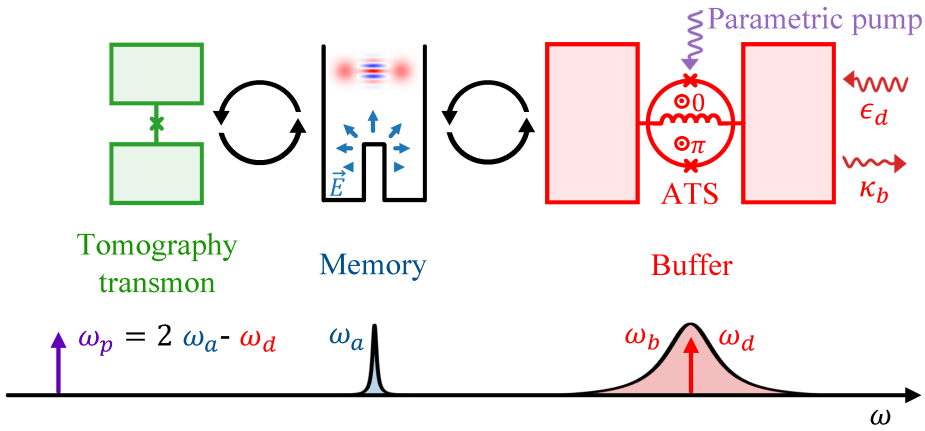

Figure 5.9 Experiment diagram of the 3D cat. The 3D cavity, referred to as the memory (blue), is capacitively coupled (solid black arrow) to a transmon qubit (green) and the buffer resonator (red), which features an ATS at its center. The transmon facilitates the tomography of the memory, while the ATS and the buffer are designed to enable the stabilization of coherent cat states within. The memory exhibits a narrow linewidth (blue Lorentzian), in contrast to the buffer, which has a very broad linewidth (red Lorentzian) due to its high loss rate. The capacitive resonant drive on the buffer (red arrows) and a parametric flux pump (purple arrows) are applied to activate the two-photon dissipation and injection mechanisms.

that, from one experiment to another, the participation of the memory can vary by approximately $\pm 50\%$, depending on the positioning of the buffer's antenna within the cavity, and the exact frequency of the buffer.

5.2 Experiment results

In its sample holder, a superconducting circuit humbly fits in the palm of the hand. Yet it is hiding a complexity the intuition may not expect. Its operation is done through tens of control knobs, and hundreds of parameters are to be estimated.

Nested in the cryostat, a place among the most remote on the Earth's surface, the device is hidden from sight and touch. Gaining progressive control over the experiment is a long process that requires successive calibrations, piecing information together, and leveraging our increasing knowledge to

	$\kappa_a/2\pi$	$\kappa_b/2\pi$	$E_J/2\pi$	$E_L/2\pi$
①  2021-05-17 | 2021-05-20	2.7 kHz			
② 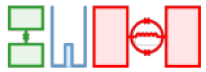 2021-08-30 | 2021-10-27	583 Hz	17 MHz	45 GHz	32 GHz
③  2022-04-04 | 2022-05-16	3.7 kHz	16 MHz	18 GHz	28 GHz
④ 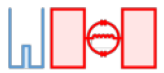 2022-07-29 | 2022-09-27	340 Hz	5 MHz	35 GHz	23 GHz

Figure 5.10 3D cat project history: Throughout the 3D cat project, more than 50 assemblies have been characterized. We focus the discussion around four, summarized with small pictograms. In the unfolding chapter, when a piece of data is displayed, the corresponding assembly pictogram is shown. The first assembly features a cavity and a transmon qubit. In the second and third, a buffer mode with an ATS at its center is added. In the fourth assembly, the transmon is removed. The loss rates of the cavity and the buffer are obtained via measurements. The energies of the ATS junctions and their inductance are obtained via room temperature measurements of their resistances.

perform measures of growing complexity.

Operating the device discussed in this chapter requires about 5×10^2 adjustable parameters. Encountering this space for the first time, it feels like a pitch black immensity of nothingness, dotted with forests of randomness hiding clearing of order. Walking in the dark, we have the foolish ambition to chart the place with a simple flashlight at our disposal. We can start everywhere, an overwhelming feeling, only cold thinking and rational analysis to comfort use.

5.2.1 The tomography transmon qubit

Monday morning, the cryostat is at 11 mK, and everything is in the green light. Taking advantage of the flexibility offered by our modular design, we exclude the ATS sample from the assembly. Later, after convincing ourselves that everything behaves nominally in its absence, the two-photon stabilizing apparatus will be added to the experiment.

For now, as we embark on a fresh measurement campaign, our journey begins with the calibration of the transmon qubit. This building block is currently favored by Google and IBM for constructing quantum computers [Arute et al. 2019] as it well understood, requires only a few incoming lines to operate and offers lifetimes in the range of 500 μs [Wang et al. 2022a]. However, in our experiment, the presence of the transmon is solely motivated by the need to perform a full quantum tomography protocol of the memory resonator. It serves as a measuring tool, functioning as an ancilla qubit that allows us to answer questions about the system's state.

The transmon is formed by connecting the two pads of a capacitor (of energy E_C) through a single Josephson junction (of energy E_J). The two pads define two isolated metallic islands. In order to model the system, the Coulomb interaction associated with their relative electric charge must be taken into account. Again, we denote $\mathbf{n}$ the number of Cooper pairs that traveled from one pad to another through the junction [Koch et al. 2007]:

$$\mathbf{H}_{\text{TRM}} = 4E_C(\mathbf{n} - n_g)^2 - E_J \cos(\boldsymbol{\varphi}) \tag{5.4}$$

In this expression, E_J is the energy of the junction defined in Eq. 4.19, and E_C the energy of the capacitor:

$$E_C = \frac{e^2}{2C} \tag{5.5}$$

The prefactor in Eq. 5.4 comes from the fact the transfer of a single Cooper pair costs 4 times the energy of the one of a single electron. n_g is called the dimensionless "charge offset" and reflects the effect of an external electrical potential applied to the system [Girvin 2012, p. 46]. Since the capacitance of a transmon is quite large and acts as antennas, this offset is subject to important noise. To overcome this parasitic effect, the transmon is operated in the regime where its dynamics is primarily governed by the energy of the junction $E_J \gg E_C$. Under this condition, the first two energy levels of the Hamiltonian become insensitive to $\mathbf{n}$ all together. This feature is the hallmark of the transmon architecture and has contributed to its remarkable success, allowing for unprecedented lifetimes shortly after its ini-

tial implementation [Schreier et al. 2008]. Introducing the ladder operators $\mathbf{t}$ and $\mathbf{t}^\dagger$ associated with $\mathbf{n}$ and φ.

$$\varphi = \sqrt[4]{\frac{2E_C}{E_J}}(\mathbf{t} + \mathbf{t}^\dagger) \quad \mathbf{n} = i\sqrt[4]{\frac{E_J}{32E_C}}(\mathbf{t} - \mathbf{t}^\dagger) \tag{5.6}$$

The Hamiltonian is well approximated by dropping the charge offset, and performing a Taylor expansion of the cosine followed by a rotating wave approximation:

$$\mathbf{H}_{\text{TRM}}/\hbar \approx \frac{\sqrt{8E_C E_L}}{\hbar}\mathbf{t}^\dagger\mathbf{t} - \frac{E_C}{12\hbar}(\mathbf{t} + \mathbf{t}^\dagger)^4 \underset{\text{RWA}}{\approx} \omega_t \mathbf{t}^\dagger\mathbf{t} + \frac{\alpha_t}{2}\mathbf{t}^{\dagger 2}\mathbf{t}^2 \tag{5.7}$$

Where $\omega_t = \sqrt{8E_C E_L}/\hbar$ is the frequency of the transmon qubit. Thanks to its anharmonicity $\alpha_t = -E_C/\hbar$, the energy spectrum of the transmon presents discreet transitions which are individually addressable. We use the lowest energy pair of states as a qubit. In our design, the transmon is dispersively coupled to a $\lambda/2$-resonator $\mathbf{r}$ at a rate $\chi_{r,t}$, allowing to perform its efficient and single shot readout. This resonator is tailored to exhibit high coupling to the drive line without compromising the lifetime of the qubit [Walter et al. 2017]:

$$\mathbf{H}/\hbar = \omega_r \mathbf{r}^\dagger\mathbf{r} + \omega_t \mathbf{t}^\dagger\mathbf{t} + \frac{\alpha_t}{2}\mathbf{t}^{\dagger 2}\mathbf{t}^2 + \chi_{r,t}\mathbf{r}^\dagger\mathbf{r}\mathbf{t}^\dagger\mathbf{t} \tag{5.8}$$

Investigating the transmon qubit begins with the spectroscopy of the readout resonator (see Fig. 5.12a), as its resonance is easily distinguishable by design. Next, we perform a two-tone spectroscopy of the transmon. During this procedure, we continuously monitor the signal reflected by the readout resonator using a weak probe while simultaneously applying a second tone to the qubit. By varying the power and frequency of this second tone, a distinctive map is constructed.

Transitions are revealed by vertical lines, widening as the power of the probe increases (see Fig. 5.11a). Beyond a certain power threshold, the overlapping the transtions becomes too important and the system behaves in a chaotic manner. Operating below this dense roof, individual transitions are addressable. Notably, due to the negative anharmonicity of the transmon, the transition between its ground state and the first excited state exhibits the highest frequency. This transition is measured to be around $\omega_t/2\pi \approx 6.51\text{GHz}$.

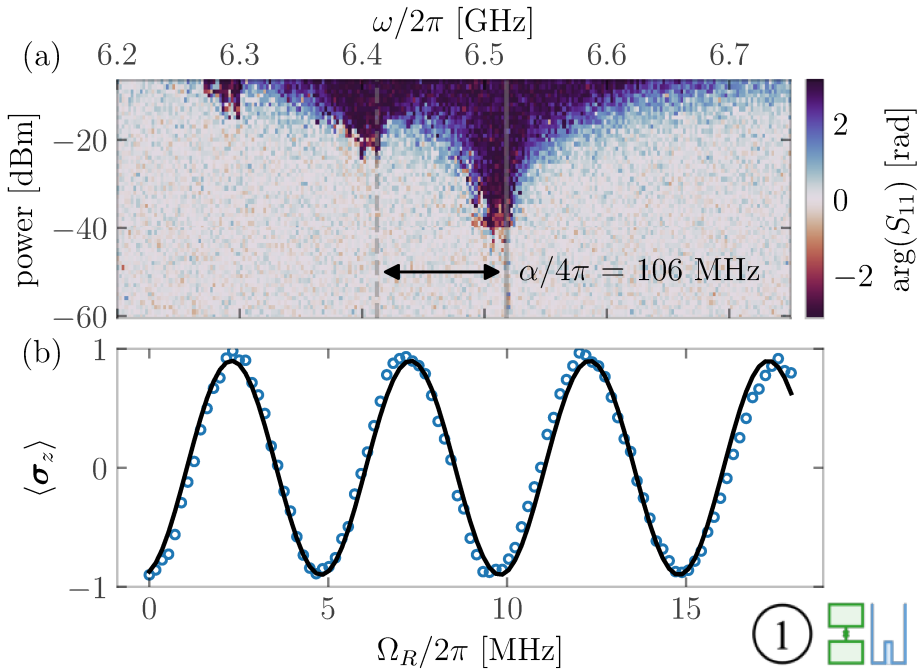

Figure 5.11 Spectroscopy and Rabi oscillations of the transmon qubit: Two-tone spectroscopy of the transmon qubit through the readout resonator (left panel) as a function of qubit drive frequency (x-axis) and qubit drive amplitude (y-axis) enables the determination of the transition frequency between the ground state and the excited state. By applying short drive pulses lasting 100 ns, Rabi oscillations (right panel) are observable in the state population (y-axis) as a function of drive amplitude (x-axis).

Focusing on this transition, a time measurement reveals Rabi oscillations (see Fig. 5.12b), the smoking gun evidence that the system operates in the quantum realm. For the sake of simplicity, we will disregard the other transitions and work within a two-level approximation. Denoting $|g\rangle$ the ground state, and $|e\rangle$ the first existed state, creation and annihilation operators are replaced by Pauli matrices. Accounting for the qubit's thermal population n_{th}^t, the coupling of the readout to the transmission line κ_r, as well as the qubit's lifetimes T_1 and its dephasing time T_φ, we can provide a comprehensive view of the operators governing the master equation:

$$
\begin{aligned}
\mathbf{H}/\hbar &= \omega_r \mathbf{r}^\dagger \mathbf{r} + \frac{1}{2}\omega_q \boldsymbol{\sigma}_z + \frac{1}{2}\chi_{r,q}\boldsymbol{\sigma}_z \mathbf{r}^\dagger \mathbf{r} \\
\boldsymbol{L}_r &= \sqrt{\kappa_r}\mathbf{r} \\
\boldsymbol{L}_{t\uparrow} &= \sqrt{n_{\text{th}}^t/T_1}\boldsymbol{\sigma}_+ \\
\boldsymbol{L}_{t\downarrow} &= \sqrt{(1-n_{\text{th}}^t)/T_1}\boldsymbol{\sigma}_- \\
\boldsymbol{L}_{t,\varphi} &= 1/\sqrt{2T_\varphi}\boldsymbol{\sigma}_z
\end{aligned}
\tag{5.9}
$$

We have essentially described the Lindblad perspective on the Optical Bloch Equations (OBE) [Raimond 2016, p. 43], with the inclusion of an additional dispersive coupling. The natural progression of our investigations involves assessing the quality of the achieved qubit. To accomplish this, we proceed with the determination of the model parameters that we've just outlined.

Revisiting the Rabi oscillations in Fig.5.11b, we can deduce the required energy and pulse duration to invert the population between the ground and excited states. Such a pulse executes a half-rotation of the transmon qubit's Bloch sphere around its X-axis and is referred to as a π-pulse. By comparing the readout of both states, we optimize the readout pulse to achieve single-shot fidelity measurements, as shown in Fig. 5.12b.

Examining the readout of the qubit prior to the π-pulse, two Gaussians distributions are visible. The transmon spends most of its time in the ground state, occasionally experiencing thermal excitations, briefly occupying its first excited state. Comparing the relative counts of these distributions enables calibration of the thermal population. The most challenging part of the work is now in the rearview mirror. The remaining calibrations of the transmon are done by navigating the Bloch sphere formed by the qubit. Entering this space for the first time, it's textbook structure gave me a fresh and immediate sensation that quantum physics was grounded in the reality. Determining the remaining parameters of the qubit is a game, with moves as rotations. Intertwining these maneuvers with pauses of variable lengths allows for the plotting of Bloch sphere coordinates over time, enabling us to benchmark the quality of our physical qubit.

Among all the pulse sequences at our disposal, one of the simplest protocols involves measuring the qubit's relaxation time T_1. Here's how it works: we start by preparing the qubit in its excited state using a π-pulse. Then, we wait for a variable amount of time before measuring the qubit's state Fig. 5.13a. By fitting the resulting curve with an exponential decay, we can quickly determine the qubit's lifetime.

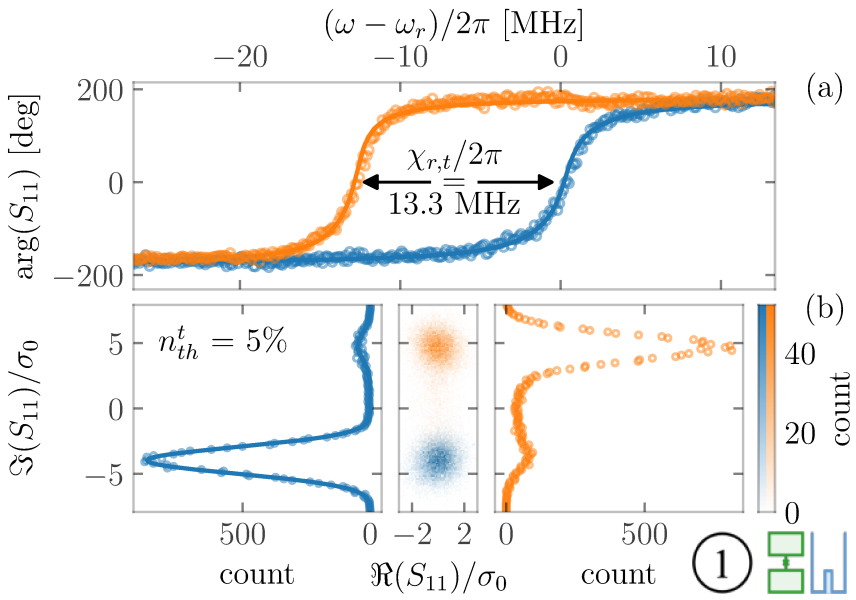

Figure 5.12 Qubit readout fidelity.(a) A spectroscopy of the readout resonator is performed with the qubit in both the ground state (orange) and the excited state (blue), allowing us to measure the qubit-readout cross-Kerr parameter, denoted as $\chi_{r,t}$.(b) Reflected signal histograms are acquired at the frequency of maximum contrast for both qubit states (center panel). When examining the real quadrature histogram with the qubit in the ground state (left panel), one can compute the thermal occupation of the excited state. Conversely, observing the same histogram with the qubit in the excited state reveals T_1 errors during the readout (right panel).

From there, we proceed with a detuned Ramsey sequence, a measure that serves a dual purpose. It provides us with two valuable pieces of information: the qubit's dephasing rate $1/T_\varphi$, and a finer measurement of the qubit frequency $\omega_t/2\pi$.

This sequence starts with the preparation of the plus state $|+\rangle$ by the mean of a slightly off-resonant $\pi/2$-pulse. Subsequently, the X-coordinate of the qubit is measured after a variable time Fig. 5.13b. During this process, we observe oscillations whose frequency is directly proportional to the detuning, enhancing our knowledge of $\omega_t/2\pi$. On the other hand, these oscillations are enveloped by an exponential decay with a rate denoted as T_2^R. Its link with the dephasing rate in obtained by solving the OBE.

$$\frac{1}{T_2^R} = \frac{1}{2T_1} + \frac{1}{T_\varphi} \tag{5.10}$$

In our experimental setup, we typically achieve values of approximately $T_1 \approx 25\ \mu s$ and $T_\varphi \approx 12\ \mu s$. While these numbers may not be considered state-of-the-art [Wang et al. 2022a], they are well-suited for our specific purposes. Notably, we do not invest extensive time in gate optimization, as the appropriate metric for this task is provided by a randomized benchmarking [Barends et al. 2014], which can be a bit tedious and is not the heart of our experiment.

However, this is not to say that more data should not be extracted from the ancilla transmon. In particular, very insightful information can be obtained by performing an echo Ramsey sequence. Simulating its outcome based on Eq. 5.9, an exponential decay is found. On paper, its characteristic time T_2^E should not differ from T_2^R. In practice, a higher value is almost always measured.

The reason for this discrepancy arises from the fact that the Markovian hypothesis, which underlies the master equation formalism, is not entirely accurate in real life. If the noise spectrum responsible for dephasing is primarily low-frequency in nature, then the echo protocol partially mitigates its effects. Going in the limit case where the noise is exclusively low frequency, the effect of the dephasing operator is entirely cancelled, but the finite value of the qubit lifetime still imposes an hard limit on the exponential decay observable with this sequence protocol:

$$T_2^E \leqslant 2T_1 \tag{5.11}$$

As the transmon qubit is a rather fragile system, it can also be seen as a good probe of its environment quality. As keen readers may already know, throughout the various stages of signal transmission, from room temperature equipment to the near absolute zero of the superconducting circuit, filtering plays a pivotal role. Among these stages, one of the most crucial involves filtering out infrared radiation. To accomplish this, in-house Eccosorb filters [Paquette et al. 2022] are manufactured in the lab. These filters are positioned between the sample holder and the external lines, operating at the same temperature as the circuit itself. Fig. 5.17 presents data from another measurement campaign. A comparison of the transmon qubit's lifetimes is performed with and without these filters in place, with all things aside equal. Looking at this table, the necessity of these filters becomes

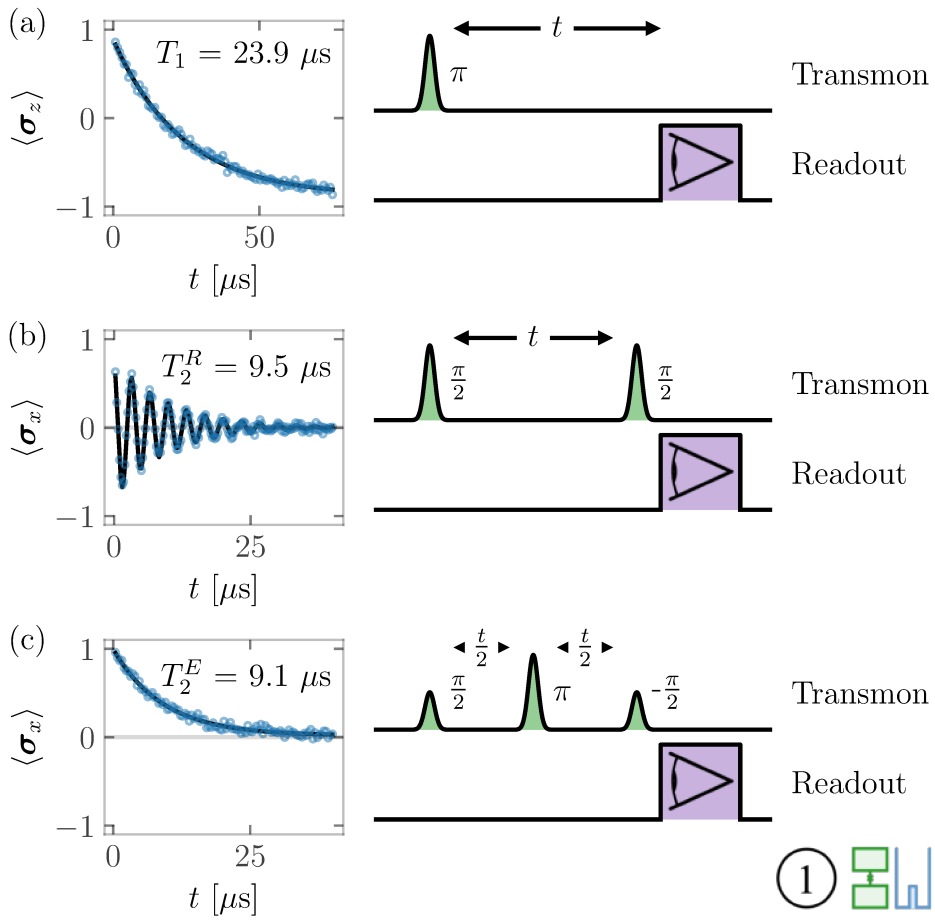

Figure 5.13 Transmon qubit life times. (a) Measure of the qubit lifetime. After a preparation in the excited state, the population (y-axis) decays toward the ground state as time goes (x-axis). (b) Detuned Ramsey sequence. In this case, the X-axis of the Bloch sphere is displayed (y-axis) versus time (x-axis). Oscillations with an exponential decay toward the center of the Bloch sphere are observed, allowing to calibrate the detuning and the dephasing rate. (c) Resonant Ramsey echo sequence, akin to (b) but with resonant $\pi/2$-pulses and population inversion in the middle. Comparison of the exponential rates with (a) aids in determining the nature of the noise spectrum, whether low or high frequency.

self-explanatory, as the qubit performs better on every conceivable metric with them in place.

5.2.2 Dispersive coupling between the cavity and the transmon

As stated earlier, we intend to use the transmon as a measuring apparatus to reveal the state of the field confined within the 3D-cavity. Before introducing further complexity to the setup, it is crucial to ensure that the coupling between the two systems manifests as expected.

The transmon interacts with the electric field of the memory mode through the pad diving within (see Fig. 5.8a), effectively creating a coupling capacitance between the two. Given the significant detuning their natural frequencies ω_a and ω_t, we are in the dispersive limit of the interaction. This regime is well described by a Hamiltonian of the form $\frac{1}{2}\chi_{a,t}\boldsymbol{\sigma}_z\mathbf{a}^\dagger\mathbf{a}$, where $\chi_{a,t}$ is referred to as the cross-Kerr. Taking into account the lifetimes, the dephasing rates, and the thermal populations of both, we have:

$$\mathbf{H}/\hbar = \omega_a\mathbf{a}^\dagger\mathbf{a} + \frac{1}{2}\omega_t\boldsymbol{\sigma}_z + \frac{1}{2}\chi_{a,t}\boldsymbol{\sigma}_z\mathbf{a}^\dagger\mathbf{a} \tag{5.12}$$

$$\begin{aligned}
\boldsymbol{L}_{a\uparrow} &= \sqrt{n^a_{\text{th}}\kappa_a}\mathbf{a}^\dagger & \boldsymbol{L}_{t\uparrow} &= \sqrt{n^t_{\text{th}}/T_1}\boldsymbol{\sigma}_+ \\
\boldsymbol{L}_{a\downarrow} &= \sqrt{(1+n^a_{\text{th}})\kappa_a}\mathbf{a} & \boldsymbol{L}_{t\downarrow} &= \sqrt{(1-n^t_{\text{th}})/T_1}\boldsymbol{\sigma}_- \\
\boldsymbol{L}_{a,\varphi} &= \sqrt{\kappa_{a,\varphi}}\mathbf{a}^\dagger\mathbf{a} & \boldsymbol{L}_{t,\varphi} &= 1/\sqrt{2T_\varphi}\boldsymbol{\sigma}_z
\end{aligned}$$

We are now interested in characterizing this model. Since the memory is extremely narrow band, finding its frequency is tedious. In addition, the antenna that dives in the cavity is now undercoupled, in order to maximize its quality factor - a direct spectroscopy cannot be made. This is why a two-tone spectroscopy approach is employed. While varying the frequency of a tone sent through the antenna of the memory, selective π-pulses are applied to the transmon. Let us detail the intuitive picture that motivates this procedure.

A pulse is said selective when it is long in the time domain, which means its spectral envelope is narrow. Therefore, for the transmon to be affected by a selective pulse, its frequency must be well adjusted - to a precision that is inversely proportional to its length. By sweeping the frequency of a selective π-pulse (see Fig. 5.14b), we can obtain a high-resolution image of

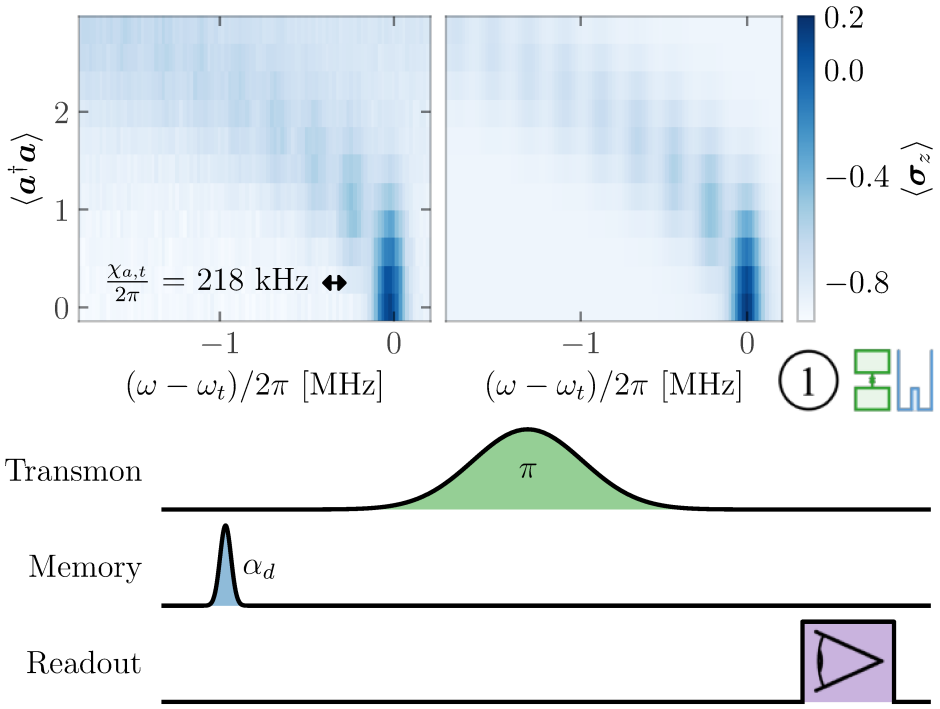

Figure 5.14 Number Splitting of the Memory. (a) The occupation of the transmon (color) is depicted versus the detuning of the pulse applied to it (x-axis) and the amplitude of the displacement applied to the memory through its antenna (y-axis). Blue spots are visible, indicative of the individual occupations of the memory Fock states. Their separation directly measures the cross-Kerr between the memory and the transmon. (b) Pulse sequence of the measurement. After a short Gaussian displacement pulse sent to the memory, a selective π-pulse is performed on the transmon qubit, followed by its state measurement.

the transmon's spectral response. Now, as described in [Schuster et al. 2007], thanks to the cross-Kerr effect, the frequency of the transmon is shifted by the population of the memory mode (see Fig. 5.14a). Therefore, when we send a pulse through the antenna that is resonant with the memory, it gets populated, and the selective π-pulse applied to the transmon becomes ineffective. By monitoring the response of the transmon, the frequency of the memory mode can be accurately determined.

One of the most striking predictions of quantum physics is the discretization of the energy levels of the harmonic oscillator, an effect we can now directly observe. By applying a short resonant Gaussian pulse through the antenna, a coherent state is injected into the memory, which takes the shape of a Poisson distribution over the Fock state basis Eq. 2.4. sweeping the frequency of the selective π-pulse applied to the transmon allows to interrogate the occupation of each Fock state. In the color plot shown in Fig. 5.14a, each blue spot corresponds to one of them. Separated thanks to the cross-Kerr effect, their occupation varies according to Eq. 2.4 as the amplitude of the memory displacement increases. Solving the master equation for Eq. 5.12, a full simulation of the situation is performed. Very good agreement between theory and experiment is observed, leading to a first calibration of $\chi_{a,t}$, as it can be directly red from the separation between the spots in this picture.

As Fock states can be individually observed thanks to this protocol, it is colloquially referred to as "number splitting" of the memory. To further confirm our understanding of the system is correct, it is also useful to introduce the time-domain analogous of this measurement. We execute a resonant Ramsey sequence on the transmon, preceded by a displacement of the memory Fig. 5.15. Aligned on the equator of the Bloch sphere, individual Fock states can be envisioned as arrows rotating around the Z-axis at rates that are integer multiples of the cross-Kerr. Expanding on this depiction, we can anticipate the occurrence of revivals every $t = 2\pi n/\chi_{a,t}, n \in \mathbb{N}$. Given the expression of a coherent state Eq. 2.4, and neglecting the losses of the memory, the analytical expression of the protocol outcome follows naturally from this picture.

$$\langle \boldsymbol{\sigma}_x(t) \rangle = e^{-|\alpha|^2} \sum_{n=0}^{\infty} \frac{\alpha^{2n}}{n!} e^{-t/T_2^R} \cos\left(tn\chi_{a,t}\right) \tag{5.13}$$

Once more, Fig. 5.15a showcases a textbook agreement between experimental results and simulations. To the extent of our requirements, it is therefore proven that the current setup behaves as expected. It is time to move on.

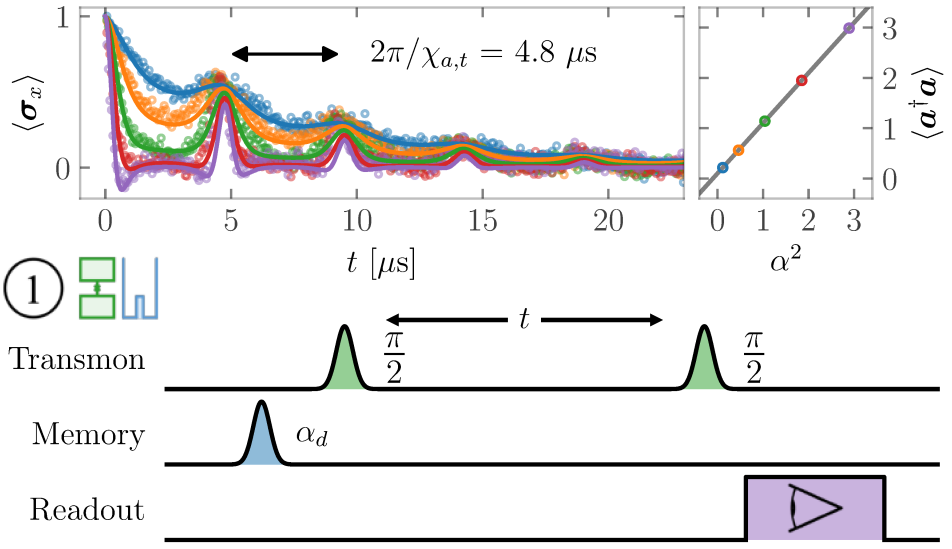

Figure 5.15 Revivals of the Memory. (a) A Ramsey sequence is performed on the transmon (left panel). The X quadrature of the Bloch sphere (y-axis) is plotted as a function of time (x-axis). As the amplitude of the displacement applied to the memory increases (color), revivals evenly spaced appear within an envelope delineated by an exponential decay. Their spacing is inversely proportional to the cross-Kerr coupling strength. A precise fit of the data allows calibration of the number of photons in the memory (right panel). (b) Pulse sequence of the measurement. After a displacement of the memory, a resonant Ramsey sequence is performed, followed by measurement of the transmon qubit.

5.2.3 Flux point calibration

The last ingredient is added to the experiment: a chip hosting the buffer resonator (see Fig. 5.8b) with an ATS at its center (see Fig. 5.8c) is introduced into the second tunnel of the sample holder. Installed in the cryostat before the weekend, this new assembly has been cooling down for two days. With the complexity of the device now increased, deciding where to begin is not a trivial matter. Since the memory, the buffer, and even the transmon modes contribute to the flux going through the central inductor of the ATS, their frequencies are highly sensitive to the flux threading through its loops. Hence, our initial task is to set the ATS to the correct flux point and maintain

it throughout the measurement campaign.

As mentioned in the previous section, the flux point that maximizes the two-photon exchange rate g_2 corresponds to a saddle point of the buffer mode frequency flux-map. This operating point also minimizes the amplitude of processes involving the conversion of an odd photon number in the memory, an essential feature for parity conservation. Its location can be achieved through various methods. One relatively expedient approach involves examining the signal reflected by a probing tone sent to the buffer at a constant frequency. By plotting the phase of the reflected signal against the DC current injected to bias the ATS, we generate a 2D map (see Fig.5.16a), akin to a slice of the ATS flux map (see Fig.7.4a). Abrupt changes in color signify the buffer mode frequency is traversing the isofrequency defined by the probing tone.

One of the greatest nightmares for an experimentalist is the prospect of cooling down a sample only to find it non-functional. Between the junctions' probing stage and insertion into the cryostat, even the smallest electrostatic disturbances can irreparably damage Josephson junctions, rendering days or weeks of painstaking work futile. Fortunately, the periodicity observed in the common and differential flux of the obtained map signals the health of the device.

In practice, the axes of the raw 2D-map we obtain may not be perfectly orthogonal. This discrepancy arises from the fact that the lines used to address the left and right loops of the ATS affect both. However, by fitting the axes of the flux map, we derive a 2 by 2 current-to-flux matrix. Inverting it allows precise control over the flux passing through each loop of the ATS.

By zooming in on the spot where the isofrequency lines almost touch on the flux map and adjusting the frequency of the probing tone, we can locate a saddle point, characterized by an X pattern (see Fig. 5.16b). Setting the DC bias precisely at the center of this X and sweeping the frequency of the probing tone allows us to acquire a spectroscopy of the buffer resonance. By fitting its shape, we can recover the value of the buffer one-photon loss rate κ_b (see Fig. 5.16c).

The flux point of the ATS can be operated, and the loss rate of the buffer is on paper large enough to enable efficient two-photon stabilisation in the memory mode. Therefore, two out of the three constraints of our little triangle are satisfied (Fig. 5.3). However, this was the easier part. Now, we must ascertain that the quality factor of the memory mode remains unaffected. If its degradation is too significant, it will be impossible to enter a regime where $\kappa_a \ll \kappa_2$, which would impede our ability to measure coherent cat states.

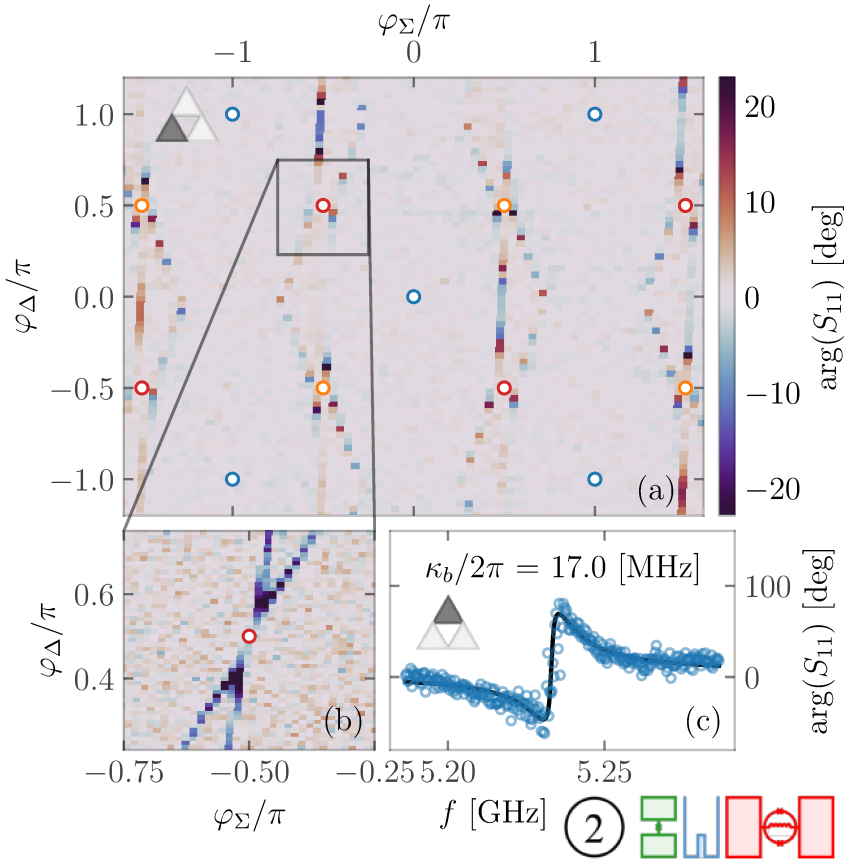

Figure 5.16 Buffer Flux Map and Spectroscopy. (a) A probing tone of constant frequency is sent to the buffer, and the phase of the reflected signal (color) is displayed against common (x-axis) and differential (y-axis) flux. The zero flux points (blue circles) as well as two classes of saddle points (orange and red circles) are indicated. The map is periodic in φ_Σ and φ_Δ, and the triangle diagram (top-left) indicates that the flux controllability constraint from Fig.5.3 is met. (b) Inset zooms in on the saddle point, where the frequency of the probing tone is adjusted to reveal the crossing of two isofrequencies. (c) Spectroscopy of the buffer mode at the saddle point. The phase of the signal (y-axis) is plotted against frequency (x-axis). We find a one-photon loss rate of approximately 17MHz. The second constraint of Fig. 5.3 is satisfied, indicated by another triangle diagram (top-left).

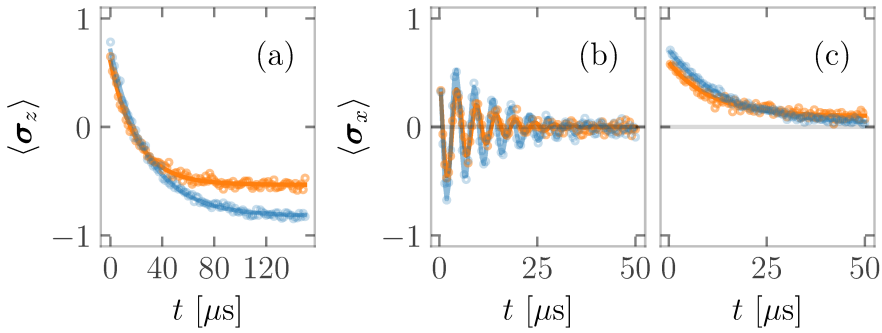

(d)	n_{th}^t	T_1	T_2^R	T_2^E
no ecosorb filter	23.4%	20.4 μs	9.3 μs	10.9 μs
ecosorb filter	8.9%	28.7 μs	13.5 μs	14.4 μs

Figure 5.17 Effect of Eccosorb filter on qubit lifetime. The same sample is measured twice, once with the Eccosorb filter (blue) and once without (orange). (a) Measurement of qubit lifetime. (b) Detuned Ramsey sequence. (c) Resonant Ramsey echo sequence. (d) Summary table recapping the lifetimes of the transmon and its thermal population. Incorporating the Eccosorb filter enhances performance across all parameters.

5.2.4 Wigner tomography of the memory mode

After the flux point of the ATS is set, we calibrate the transmon qubit again. However, we skip this step as it has already been detailed extensively, and a summary of the transmon properties for this measurement campaign can be found in the blue entry of the table in Fig. 5.17d.

Next, we proceed again with the measure of $\chi_{a,t}$. This time around, we perform a resonant Ramsey sequence on the qubit where no coherent state is injected into the memory. As depicted in Fig. 5.18, subtle oscillations persist, serving as clear evidence of thermal occupation of the 3D-cavity. In order to perform precise photon number calibration, this population must be measured. Assuming the memory has reached thermal equilibrium, its state is a decoherent mixture of Fock states which spreading can be retrieved from its effective temperature:

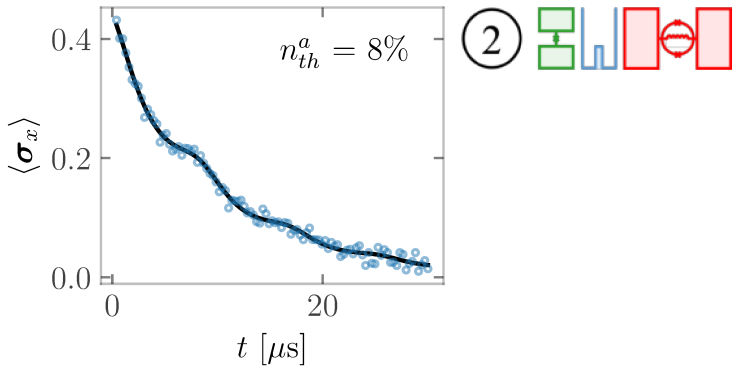

Figure 5.18 Thermal occupation of the memory. A resonant Ramsey sequence on the transmon reveals on close inspection the appearance of slight oscillations in the data (open circles). This time, no displacement is applied to the memory, therefore their presence is a smoking gun evidence of the thermal occupation of the memory. By fitting these data (solid black line) with Eq. 5.15, its value can be computed.

$$\boldsymbol{\rho} = \sum_{n=0}^{\infty} \frac{(n^a_{\text{th}})^n}{(1+n^a_{\text{th}})^{n+1}} |n\rangle\langle n| \tag{5.14}$$

Taking inspiration from Eq. 5.13, we can produce the analytical expression for the shape of the revivals that are induced by this thermal population during the Ramsey sequence:

$$\langle \boldsymbol{\sigma}_x(t) \rangle = \sum_{n=0}^{\infty} \frac{(n^a_{\text{th}})^n}{(1+n^a_{\text{th}})^{n+1}} e^{-t/T_2^R} \cos\left(tn\chi_{a,t}\right) \tag{5.15}$$

Assuming that the occupation of Fock states is accurately captured by Eq. 5.14, the temperature of the memory mode can be determined by fitting the data Fig. 5.18.

We will now explain how the dispersive coupling between the transmon and the memory can be employed to measure the Wigner function of the field hosted in the cavity. Looking at the resonant Ramsey Fig. 5.15, the revivals are striking features. For quantum state tomography however, one should look passed their catchiness as the most useful information is hidden precisely in between them. Writing the value output by the Ramsey sequence at $t = \pi/\chi_{a,t}$ reveals its link with the parity operator.

$$\langle \boldsymbol{\sigma}_x(\pi/\chi_{a,t})\rangle = e^{-\pi/\chi_{a,t}T_2^R}\sum_{n=0}^{\infty}(-1)^n \langle n|\,\boldsymbol{\rho}\,|n\rangle \propto \langle e^{i\pi \mathbf{a}^\dagger \mathbf{a}}\rangle \tag{5.16}$$

Direct detection of the Wigner function becomes feasible, as a very non trivial fact is that its value aligns precisely with that of the displaced parity operator [Lutterbach and Davidovich 1997].

$$\frac{\pi}{2}W(\gamma) = \langle \mathcal{D}^\dagger(\gamma) e^{i\pi \mathbf{a}^\dagger \mathbf{a}} \mathcal{D}(\gamma)\rangle \tag{5.17}$$

In this expression, we deliberately maintain the $\pi/2$ on the left to maintain consistency with the parity operator, which ranges between -1 and 1. While this equality is fundamental and can be formally proven, a deeper understanding of the underlying principle eludes the writer. As such, we refrain from further elaboration and accept it as a foundational concept for the remainder of the text. Historically, one of the earliest direct measurements of the Wigner function utilizing this relation occurred in the context of a resonant cavity coupled to traveling Rydberg atoms [Bertet et al. 2002]. Later, it was implemented in the context of CQED to acquire the full 2D image of the Wigner function [Vlastakis et al. 2013].

Reflecting on Fig.5.15a, one can envision colors as representing a dimension, allowing the vacuum Gaussian state to be sensed between the first two revivals. To unveil it in its full splendor, the duration between the two $\pi/2$-pulses is fixed, while the amplitude of the initial displacement is varied, as illustrated in Fig.5.19a. Pixel by pixel, one can reconstruct a photographic representation of the phase space. Initially sparsely populated, it hosts a solitary wave function resembling a Gaussian centered within. Later, it will be our task to shape it in a beautiful cat. For now, we are satisfied with its observation.

To avoid any systematic offset in the measure of the Wigner, the sign of the second $\pi/2$-pulse in the Ramsey sequence is alternated. Two signals are produced; subtracting them reveals a slice of the thermal vacuum state. Starting with Eq. 5.14, one can compute its corresponding Wigner distribution [Curtright et al. 2013].

$$\frac{\pi}{2}W(\gamma) = \frac{1}{1+2n_{\text{th}}} e^{-2|\gamma|^2/(1+2n_{\text{th}})} \tag{5.18}$$

In the absence of thermal population (i.e., when $n_{\text{th}} = 0$), only the vacuum state $|0\rangle$ is occupied, and the Wigner function is a Gaussian with a

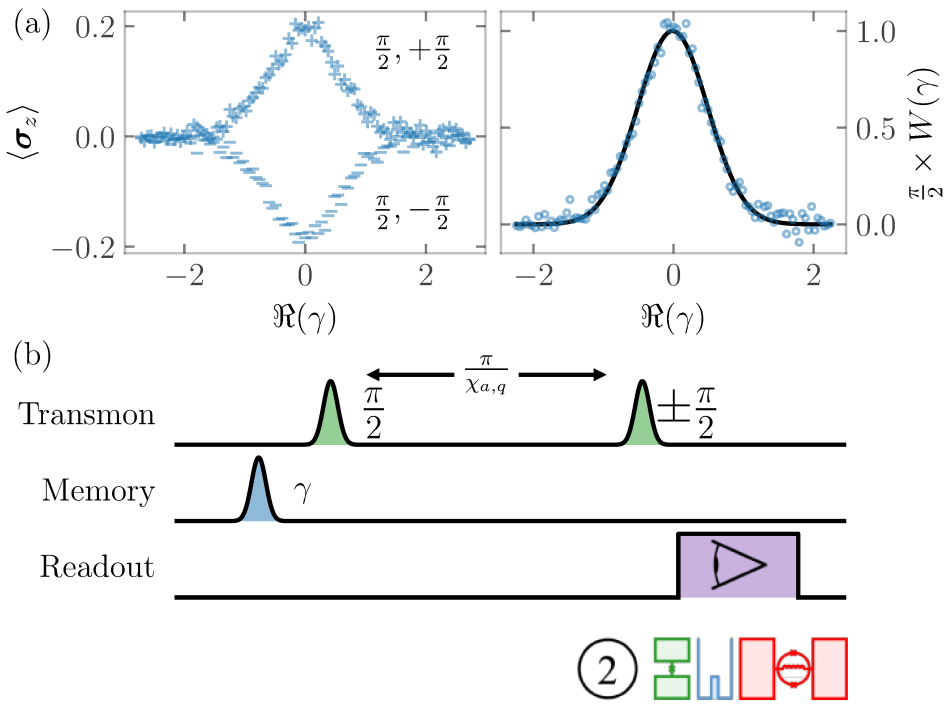

Figure 5.19 Wigner function of the vacuum state. (a) A Ramsey sequence preceded by a displacement of the memory is performed, where the time between the two $\pi/2$-pulses is set to $\pi/\chi_{a,t}$. For each displacement amplitude (x-axis), the sign of the second $\pi/2$-pulse is alternated (plus and minus blue markers). The resulting signals (y-axis) are displayed in their raw data form (left panel). Subtracting one from the other, the value of the displaced parity operator can be assessed (right panel). (b) Pulse sequence of the measurement.

standard deviation of $1/2$. However, when there is non-zero thermal occupation, the state spreads. Its Wigner representation still corresponds to that of a Gaussian but with a larger standard deviation. By fitting the data with Eq. 5.18 given a prior knowledge of the thermal population in the memory, we can calibrate the amplitude of the displacement pulses. Ultimately, this allows us to plot the Wigner function of the memory state with axes given in units of square root photon number.

With that, we finally are able to resolve the Wigner tomography of memory mode. Its quantum state is now measurable. In a fashion somewhat similar to what we did with the transmon, we are now going to assess the quality of our resonator by monitoring its response over time after different pulses.

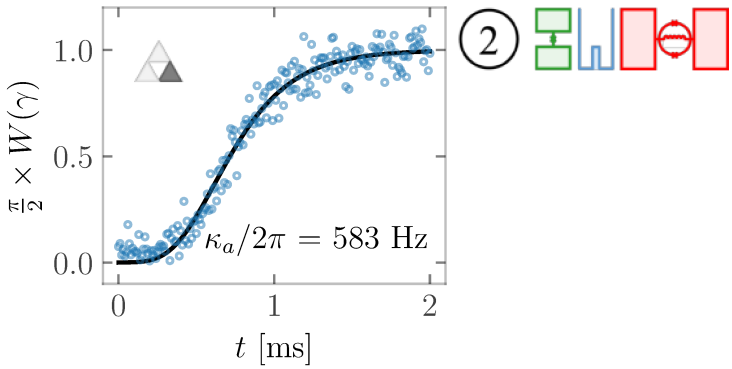

Figure 5.20 Lifetime of the memory. Starting from a displaced thermal state, the value of the parity operator (y-axis) is monitored over time (x-axis). Data (blue open circles) are compared to theory (solid black line). We find that the one-photon loss rate of the memory mode is about 600 Hz. The triangle diagram (top-left) indicates that the third constraint of Fig. 5.3 is satisfied.

For instance, we can displace the memory by an amplitude α_0, effectively preparing a displaced thermal state, and monitor its decay. Since the amplitude of the field goes as $e^{-\kappa_a t/2}$, we can deduce an analytical expression for the parity operator over time.

$$\langle \mathbf{a}^\dagger \mathbf{a}(t) \rangle = \frac{1}{1+2n_{\text{th}}^a} \exp\left(\frac{-2}{1+2n_{\text{th}}^a} \alpha_0^2 e^{-\kappa_a t}\right) \tag{5.19}$$

In this formula, everything is known, apart from the one-photon loss rate κ_a which appears in the double exponential. Performing this measurement allows to retrieve its value Fig. 5.20. The effective loss rate we find is about 0.6 kHz, a four fold degradation compared to its bare quality factor Fig. 5.2, yet significantly greater than the typical 5 kHz achieved with 2D resonator geometry. The quality factor we achieve is about $Q = 8 \times 10^6$, giving us a lifetime of 265 μs. In a fashion analogous to a two-level system, κ_a is not the only quantity characterizing the quality of a mode, and the dephasing rate $\kappa_{a,\varphi}$ is also to be calibrated.

Assuming the dynamics is uniquely governed by the action of the dephasing operator $\boldsymbol{L}_{a,\varphi}$, an initial displaced state gets spread over a circle. The dispersion over the circle grows over time at a rate given by $\kappa_{a,\varphi}$, and its angle follows a Gaussian distribution.

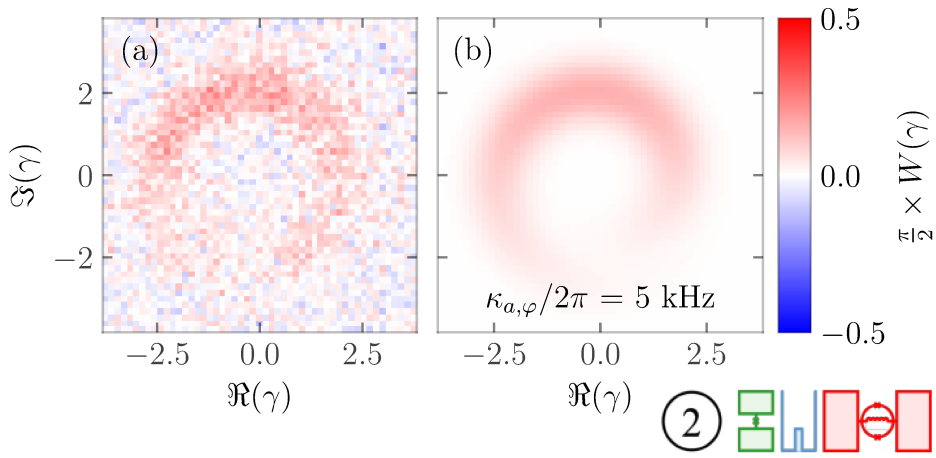

Figure 5.21 Dephasing rate of the memory. (a) Starting with a displaced thermal state, the Wigner tomography (color) is performed after 50 μs. Under the effect of the dephasing induced by flux noise, the state spreads across a circle in the phase space (x and y axis). (b) A numerical simulation shows the dephasing rate is about 5 kHz.

After a fraction of a millisecond, the 2D image of the Wigner function is acquired Fig. 5.21. By the mean of a numerical simulation, the spread of the initially coherent state is reproduced when the dephasing rate is about $\kappa_{a,\varphi}/2\pi = 5$ kHz. This figure probably finds its origin in the flux noise to which the ATS is submitted. As the coupling obtained thanks to the antenna extending from the buffer mode in the cavity is important, any frequency drift of the buffer is repercuted onto the memory, leading to dephsaing. Through the cross-Kerr effect, the thermal occupation of the transmon can also be a source of dephsaing. However, the action of this parasitic term remains local in the phase space. In a fashion similar to the one-photon loss operator, its action should be exponentially suppressed by the two-photon stabilisation. And we should still be able to recover an exponential scaling of the bit-flip time.

5.2.5 Calibration of the two-photon exchange

After considerable effort, we are finally poised to stabilize a Schrödinger cat in the 3D-cavity. Employing the ATS as a 4-wave mixer, we must carefully select the frequencies ω_p of the parametric pump applied to the common flux φ_Σ and ω_d of the capacitive drive sent to the buffer for two-photon

injection into the memory. Although the frequencies of the memory and buffer are known from prior calibrations, satisfying the frequency matching condition Eq. 4.30 is not without its complications.

When the parametric pump is enabled, the frequency of the buffer and the memory get dressed, effectively shifted from their natural frequency by an effect known as AC-Stark shift [Leghtas et al. 2015].

Two approaches are possible. The optimal method involves sending additional drives to the memory and buffer at the pump frequency to destructively interfere with this shift. However, due to its complexity, we opt for a more expedient solution. When the pump is active, we simply observe the frequency shift of the buffer and the memory and adjust accordingly.

Fig. 5.22 presents a two tone spectroscopy of the buffer. In each panel, as the frequency of the pump is swept (y-axis), spectroscopies of the buffer are made (x-axis) and the amplitude of the reflected signal is displayed (color). When the drive probing tone hits the buffer $\omega_d = \omega_b$, The resonance appears as a dark vertical line. When the amplitude of the pump tone ε_p is small (upper panels), the resonance frequency remains largely unaffected. However, when we increase the amplitude (down panels), the response of the buffer gets distorted to the right by the AC-Stark shift.

Upon closer examination of these color plots, fine features emerge. Several lines intersect with the buffer's resonance, indicating the impact of resonant parametric processes on the buffer's content. By fitting line equations to these features, we can infer the underlying wave mixing phenomena. Among the observed lines, one stands out, associated with the two-photon exchange $(\omega_p + \omega_d)/2 = \omega_a$ (orange). I was unable to identify intercepts for the other three lines, leaving their associated photon conversion processes unclear. To cite only a few possible mechanisms, these processes could potentially involve the $3\lambda/4$ mode of the 3D cavity or even a TLS residing in the substrate of the buffer chip.

Now, let us zoom on the two-photon exchange line. Referring back to the analysis in Eq. 2.40 from the second chapter, we intuit that when this process is resonant, the buffer becomes depopulated. Hence, in the color plot Fig. 5.22, the two-photon process manifests as a line with lighter color passing through the dark resonance of the buffer. To create a cat state in the memory, our goal is to select the point (ω_d, ω_p) where the two-photon line intersects with the center of the buffer resonance.

However, determining this point is contingent upon the pump amplitude ε_p, and it is initially unclear how to select it based solely on this two-tone spectroscopy. In practice, making this choice is more akin to an art than a science, often involving numerous trials and errors. Our objective is to

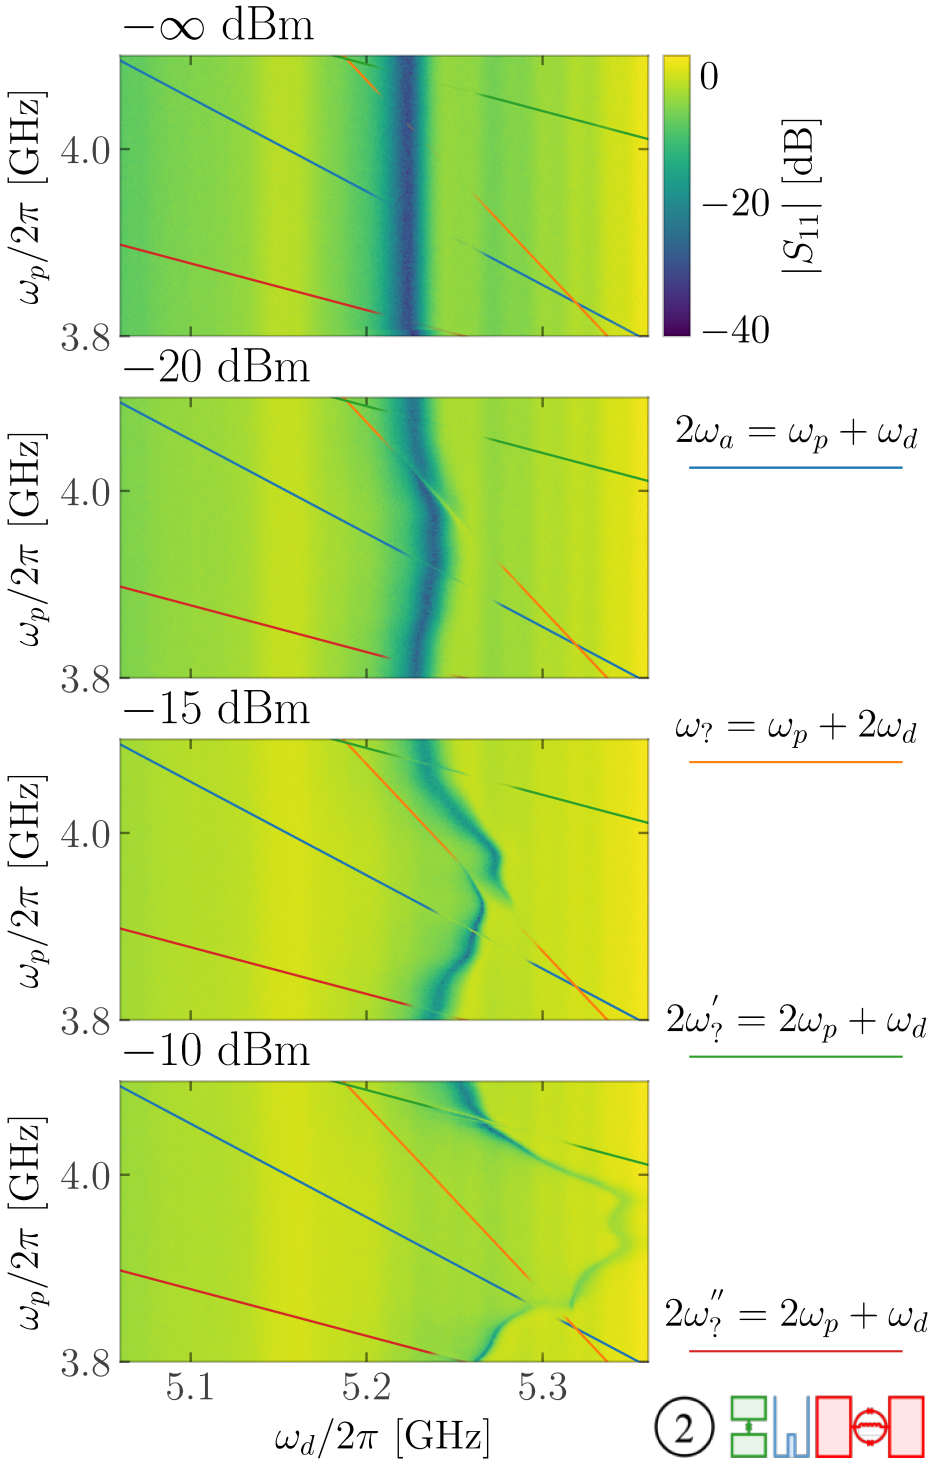

Figure 5.22 Two-tone spectroscopy of the buffer.

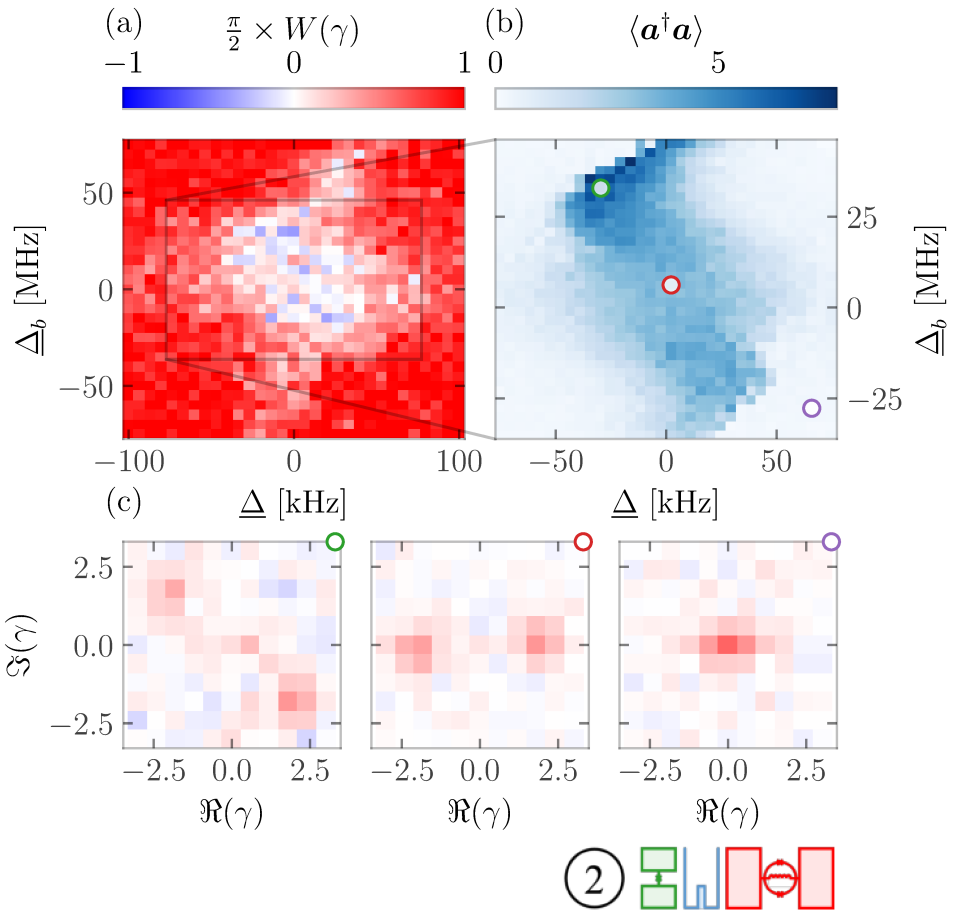

Figure 5.23 Two-photon anti-crossing from Wigner tomography. (a) As we go through several frequencies for the pump and the drive in the coordinate system defined in Eq. 5.20 (x and y axis), a Wigner tomography of the memory is made, and the value of the parity operator is displayed (color). (b) In the same coordinate system, the number of photons in the steady state (color) is computed from a double Gaussian fit. (c) The Wigner of three data points are displayed.

maximize the pump amplitude to achieve a high two-photon exchange rate g_2. However, excessively high ε_p can lead to unexpected effects such as heating or degradation of the lifetime. A practical guideline is to operate within a regime where the AC-Stark shift remains manageable.

After deciding on a pump amplitude, the pump and drive frequencies must be chosen precisely. However, in Fig. 5.22, the two-photon exchange

feature is extremely narrow. This is why, following [Lescanne et al. 2020b], we introduce a coordinate system which axes are parallel and orthogonal to the two-photon exchange line.

$$\begin{aligned} \underline{\Delta} &= \frac{\omega_p + \omega_d}{2} - \omega_a \\ \underline{\Delta}_b &= \frac{\omega_p - \omega_d}{2} - \omega_a + \omega_b \end{aligned} \tag{5.20}$$

This transformation allows us to navigate the spectral landscape more effectively. Now, instead of focusing solely on the response of the buffer, for each point $(\underline{\Delta}, \underline{\Delta}_b)$, we conduct a low-resolution Wigner tomography of the memory mode in its steady state. With this piece of data, we can for instance generate a 2D map of the parity operator, revealing a characteristic diamond shape as shown in Fig. 5.23a. In this representation, the vacuum state of the memory appears as a red pixel, while the stabilization of a decoherent cat state outputs white. Intuitively, we aim to operate the system at the center of this diamond.

Furthermore, these data allow us to extract the number of photons in the steady state, as illustrated in Fig. 5.23b, which exhibits excellent agreement with theoretical predictions depicted in Fig. 7.2. Remarkably, we find that the center of the diamond corresponds to a saddle point of the photon number. Examining its corresponding Wigner function unveils a low-resolution depiction of the decohered cat state.

I vividly remember the first time I saw the two red spots of a decoherent cat state appearing on my screen. It was well past 10 PM in my flat, where I was operating the experiment remotely. My smart lights bathed my bedroom in a soothing shade of red, and I felt as though I were piloting a submarine in the serene depths of the ocean. My excitement was immense, and I eagerly anticipated sharing this milestone with others.

Scientific literature is the output product of researchers' efforts. A well-crafted paper is appealing and comprehensible, often requiring rearrangement of the narrative flow in hindsight. However, this approach overlooks the inherent human nature of research, hiding the underlying decision-making processes that gave birth to the paper in the first place. In the following paragraphs until the end of this section, I will shift the narrative to the past tense, aiming to immerse the reader in the decision-making context we experienced at the time.

After several cooldowns, the moment of truth arrived: were we able to prepare a coherent Schrödinger cat state simply by enabling two-photon injection and dissipation? As the pump and drive tones were activated, we

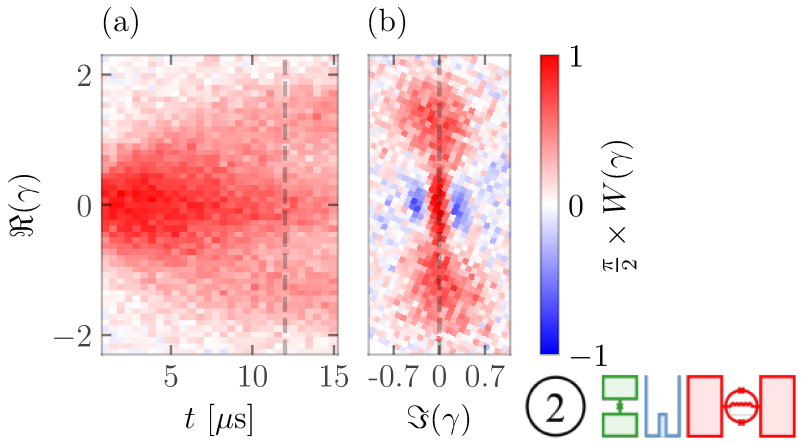

Figure 5.24 Squeezing a cat out of the vacuum. (a) Starting from the vacuum state, the two-photon drive is enabled. The value of the Wigner function (color) is displayed against time (x-axis) and the real axis of the phase space (y-axis). The Initial Gaussian splits in three: two coherent states that compose the Shrodinger cat state $|\pm\alpha\rangle$, as well as its interference pattern which lies in the middle. (b) Plotting the Wigner function (color) in the phase space coordinates (x and y-axis) 12 μs after we started the two-photon injection, the image of a coherent cat appears on the screen.

recorded a movie of the memory state. After more than a year of hard work, the Wigner function of a cat state with magnificent negativities appeared on the screen (see Fig. 2.4). Out of a purely dissipative effect, we produced a quantum superposition. This was not the first time that reservoir engineering had been used experimentally as a quantum resource. In [Kienzler et al. 2015], for example, with trapped ion qubits, it was already demonstrated that squeezed states could be obtained. However, I was struck by the peculiar shape of the coherent cat state, and I was extremely enthusiastic about this result, but I think I didn't realize how important it was. However, the joy was short-lived.

As we progressed with finer and finer characterizations, more and more problems emerged. First, the rate at which the fringes of our cat states faded was too significant. As discussed earlier, the phase-flip rate should be given by $\kappa_{\text{pf}} = 2\kappa_a\alpha^2$. Inverting this formula using the data from Fig.5.25, we found an effective one-photon loss rate of the memory κ_a^{eff} of about 4.4kHz, a seven-fold degradation compared to the direct measurement of its lifetime

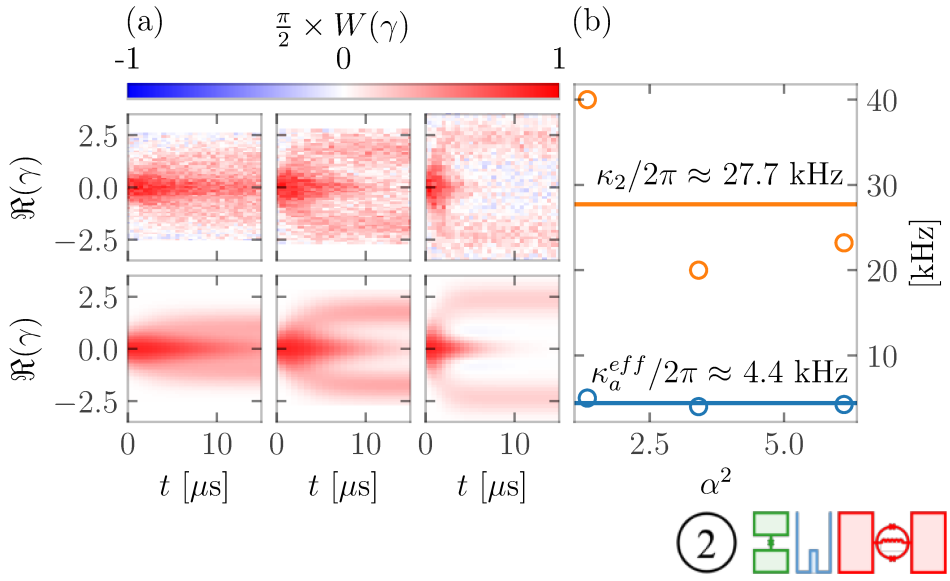

Figure 5.25 Estimating two-photon coupling from cat state inflation. (a) Initially prepared in the vacuum state, the evolution of the memory is shown after activating the two-photon stabilization. The time evolution (x-axis) of the Wigner function (color) is displayed along the real axis of the phase space (y-axis). The experiment is repeated for values of $\alpha^2 = 1.7, 3, 6$. Data (top) and simulations (bottom) are illustrated. (b) Effective values of the two-photon dissipation rate (orange) and the one-photon loss rate (blue) are derived from the simulations.

(Fig. 5.20). To this day, we have no clear explanation for this discrepancy. The figure of merit $\kappa_2/\kappa_a^{\text{eff}} \approx 6.3$ was greater than 1, allowing us to observe coherent cats, but it was still disappointing compared to the expected value of $27.3/0.6 \approx 46$, which should have been possible thanks to the extensive effort put into filtering.

However, the greatest disappointment came when we characterized the scaling of the bit-flip time (Fig 5.26). Initially, it appeared exponential, but it saturated extremely early, around 20 photons, at $T_{\text{bf}} = 20$ ms. While this was a first for this technology at the time, it was extremely underwhelming considering the platform's potential to achieve "macroscopic" bit-flip times.

At this exact moment, we made a mistake. We could have published these results, but our judgment was obscured. Alice&Bob was under immense pressure to demonstrate the reliability of the technology on which

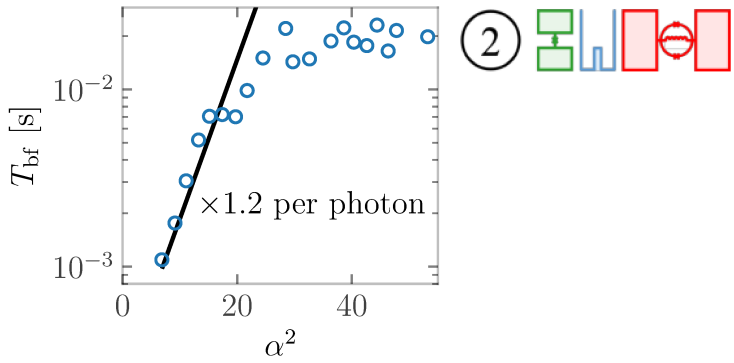

Figure 5.26 Scaling of the bit-flip time with respect to cat size. The bit-flip time (y-axis) is shown as a function of the number of photons in the steady state (x-axis). For a photon number smaller than 20, the scaling looks exponential. However, after this point, it quickly saturates around 20 ms.

the company's promises were built. Consequently, we chose to continue iterating on the design of the experiment rather than collecting more data. Looking back on it now, this decision seems shortsighted. Ultimately, on the most promising sample of the 3D cat project, I have less than a few week's worth of data, which is ridiculous. We should have taken a step back, taken a vacation for one or two weeks, and returned with clear minds and fresh ideas.

Meanwhile, in the lab, another experiment was yielding exceptionally promising results. On a superconducting 2D geometry, we achieved a dissipative a bit-flip time exceeding one minute, a remarkable feat. Consequently, I decided to shift my focus to a supportive role, performing data analysis for the team on this sample. What set this sample apart was it was lacking a transmon qubit. This was deliberate, a choice made in an effort to simplify the design. However, in this context, it was unclear how to make the tomography of the memory mode. My contribution involved proposing a protocol for in situ calibration of the photon number, which led to a subsequent publication [Berdou et al. 2023]. Unfortunately, the exceedingly low two-photon exchange rate in this sample prevented the preparation and observation of coherent cat states. The robustness of the two-photon dissipation was demonstrated, but what we had was not a qubit.

Our mistake regarding the 3D cat project was not in failing to publish; it was failing to capitalize on the modularity offered by the 3D cat experiment.

Taking inspiration from [Berdou et al. 2023], we should have removed the transmon from the equation! In hindsight, I am still baffled by our oversight. All the pieces of the puzzle were there, and we had a working sample. At the time, however, we were blinded by the necessity of performing the Wigner tomography to characterize our quantum system, and this option was not even considered.

On the other hand, parallel to our experimental work, a theoretical proposal developed. It suggested the apparition of dynamical instabilities [Burgelman et al. 2022] in our system, which we thought could be responsible for the degradation of the lifetimes under activation of the pump. Applied to the ATS, this theory suggested that the ratio $E_J/E_L \approx 1.4$ we had chosen was too aggressive, and that cranking it down could solve our problems.

Unfortunately, the outcome of this choice was exceedingly disheartening. As far as I can tell, reducing E_J did not improved the effective one-photon loss rate of the memory under pump κ_a^{eff}. On the contrary, since g_2 is proportional to E_J, it resulted in a significant reduction of the two-photon loss rate κ_2 in subsequent iterations of the experiment. To compound matters, when I returned later to remeasure the sample that had produced coherent cat states, it had been damaged by ElectroStatic Discharges (ESD). The combination of junction fabrication unreliability and the poor repeatability of the mechanical fixation system within the 3D cavity prevented us from ever measuring a coherent cat state again over the course of this project.

About one year later, another crucial piece of evidence convinced us that removing the tomography transmon qubit was the right decision. While measuring the bit-flip time of my cat qubit, I opted to continuously monitor the state of the transmon. Specifically, the state $|+\alpha\rangle$ was prepared in the memory, and after a period t, the populations in $|+\alpha\rangle$ and $|-\alpha\rangle$ were recorded. During the interval t, the state of the transmon qubit was measured every $3\ \mu$s. From these observations, I was able to distinguish the raw bit-flip time of my cat qubit from the one calculated when trajectories showing the transmon leaving the ground state were excluded. In Fig. 5.27, these data points are shown in orange and blue, respectively. The conclusion is clear: excitations of the transmon induced bit-flips.

5.2.6 Removing the tomography transmon qubit

Since it has been established that the transmon is likely detrimental to the stability of the system, our plan of action is as follows: We are going to remove the transmon qubit sample from its tunnel. Then, to gather information about the state of the memory, the antenna coupled to the 3D cavity

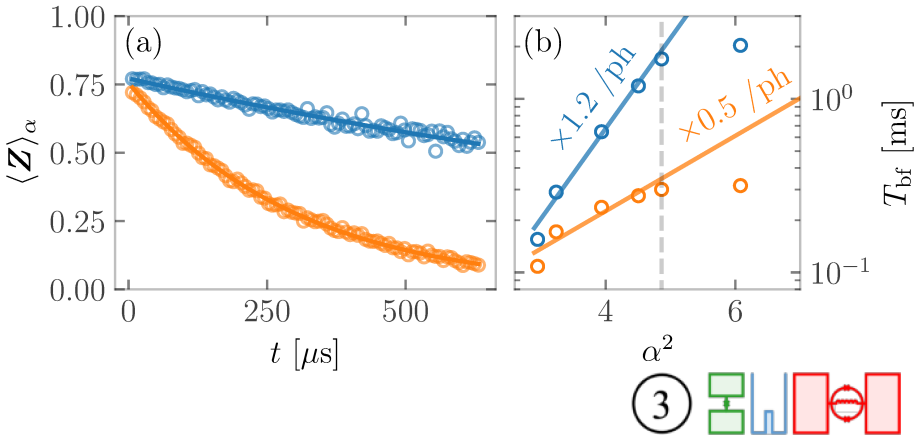

Figure 5.27 Scaling of the bit-flip rate conditioned on the transmon being in the ground state. (a) The state $|+\alpha\rangle$ is prepared in the memory (with somewhat low fidelity, as at $t = 0, \langle \mathbf{Z} \rangle_\alpha \approx 0.75$ instead of 1). As time progresses (x-axis), the Z quadrature of the cat qubit code space is measured (y-axis). Every 3 μs, the state of the transmon qubit is recorded for each measured trajectories. Orange circles represent raw data, while blue circles denote data obtained by excluding trajectories where the transmon qubit exits its ground state. An exponential fit (solid lines) is used to calculate the bit-flip time in both cases. (b) The measurement is repeated for different cat amplitudes (x-axis), and the resulting bit-flip times are displayed (y-axis). The data point in the left panel corresponds to the dashed vertical line. By post-selecting, the exponential scaling (indicated along the solid colored lines) is enhanced, and the saturation of the bit-flip time increases more than sixfold.

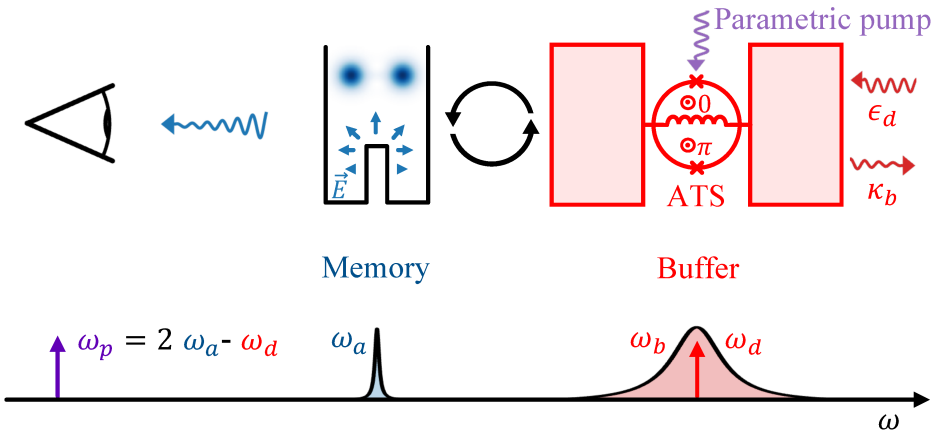

Figure 5.28 Transmon free experiment diagram. Similar to the setup in Fig. 5.9, the 3D cavity is capacitively coupled to the ATS sample. However, in this configuration, the tomography transmon qubit is absent. To measure the state of the cavity, an antenna is critically coupled to it, enabling the heterodyne detection of the field that leaks through.

will be adjusted to achieve critical coupling. This adjustment will allow us to measure the field of the memory mode directly via fluorescence, using heterodyne detection, without significantly degrading its lifetime. We summarize this situation in the diagram shown in Fig. 5.28.

Once again, the ATS is set to the frequency saddle point, and the frequency of the parametric pump and the capacitive drive are calibrated. However, this time, the calibration is performed over a wide range of pump amplitudes. By activating the drive and the pump, we may be able to stabilize cat states in the 3D cavity, but measuring them now presents an extremely challenging task.

The heterodyne detection of the field radiated by the memory results in two signals, denoted I and Q, which are integrated over an integration time T_m. These signals can be essentially envisioned as noisy measurements of the real and imaginary axes of the memory phase space:

$$\begin{aligned} I_t &= \sqrt{G}\int_t^{t+T_m} \left(\sqrt{2\kappa_a^c\eta}\,\mathrm{Tr}(\boldsymbol{\rho}(\tau)(\mathbf{a}+\mathbf{a}^\dagger)/2)d\tau + dW_I\right) \\ Q_t &= \sqrt{G}\int_t^{t+T_m} \left(\sqrt{2\kappa_a^c\eta}\,\mathrm{Tr}(\boldsymbol{\rho}(\tau)(\mathbf{a}-\mathbf{a}^\dagger)/2i)d\tau + dW_Q\right) \end{aligned} \tag{5.21}$$

In this formula, G is the gain of the amplification chain, κ_a^c is the coupling rate of the antenna to the memory mode, and η is the quantum efficiency of our detection setup. $dW_{I/Q}$ represents a Wiener process that introduces noise to the time traces, importantly verifying $dW_{I/Q}^2 = dt$. The statistics of the distribution of the (I, Q) pairs collected over time provide insights into the state of the memory. A key point, though difficult to prove, is that the quantity $\overline{I^2 + Q^2}$, where the overline represents a statistical average, reflects the number of photons in the memory. Assuming momentarily that $\boldsymbol{\rho}$ is a coherent state, and applying this assumption to Eq. 5.21, we derive:

$$\overline{I^2 + Q^2} = 2GT_m + 2G\kappa_a^c \eta T_m^2 \langle \mathbf{a}^\dagger \mathbf{a} \rangle \tag{5.22}$$

Now, it happens that this assertion generally holds true [Barchielli and Gregoratti 2009, Tilloy 2018] when the integration time T_m approaches zero. In this equation, the first term can be calibrated by measuring the output signal from the memory when it is in the vacuum state, yielding a measurement that is proportional to the steady-state number of photons. After this calibration, we can stabilize cat states of varying sizes in the strongly dissipative regime, that is to say $\kappa_2 \ll \kappa_a$, and observe the memory response. In Fig.5.29a, the blue curve illustrates a pronounced dissipative phase transition. Using the methodology outlined in Sec.2.1.2, we can calibrate the axes of the plot. As the pump amplitude is increased, the visibility of the dissipative phase transition decreases, yet we can extrapolate the value of the pump amplitude ε_p and the calibration of the photon number in the memory remains accurate.

For the highest pump power achievable with this system (15 dBm, as shown in Fig.5.29b), we can now analyze time traces of I and Q for various drive amplitudes. This process is akin to conducting a continuous measurement along the Z axis on the Bloch sphere in the cat qubit code space. The resulting plots illustrate how the system toggles between $|+\alpha\rangle$ and $|-\alpha\rangle$ over time. Fig.5.30 presents four trajectories corresponding to increasing cat sizes ($\langle \mathbf{a}^\dagger \mathbf{a} \rangle = 11.3, 14.5, 23.0,$ and 149.0) with an integration time of $T_M = 2$ ms. The left panels display histograms of the (I, Q) pairs, where we can recognize two blobs, while the right panels show the evolution of the I quadrature (y-axis) over time (x-axis). Clear bit-flip events are visible. By measuring the average time each state persists before a bit-flip occurs, we can determine the bit-flip time T_{bf} directly from the data as half this quantity. The last trajectory, lasting two hours, shows no visible bit-flips, demonstrating the system's extreme robustness. To this date, I think it is a record for dissipative cat qubits.

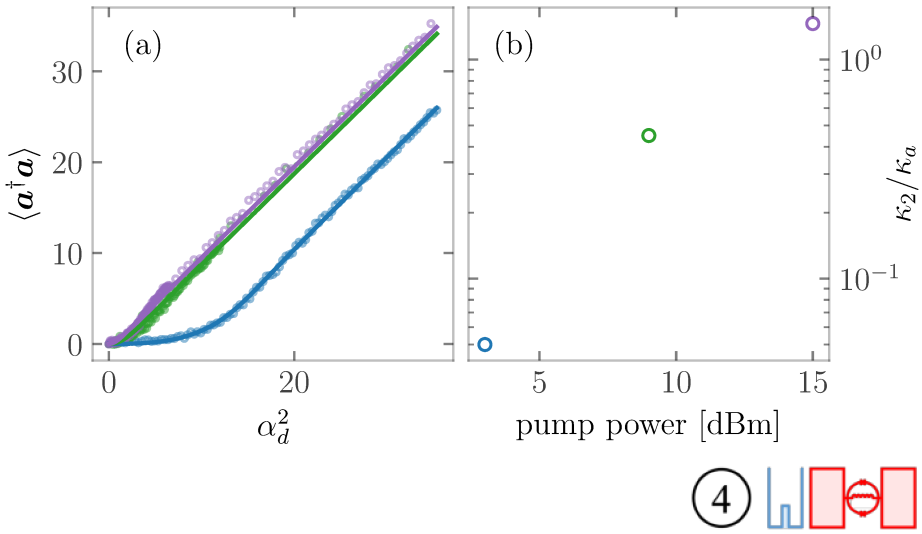

Figure 5.29 Cat size as a function of applied drive amplitude for different pump powers. (a) For various pump powers (color), the drive amplitude (x-axis) is swept. The field leaking from the cavity through the antenna is collected via heterodyne detection, and the average value of I^2+Q^2 (open circles) is displayed (y-axis). After subtracting the offset caused by amplification chain noise, a signal proportional to the steady-state photon number is obtained. Using the calibration procedure described in the second chapter Sec. 2.1.2, the axes of the plot are calibrated, and a fit is produced (solid lines). (b) The resulting values for the ratio κ_2/κ_a (y-axis) are plotted against the pump amplitude (x-axis).

Now, by adjusting the pump amplitudes, we can measure how the bit-flip time scales with the photon number in the steady state. As shown in Fig. 5.31, increasing the pump amplitude extends the saturation point of the bit-flip time, enabling the achievement of macroscopic lifetimes.

Compared to [Berdou et al. 2023], the bit-flip time we achieved extends well beyond a minute, a significant improvement that warrants explanation. We can hypothesize several mechanisms that could explain this increase. First, for the two-hour time trace exhibiting no bit-flips in Fig. 5.30, we estimate that the ratio κ_2/κ_a is just under 2. In [Berdou et al. 2023], this figure of merit is only 0.07.

Additionally, there may be reasons, not entirely clear to me, why the 3D configuration provides better protection against energetic events. In se-

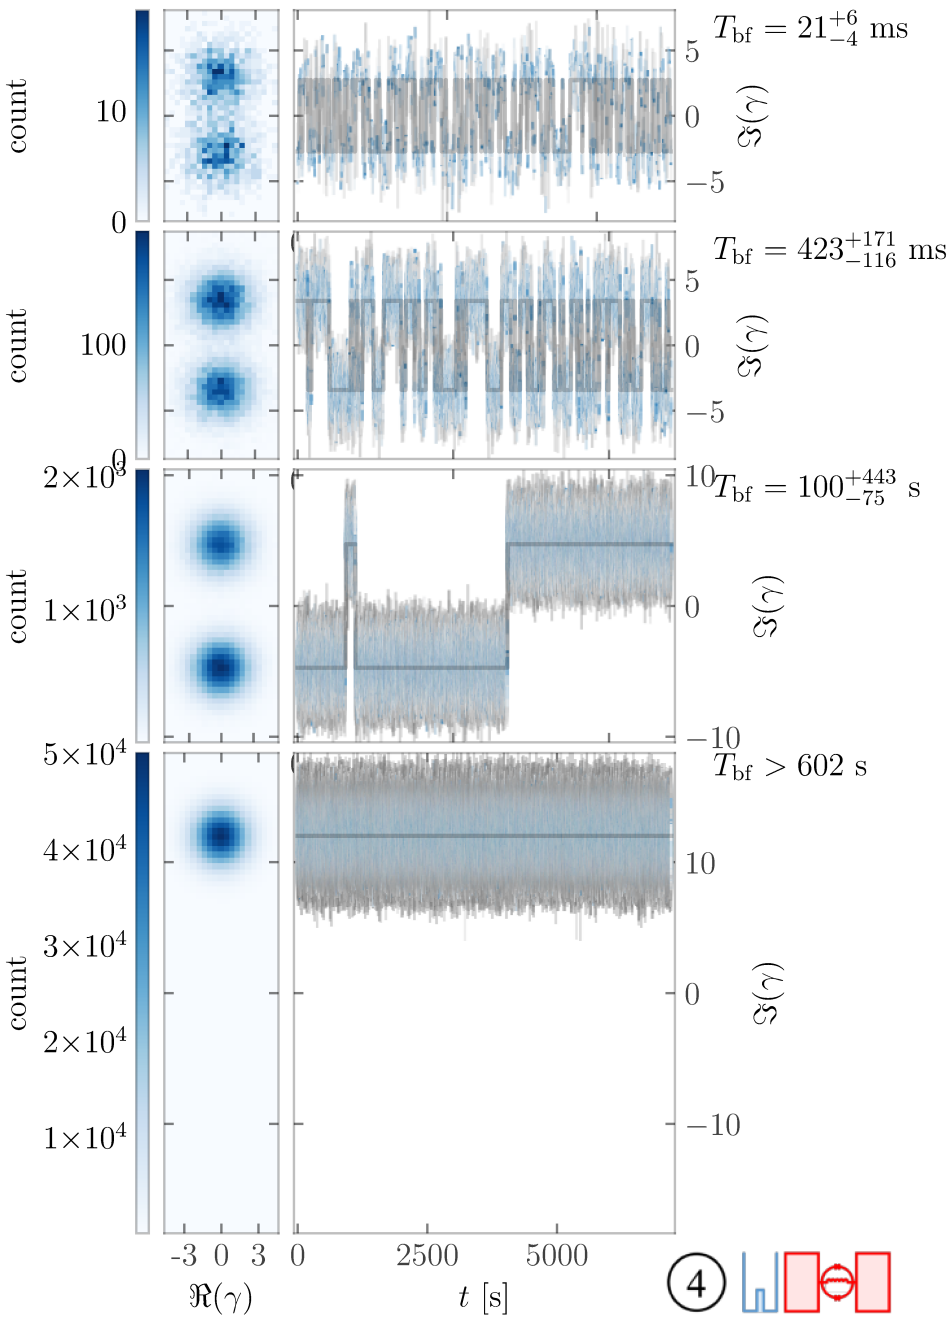

Figure 5.30 Trajectories of the memory for varius cat sizes.

tups involving multiple transmon qubits, it has been established that bit-flip events are often correlated. For instance, [McEwen et al. 2022] reports that within a surface code, single bit-flips can trigger a 'wave' of subsequent bit-flips spreading across the sample surface. These events are likely induced by energetic particles impacting the sample, possibly leading to the creation of quasiparticles. When quasiparticles cross Josephson junctions, they introduce additional dissipation channels [Dubovitskii et al. 2024]. It is conceivable that the bulk of the 3D cavity more effectively disperses the impact of these energetic particles, thereby mitigating the generation of quasiparticles.

5.3 Conclusion and Outlook

This last piece of data concludes our discussion on the 3D cat project. Although the results were not entirely conclusive or completely satisfactory, the project yielded several accomplishments. In a single iteration of the experiment (assembly 2 in Fig.5.10), we successfully coupled a lossy buffer mode ($\kappa_b/2\pi = 17$MHz) to a 3D cavity while maintaining its bare lifetime above $T_a = 250\ \mu$s, and demonstrated total control over the common and differential flux threading through the ATS at the center of the buffer. Thanks to the coupling between the memory and the buffer, the effective two-photon dissipation we achieved was $\kappa_2/2\pi \approx 28$ kHz. However, under parametric pumping, the effective lifetime of the memory degraded to $\kappa_a^{\text{eff}}/2\pi = 4.4$ kHz. In another iteration (assembly 3), we produced a dataset demonstrating the detrimental effect of the tomography transmon qubit on the bit-flip time of the cat qubit. Finally, by removing the transmon (assembly 4), we were able to demonstrate bit-flip times exceeding two hours, at the cost of being unable to confirm the quantum properties of our system.

The insights gained have laid a fertile foundation for designing a new experiment, which is the subject of the next chapter. At this stage, the future of the 3D cat project remains uncertain. If it continues, the next steps might be undertaken by Alice&Bob or through a collaboration with another lab, such as the groups led by Serge Rosenblum and Fabien Lafont at the Weizmann Institute.

Here are some insights regarding a possible continuation of the project. First, I think the rectangular tunnel design should be abandoned. It is too complex to manufacture, and working with multiple lines in the same tunnel with a microstrip line geometry introduces excessive complications related to filtering. Therefore, I suggest reverting to cylindrical tunnels, one per line.

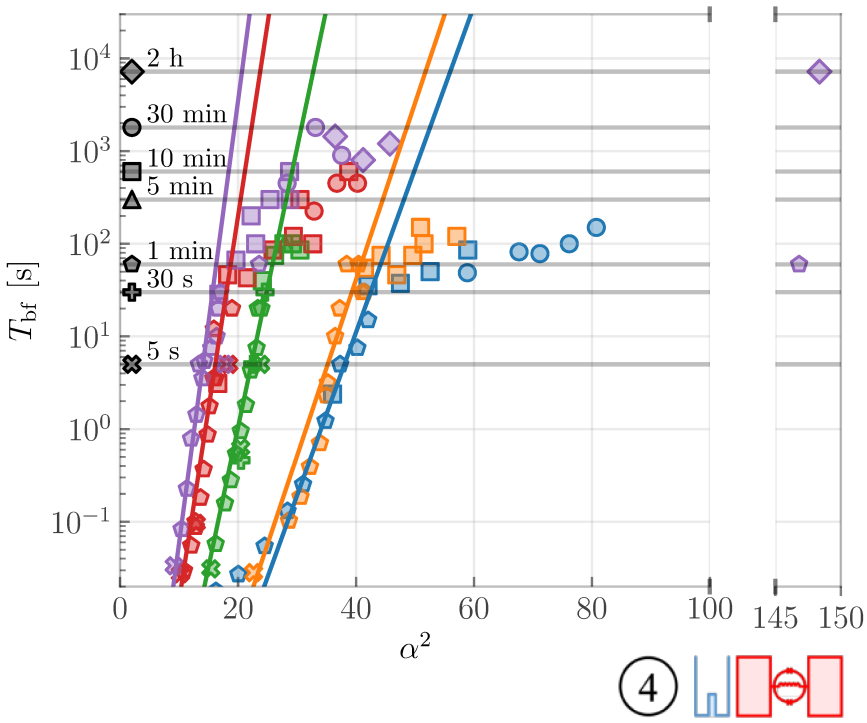

Figure 5.31 Scaling of the bit-flip rate versus pump power. For pump amplitudes of 3, 6, 9, 12, and 15 dBm (represented by blue, orange, green, red, and blue colors respectively, which coincides with the convention of Fig. 5.29), the bit-flip time (y-axis) is measured against the photon number (x-axis). This time is calculated from time traces of varying lengths, from 5 seconds to 2 hours, indicated by different symbols on the plot. As the pump amplitude increases, the saturation point of the bit-flip time extends from 1 minute to over 2 hours.

By slicing through these three cylindrical tunnels, we could create a slot where a single chip holding the three flux lines could sit. By spacing the tunnels sufficiently, each line could be filtered independently, for instance with stepped impedance filters [Hong and Lancaster 2004].

Now, these filters are going to occupy a lot of space on the wafer they are fabricated on. Therefore, having the ATS on a standalone chip is advisable. To achieve this, it is necessary to develop a mechanical fixation that allows the ATS chip to be positioned extremely close to the flux line on the filter chip. This arrangement separates the problems of junction fabrication and filtering.

Finally, to increase the two-photon coupling, boosting the junction energy seems to be the way forward. Also, as has recently been demonstrated with SQUID [Lu et al. 2023], opting for inductive coupling between the 3D cavity and the ATS might be an excellent way to increase the participation of the memory.

6

TRANSMON FREE 2D DISSIPATIVE CAT-QUBIT

Transitioning from our explorations in the 3D cat project, which despite its ambitious start eventually reached a standstill, this chapter introduces a pivotal shift towards a simpler, more direct approach in quantum circuit design.

The 3D cat project had reached its end, and personally, I found myself grappling with disappointment. However, delving into the autopsy of the data illuminated a crucial lesson: a good design is a simple design! Motivated by this insight, we devised a new circuit.

The presence of the transmon was associated with unexplained parasitic dynamical effects, potentially leading to bit-flips. Although the exact reasons were unclear, the solution seemed straightforward: eliminate the transmon. Likewise, the mechanical fixation of samples within 3D cavities was introducing variability in couplings. To counter this, we shifted back to a more predictable 2D geometry. Embracing this configuration, we hypothesized that if we could achieve a sufficiently high two-photon dissipation rate, the pathway to performing effective tomography of the memory mode would reveal itself in due course. Fortunately, this strategic gamble proved worthwhile.

The content of this chapter is largely derived from the associated paper [Ulysse Réglade and Adrien Bocquet et al. 2023]. In an effort to match the practical style of the previous chapter, I have reorganized the text to follow the chronological order of the successive calibrations. If readers are eager to get directly to the details regarding the main results, they may prefer to consult the paper itself or skip to the end of this chapter.

6.1 Principle of the experiment

Let us set the stage! In this experiment, we implement a cat-qubit with bit-flip times exceeding ten seconds, an improvement of four orders of magnitude over previous cat-qubit implementations, and six orders of magni-

tude over the lifetime of the photons composing the qubit. We observe phase-flip times greater than 490 ns, mainly limited by single-photon loss. We control the phase of coherent superpositions by rotating in a Zeno-blocked manifold [Touzard et al. 2018], performing a π rotation around Z in 235 ns. We verify that this manipulation only marginally reduces the bit-flip time, maintaining it above ten seconds. This was made possible by implementing a quantum tomography protocol that requires no additional ancillary elements [Albert et al. 2016]. Indeed, the Josephson dipole that mediates two-photon dissipation is operated to map quantum observables of the memory onto the buffer. This experiment demonstrates the tomography and control of a cat-qubit without breaking bit-flip protection at the level of bit-flip times of ten seconds. However, further improvements in state preparation, measurement fidelities and single photon loss will be necessary before scaling to a fully-protected hardware-efficient logical qubit [Guillaud and Mirrahimi 2019; Puri et al. 2020; Darmawan et al. 2021; Chamberland et al. 2022b]. The stage is now set, and we shall dive into the heart of the matter.

The principle of the experiment is illustrated in Fig.6.1a. As readers will notice, it only slightly differs from the archetypal design we presented in the first chapter (see Fig.4.4). The primary distinction is that our memory resonator is a Coplanar Wave Guide (CPW) $\lambda/4$-resonator. Several reasons underpin this design choice.

First, generally, the quality factor of distributed modes is better understood than that of lumped mode. While lumped modes occupy a smaller footprint, our objective is to demonstrate the properties of a single qubit; thus, we can accommodate the larger component in exchange for potentially better quality factors.

Second, for the four-wave mixing process we aim to implement, using a $\lambda/4$-resonator is advantageous compared to a $\lambda/2$-resonator (as used in [Berdou et al. 2023]). This is because the second mode of a $\lambda/2$-resonator has a frequency very close to $2\omega_a$, which can interfere with the frequency matching condition outlined in Eq. 4.30.

Lastly, this geometry features an open boundary condition at the end of the memory resonator, making it an antinode of the electric field. As depicted in Fig. 6.2, a large capacitor is placed at this point and coupled to the buffer. This configuration allows us to increase the participation ratio of the memory's zero-point fluctuation field φ^a_{zpf} compared to that in the 3D cat experiment.

Provided that we operate close to the first resonance of the $\lambda/4$-resonator, it can be modeled as an LC-oscillator. Consequently, our circuit is accurately described by Eq. 4.26 introduced in the first chapter. The initial step

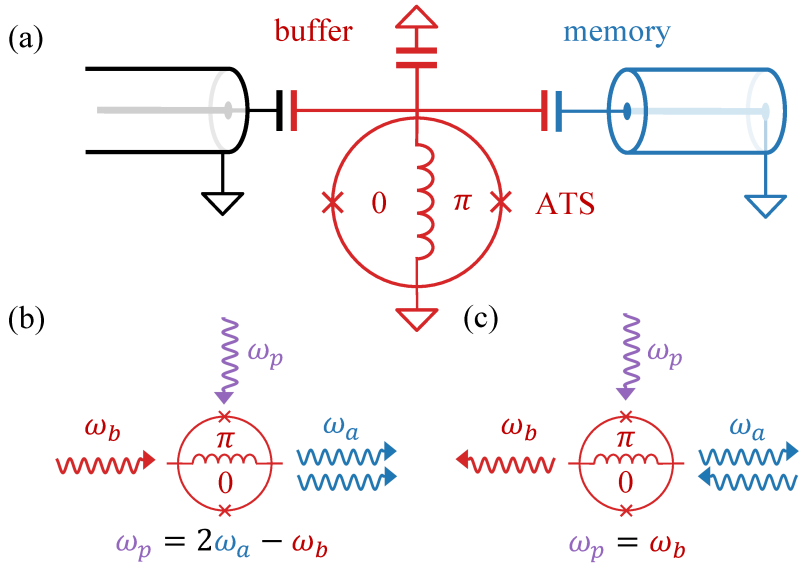

Figure 6.1 Experiment diagram of 2D cat. (a) Circuit implementation of our dynamical system. A quarter wavelength transmission line mode (memory in blue) is coupled to its environment (black) through a buffer mode (red) composed of an asymmetrically threaded SQUID (ATS) (see text). (b) A pump (purple) and a buffer drive (red) combine through the ATS to inject photon pairs in the memory (blue). The reverse process that removes photon pairs is not represented. This process allows to stabilise Shrodinger cat states in the memory. (c) A pump at frequency ω_b (purple) displaces the buffer (red) conditionally on the number of photons in the memory (blue), a key ingredient for our quantum tomography protocol.

in our exploration will be to determine the parameters of this circuit.

6.2 Calibration of the experiment

6.2.1 Flux point

In order to pin down the value of each parameter entering the Hamiltonian Eq. 4.26, we run a 3D finite element electromagnatic simulation of our device in the absence of the two SQUID junctions, and where the ATS inductance consisting in a chain of 13 junctions is modeled by an inductive lumped element of energy E_L. Although room temperature resistance mea-

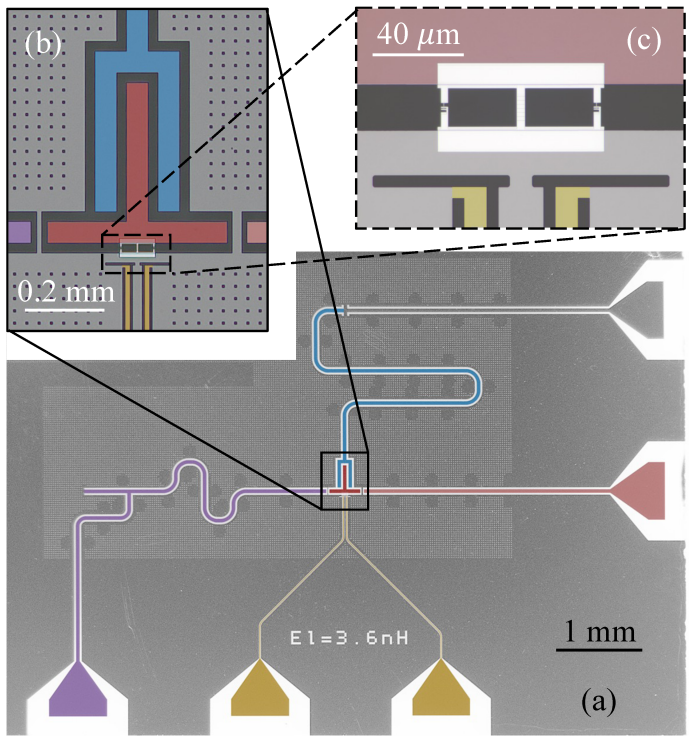

Figure 6.2 Optical micrographs of a twin sample. (a) A quarter wave-length coplanar waveguide (CPW) resonator supports the memory mode (blue). It is shunted to ground on its top end. The transmission line in its vicinity is unused in our experiment and is shunted to ground with a wirebond. On its bottom end, the memory oscillator couples capacitively to a superconducting island (red), corresponding to the buffer capacitance. The buffer is coupled from the left to a transmission line (purple) incorporating a Purcell filter. From this port we collect the buffer signals, and inject the buffer drive. A weakly coupled transmission line (pink) channels the memory drive and the cancellation tone for the longitudinal coupling into the circuit. The two symmetric lines (yellow) carry common and differential flux to bias the ATS, as well as the two-photon and longitudinal pump signals. (b) Zoom on the buffer island (red) and its surroundings. The tantalum film (light grey) is punctured with holes to avoid vortex trapping. The substrate (dark grey) is made out of sapphire. (c) Zoom on the ATS, that is fabricated in aluminum (white). The two small cross junctions (left and right) form a SQUID loop that is shunted in its center by a chain of 13 junctions.

$\omega_b/2\pi$	7.70 GHzz
φ_b	0.29
$\omega_a/2\pi$	5.26 GHz
φ_a	0.06
E_L/h	42.76 GHz
E_J/h	12.03 GHz
$\Delta E_J/h$	0.47 GHz

Table 6.1 Buffer, memory and ATS parameters entering Hamiltonian (4.26) in order to produce the fit of Fig. 6.3.

surements provide an estimate of E_L through the Ambegaokar-Baratoff formula, its precise value is determined by sweeping E_L in the simulation and extracting the corresponding $\omega_a, \omega_b, \varphi_a, \varphi_b$. We then pick the value of E_L for which ω_b matches best the mean of the measured frequencies $\omega_{b,\pm}$ at the two saddle points. The difference between the measured frequencies at these two saddle points then sets the SQUID junctions asymmetry ΔE_J. Finally fitting the measured buffer frequencies away from the saddle points sets the value of E_J. The comparison between the measured buffer frequencies and the ones extracted from Eq. 4.26 is displayed in Fig. 6.3. The parameters entering the Hamiltonian are listed in table 6.1.

6.2.2 Longitudinal readout of the memory

When we use the ATS as a mean to stabilise a cat qubit, it is well described by Eq. 2.40. However, it is important not to forget that a wave mixer can be used in more than one way. The originality of our experiment is that we are going to "hack" the ATS in a probing apparatus for the memory. Josephson circuits have been referred to as the Swiss army knife of microwave quantum optics [Flurin 2014]. Simply by switching the pump frequency, the behavior of a dipole can be dramatically changed. By setting the pump frequency to $\omega_p = \omega_b$ (Fig. 6.1c) in Eq. 4.26, the following low order processes are resonantly selected in the rotating wave approximation:

$$\begin{aligned} U_{\omega_p^l=\omega_b}(\boldsymbol{\varphi})/\hbar &\approx g_{\text{lin}}(\mathbf{b}+\mathbf{b}^\dagger) + g_l \mathbf{a}^\dagger\mathbf{a}(\mathbf{b}+\mathbf{b}^\dagger) \\ &+ g_{\text{sp}}\left(\mathbf{b}^{\dagger 2}\mathbf{b} + \mathbf{b}^\dagger\mathbf{b}^2\right) . \end{aligned} \tag{6.1}$$

Let us dissect this expression. The First term is a linear drive on the

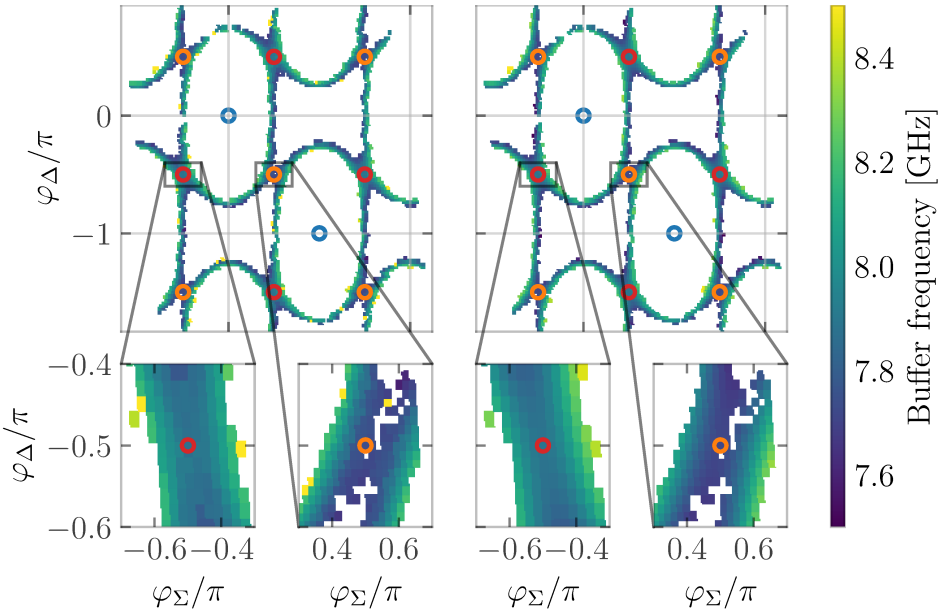

Figure 6.3 2D flux map of the buffer mode. Measured (left) and simulated (right) buffer frequency (color) as a function of differential (x-axis) and common (y-axis) flux. Open orange and red circles mark the location of the two non-equivalent saddle points. The blue open circles point to the maximum of the flux map. The bottom panels are zooms around two nonequivalent saddle points. In the blank regions, the buffer is misaligned with its Purcell filter and is therefore too undercoupled to the transmission line to be detected. For the rest of the experiment, the flux is set to $(\varphi_\Sigma, \varphi_\Delta) = (-\pi/2, -\pi/2)$.

buffer of amplitude $g_{\text{lin}} = -E_J \varepsilon_p^l \varphi_b/\hbar$. We cancel it by applying an additional tone on the buffer, with equal amplitude and opposite phase, through a capacitively coupled port. The second term is the one we desire. In the context of cavity QED, the coupling it implements is referred to as "longitudinal". Its strength is given by $g_l = E_J \varepsilon_p^l \varphi_a^2 \varphi_b/\hbar$ [Touzard et al. 2019]. Under its action, conditioned on the number of photons n_a in the memory, the buffer converges towards a coherent state of amplitude denoted $\beta_{1ph} \times n_a$. When cascaded with a heterodyne detection of the buffer, it constitutes a quantum non demolition (QND) measurement of the memory photon number.

Finally, the third term, where $g_{\text{sp}} = \frac{1}{2}\varphi_b^2/\varphi_a^2 g_l$, is a spurious term that limits the dynamical range of our detector. Future experiments will optimize

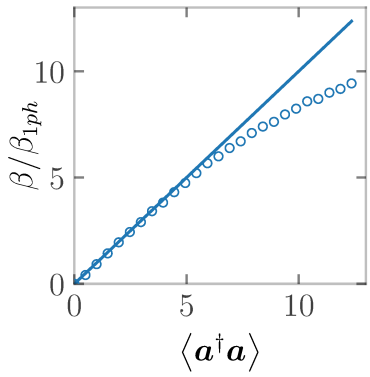

Figure 6.4 Longitudinal readout of a displaced state. We measure (open circles) the buffer amplitude (y-axis) versus the memory photon number (x-axis) reached after a displacement pulse. The deviation from the linear trend (solid line) is a signature of compression due to higher order processes.

the ratio φ_b/φ_a to maximize the readout fidelity of the memory through this longitudinal coupling to the buffer. Considering only the longitudinal interaction, when there are exactly n_a photons in the memory, the buffer converges to a coherent state of amplitude $\beta_{1ph} \times n_a$, where $\beta_{1ph} = -2ig_l/\kappa_b$.

In the following we describe the procedure we follow to calibrate (i) the pump frequency (ii) the amplitude of the cancellation tone and (iii) its phase. We sweep these parameters in a three-dimensional grid. For each combination we activate the pump and cancellation tones and record histograms of the buffer output both with and without an initial memory displacement. We select the combination of parameters that maximizes the Kullback-Leibler divergence between the two histograms. Finally, we record the buffer output as a function of the amplitude of the memory displacement (Fig. 6.4). We observe that the buffer output amplitude grows linearly with the number of photons in the memory up to about 5 photons, as expected from a longitudinal interaction. Above 5 photons, we observe a deviation from the linear dependence likely due to the spurious term $g_{\text{sp}}\left(\mathbf{b}^{\dagger 2}\mathbf{b} + \mathbf{b}^\dagger \mathbf{b}^2\right)$. The power of the pump is chosen to maximize the span covered by the buffer amplitude over a window of 0 to 10 memory photons, which is the range of interest in this experiment. For the optimal pump power, we are able to distinguish between the vacuum and a coherent state in the memory containing 10 photons with a single-shot fidelity of 89% (Fig. 6.5).

Importantly, at this stage, it is unclear how the axes of Fig.6.4 and Fig.6.5

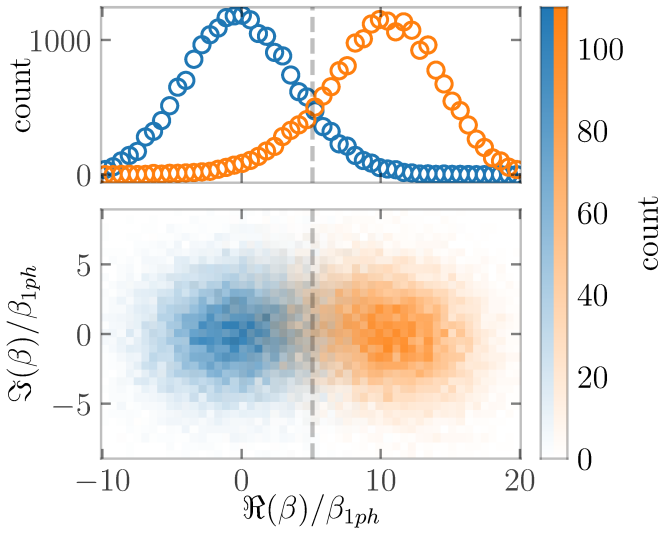

Figure 6.5 Measurement fidelity of the longitudinal readout. The memory is prepared in the vacuum (blue) or displaced to a coherent state containing an average photon number of 10 (orange). The longitudinal pump is then activated, and the buffer output field is integrated over 10 μs. This measurement is repeated 20 000 times for each memory state preparation. (Bottom) Two-dimensional histogram of the real and imaginary quadratures of the buffer field. The histogram was rotated so that all the information on the memory state is encoded in the real quadrature. The histogram was also displaced so that the origin of the axis coincides with the center of the histogram corresponding to a memory in the vacuum. Finally the real and imaginary axis are re-scaled in units of average photon number present in the memory. (Top) Histogram of the real quadrature of the buffer output field, obtained by integrating the two-dimensional histogram over the imaginary quadrature. The vertical red dashed line is used to compute a single shot fidelity. We denote $P(L|0ph)$ the probability of measuring a data point on the left of the threshold, given a memory state prepared in the vacuum. Similarly, we denote $P(R|10ph)$ the probability of measuring a data point on the right of the threshold, given a memory state prepared in a coherent state containing 10 photons. We define the measurement fidelity as $(P(L|0ph) + P(R|10ph))/2$, and find 89%.

were calibrated. The displacement of the buffer is proportional to the occupation of the memory, yet we currently lack knowledge of the "size" of a single photon. Later in this chapter, we demonstrate how we were able to access the Wigner function of the memory state, enabling us to ascertain this critical quantity.

6.2.3 Two-Photon Exchange

For now, let us concentrate on the two-photon stabilization. It is important to note that since we are now using two types of pumps - one for the two-photon dissipation and another for the longitudinal readout — we denote their associated frequencies as ω_p^{2ph} and ω_p^l, respectively. As before, the two-photon pump Stark-shifts the mode frequencies $\omega_{a,b}$. To calibrate the shifted frequencies, we activate the pump and vary its frequency ω_p^{2ph} in the vicinity of the frequency matching condition $2\omega_a - \omega_b$. This time around, since we have a mean to easily monitor the response of the memory, we change our strategy. Rather than performing a spectroscopy of the buffer, we apply a weak drive on the memory (which will later be used to perform Z rotations) of amplitude ε_Z whose frequency ω_Z and vary its frequency around ω_a. For each value of this two-dimensional frequency sweep, we measure the memory photon number through longitudinal coupling to the buffer. This defines a two-dimensional map shown in Fig. 6.6, where $(\omega_p^{2ph}, \omega_Z) = (2\omega_a - \omega_b, \omega_a)$ is a saddle point. Indeed, at this operating point, the memory photon number is a balance between the memory drive strength, and the two-photon confinement rate. Detuning the two-photon pump results in a lower confinement rate, and hence a larger memory photon number. Conversely, detuning the memory drive reduces the memory photon number. Fitting the data to a second degree polynomial, we precisely locate the saddle point and pin down ω_p^{2ph} and ω_Z. Finally, we set the buffer drive of amplitude ε_d to frequency $\omega_d = 2\omega_Z - \omega_p^{2ph}$.

With the pump and drive frequencies calibrated, the first step is to plot the trend of the photon number inside the memory against the drive amplitude, as shown in Fig. 6.7. The relationship is almost perfectly linear, suggesting that the two-photon dissipation rate likely exceeds the cavity's loss rate. Previously, we could have reduced the pump amplitude to reveal a dissipative phase transition, calibrated the pump at this operating point, and assessed the two-photon loss rate at higher pump amplitudes relative to this calibration. However, this method lacks precision. Considering that we can measure the photon number in the memory over time, we opt here for a different strategy.

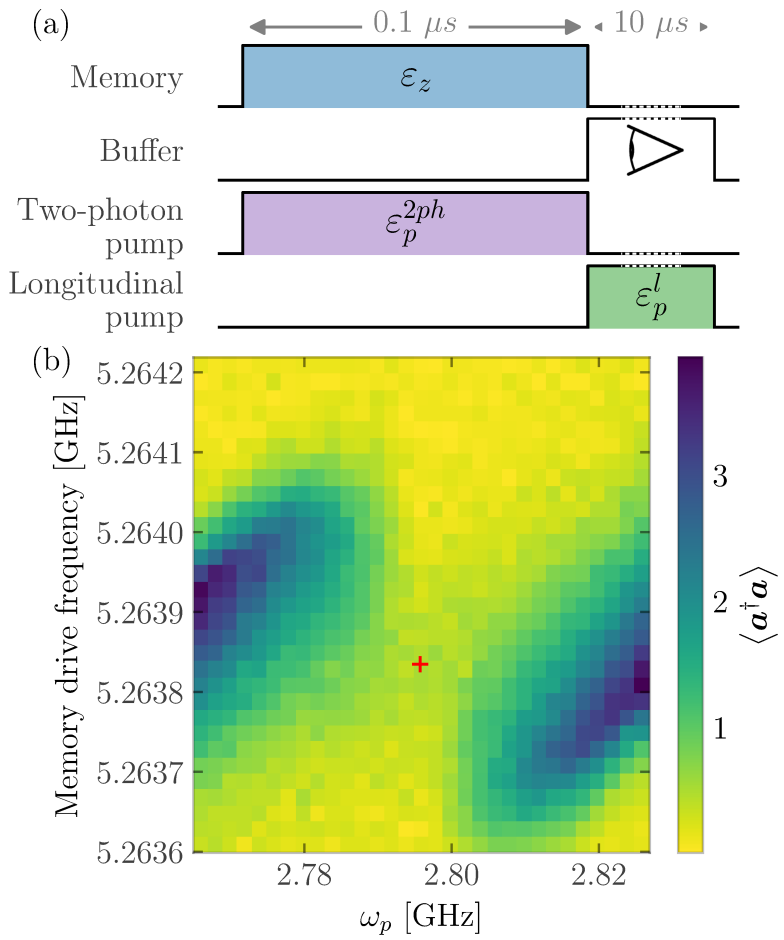

Figure 6.6 Calibration of the two-photon pump and memory drive frequencies. (a) Pulse sequence: square pulses are applied on both the memory and the two photon pump. Their duration is 0.1 μs and their frequencies are swept. Then we measure the memory photon number through longitudinal coupling for 10 μs (b) Photon number in the memory (color) as a function of the two-photon pump frequency ω_p^{2ph} (x-axis) and memory drive frequency ω_Z (y-axis). At the saddle point (red cross), the frequency matching condition $2\omega_Z - \omega_p^{2ph} = \omega_b$ is satisfied. Subsequently we set the buffer drive frequency to $\omega_d = 2\omega_Z - \omega_p^{2ph}$.

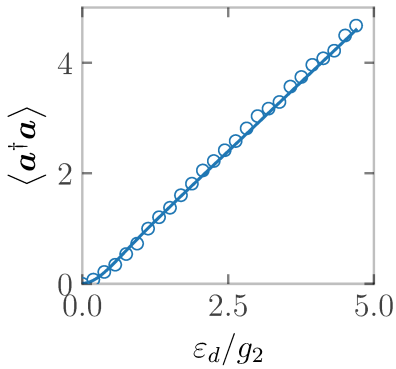

Figure 6.7 Photon number imposed by the two-photon stabilisation. We measure (open circles) a linear increase (simulation in solid line) of the steady-state memory photon number (y-axis) versus drive amplitude (x-axis). Compared to Fig. 5.29, almost no curvature is visible, meaning the two-photon dissipation should be greater than the losses of the memory mode.

To determine g_2, we measure the rate at which photons are removed from the memory mode by the two-photon dissipation mechanism. To do so, we implement the following pulse sequence (Fig. 6.8a): first the memory is displaced to a coherent state of amplitude α, then the two-photon pump is activated for a variable time ΔT, and finally, the memory photon number is measured by activating the longitudinal pump. The measured memory photon number versus time for various α is displayed in Fig. 6.8b. Two timescales are apparent. On a short timescale $1/\kappa_{\text{conf}} \approx 2/\kappa_b = 120$ ns, pairs of photons are quickly removed from the memory until it is left in a mixture of $|0\rangle$ or $|1\rangle$ Fock states. At this stage, the number of photons $\langle a^\dagger a \rangle$ ranges from $1/2$ for initial state amplitudes $\alpha \gg 1$ down to 0 when $\alpha = 0$. Then, the remaining photon is dissipated into the memory bath on a longer timescale $1/\kappa_a$. The data are fitted by solving the master equation of the memory-buffer system (Eq. (2.40)), with $\varepsilon_d = 0$, and $\rho(0) = |\alpha\rangle\langle\alpha|$. The only fit parameters are the two-photon exchange rate and the thermal population of the buffer, for which we obtain $g_2/2\pi = 763$ kHz and $n_{th}^b = 1.1\%$. The dissipation rates κ_a and κ_b are measured independently using standard techniques.

To fit the curve of Fig. 6.8, again, the number of photons in the memory (y-axis) and the thermal population of the memory n_{th}^a are to be calibrated in the first place. It is now time to explain how, by exploiting the two-photon

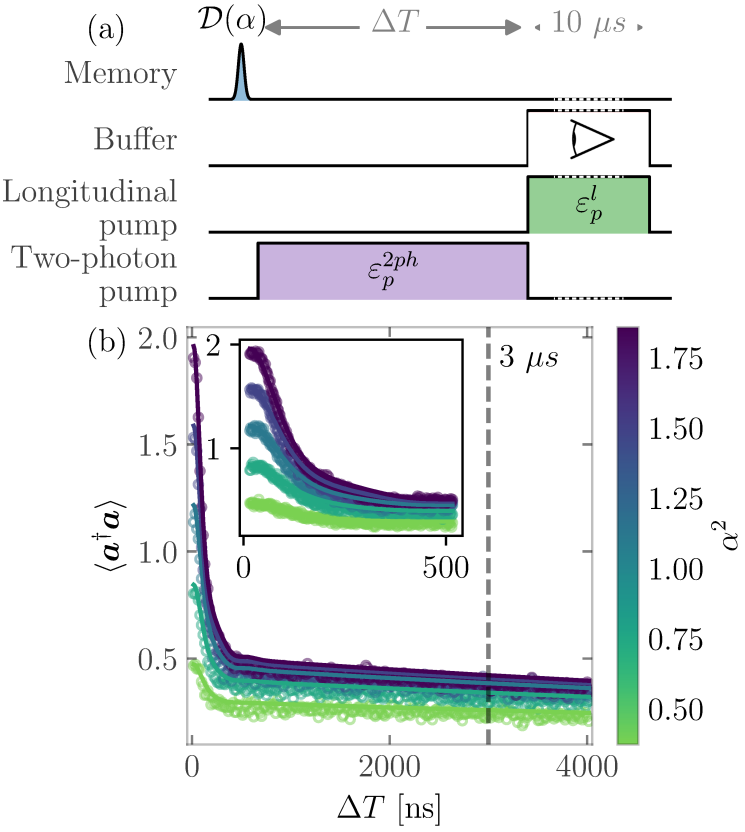

Figure 6.8 Cat deflation. Calibration of the convergence time towards the $|\pm\rangle_{\alpha=0}$ manifold, colloquially referred to as the deflation time. (a) Pulse sequence : first, a coherent state of amplitude α in the memory is created. A two-photon pump is then activated for variable time ΔT, and finally, the memory photon number is measured by longitudinal readout through the buffer. (b) Photon number (y-axis) as a function of ΔT (x-axis) for various initial coherent states $|\alpha\rangle$ (color). The data (open circles) are overlaid on theory curves obtained from solving the Lindblad master equation of Eq. 2.40 with $\varepsilon_d = 0$. The parameters entering the simulation are displayed in table 6.2. The only parameters that were varied to fit the data are g_2 and n_{th}^b. Other parameters were determined independently.

$g_2/2\pi$	0.763 MHz
$\kappa_b/2\pi$	2.6 MHz
$\kappa_a/2\pi$	9.3 kHz
n_{th}^a	10%
n_{th}^b	1.1%

Table 6.2 Parameters entering a Lindblad master equation to produce the numerical simulations of the paper, namely Fig. 6.1b, Fig.6.17b, Fig. 6.8b, Fig. 6.14c and Fig. 6.13b.

dissipation, we are able to access the value of the parity operator.

6.2.4 Measure of the parity operator

The readout protocol based on the memory-buffer longitudinal coupling yields a buffer signal proportional to the number of photons in the memory. Hence the proportionality constant must be calibrated. Even though the experiment shown Fig. 6.8 allows to calibrate it in principle, it is not very accurate. In order to access this crucial information, we elaborate a simple Wigner tomography tool. The signal to noise ratio of this method quite low, but it requires few calibrations. Later, after many more calibrations, we show how the SNR of the parity measurement can be improved.

Remarkably, since two-photon dissipation (with a buffer drive set to zero) removes photons in pairs, it maps all even parity states to the vacuum $|n=0\rangle$, and all odd parity states to Fock state $|n=1\rangle$. This maps parity to photon number, which is measurable through longitudinal coupling to the buffer. A memory displacement of amplitude γ followed by a two-photon pump applied without buffer drive during $2\ \mu$s (Fig. 6.8) therefore measures the Wigner function of the memory's initial state $W(\gamma)$ (Fig. 6.9a). Single-photon losses after mapping to 0/1 and the overlap between the histograms of the longitudinal readout corresponding to $|n=0\rangle$ and $|n=1\rangle$ lead to a global reduction of the scale of the Wigner tomography.

We apply this tomography protocol to three different initial state preparations: (i) applying the two-photon pump and buffer drive for $\Delta T = 80\ \mu$s resulting in an almost equal statistical mixture of $|\pm\rangle_\alpha$, (ii) applying the two-photon pump and buffer drive for $\Delta T = 800$ ns resulting in a mixture biased towards $|+\rangle_\alpha$ and (iii) no pump or drive at all, resulting in the vacuum state with some residual thermal population. Interestingly, note that while the thermal vacuum state displays a Wigner distribution that is broadened by

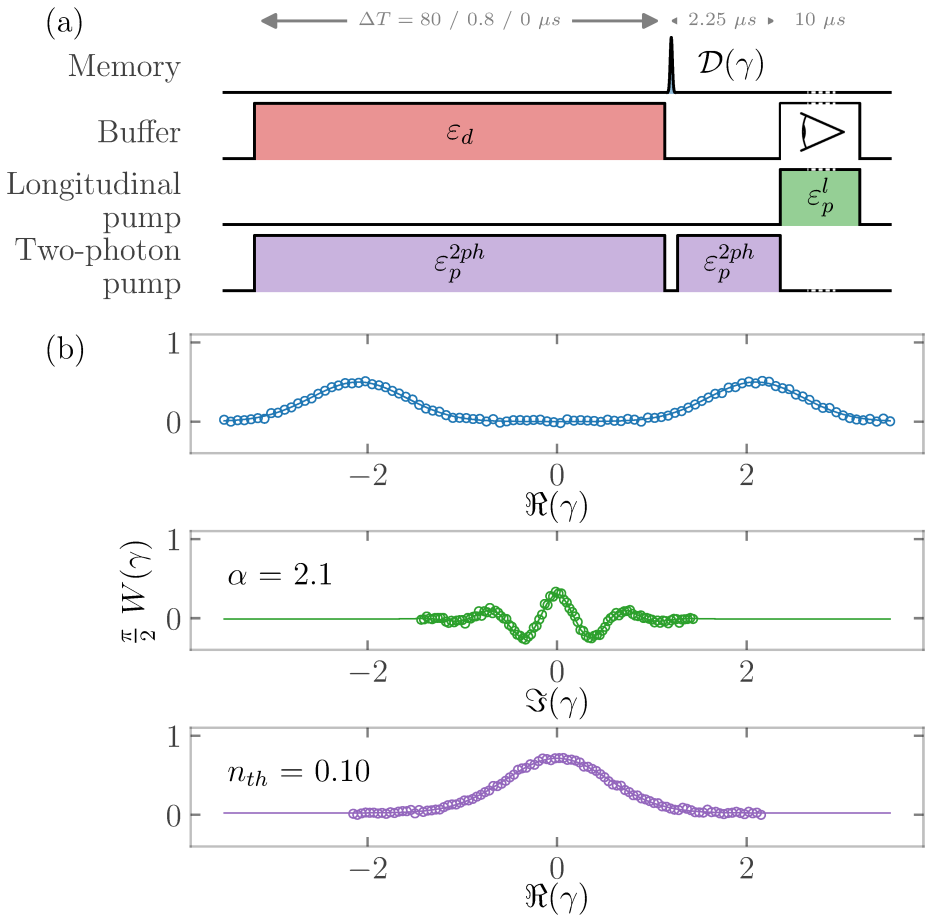

Figure 6.9 Photon number calibration. (a) Pulse sequence: a two-photon pump and buffer drive are active for $\Delta T = 80, 0.8, 0\ \mu$s, thereby preparing a statistical mixture of $|\pm\alpha\rangle$, a cat state $|+\rangle_\alpha$ and the thermal vacuum, respectively. Next, a Wigner tomography protocol is implemented (see text). (b) Cuts of the Wigner functions (y-axis) of these three states (top to bottom) along the real (top and bottom) and imaginary (middle) axis. The data (open circles) are overlaid with fits to the analytical formula displayed in Eq. (6.2) (solid lines).

thermal fluctuations, states (i) and (ii) are cooled into the cat-qubit manifold, and therefore the memory thermal occupation will impact the state reached within this manifold, but does not lead to thermal broadening. This asymmetry will be leveraged to calibrate the memory thermal occupation. Cuts of the Wigner functions of these three states are shown in Fig. 6.9.b, and overlaid with their fits to analytic formula [Haroche and Raimond 2006b]:

$$\begin{aligned} W_{\Delta T=80\ \mu s}(\gamma_r, 0) &= A\left(e^{-2(B\gamma_r-\alpha)^2} + e^{-2(B\gamma_r+\alpha)^2}\right) \\ W_{\Delta T=0.8\ \mu s}(0, \gamma_i) &= Ce^{-2(B\gamma_i)^2}\left(e^{-2\alpha^2} + \cos(4\alpha B\gamma_i)\right) \\ W_{\Delta T=0\ \mu s}(\gamma_r, 0) &= De^{-2(B\gamma_r)^2/(2n_{th,a}+1)}\ , \end{aligned} \tag{6.2}$$

where $\gamma_{r,i}$ are the real and imaginary part of γ and $A, B, C, D, \alpha, n_{th}^a$, as well as offsets, are free fit parameters. In essence, the frequency of the oscillations in (ii) sets an absolute scale B for memory displacements in units of square root of photon number; the width of the Gaussian in (iii) sets the thermal population $n_{th}^a = 10\%$ of the memory and the separation of the two Gaussians in (i) indicates the size of the cat-qubit $\alpha = 2.1$. Furthermore, we do not observe a broadening of the two Gaussians in (i), which indicates a small thermal population n_{th}^b of the buffer. Finally, the constants A, C, D encode the different fidelities of state preparation and measurement.

From now on, we have a means to measure the parity of the memory, which consequently enables us to perform its tomography. However, this protocol is relatively slow. Before analyzing the performance of our cat qubit, we will focus on accelerating this process!

6.2.5 Holonomic gate protocol

The tomography protocol described in the previous section relies on mapping the value of the parity to the occupation of states $|0\rangle$ and $|1\rangle$, which are challenging to distinguish since their energy levels differ by merely a single photon. In this section, we introduce a new tomography protocol (Fig. 6.10) based on the so-called holonomic gate proposed by Albert it et. al [Albert et al. 2016]. All odd parity states are mapped onto the vacuum $|0\rangle$, and all even parity states onto a coherent state $|\psi_5\rangle = |2i\alpha_2\rangle$ where $2\alpha_2 \approx 4.8$. The QND photon number measurement through longitudinal coupling then distinguishes between $|0\rangle$ and $|2i\alpha_2\rangle$, providing a photon number parity measurement with a much stronger contrast. Our pulse

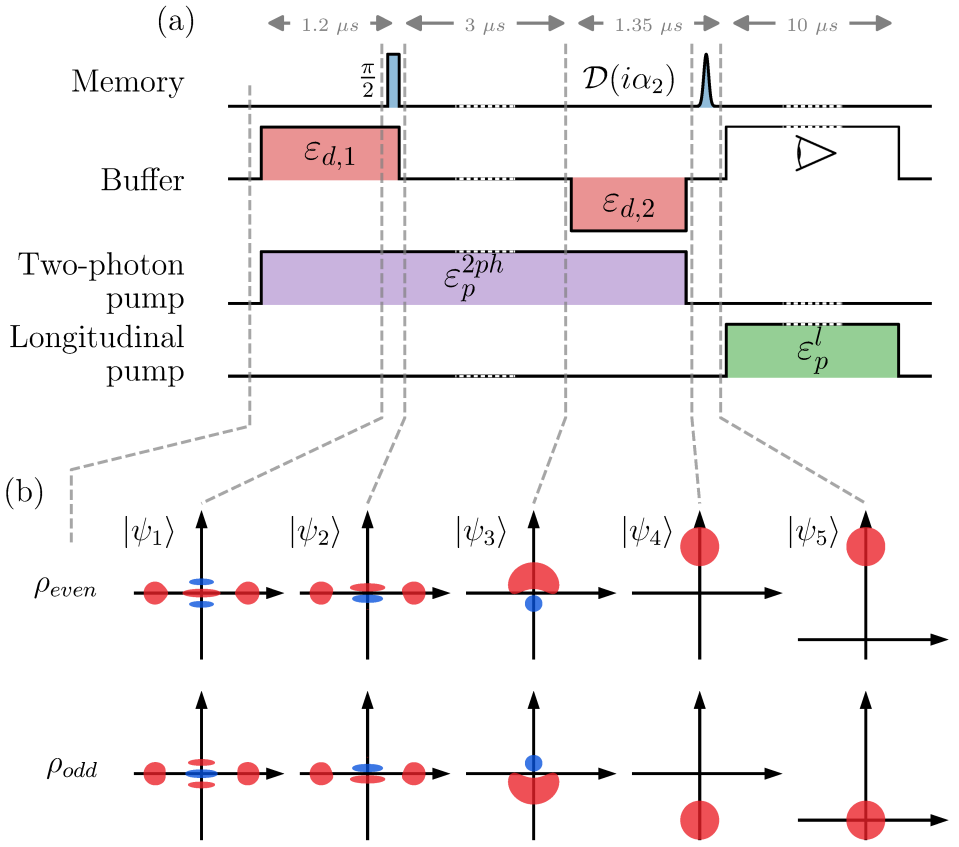

Figure 6.10 Quantum tomography protocol based on the holonomic gate [Albert et al. 2016]. (a) Pulse sequence for each control channel as a function of time. The parity of the memory is mapped onto photon number, which is measured through longitudinal coupling to the buffer (see text). (b) Cartoon of the memory Wigner function at each step (dashed lines) of the protocol for an even (top) and odd (bottom) initial state.

sequence (Fig. 6.10a) alternates between memory drives, two-photon dissipation, buffer drives of various amplitudes and longitudinal coupling.

We now describe it step by step (Fig. 6.10b). Let us assume that the memory is initially in an even parity state ρ_{even}. First, we activate a two-photon pump ε_p^{2ph} and buffer drive $\varepsilon_{d,1} = g_2^* \alpha_1^2$ where α_1 is real. By parity conservation, ρ_{even} is mapped to the even state $|\psi_1\rangle = |+\rangle_{\alpha_1}$. Next, we add a memory drive along the imaginary axis. Two-photon dissipation confines the dynamics to the quantum manifold spanned by $|\pm\rangle_{\alpha_1}$. The added memory drive induces coherent Zeno-blocked oscillations around the cat-qubit Z axis [Touzard et al. 2018]. We tune the drive length to perform a $\pi/2$ rotation reaching the parity-less state $|\psi_2\rangle = (|+\rangle_{\alpha_1} + i\,|-\rangle_{\alpha_1})/\sqrt{2}$ [Yurke and Stoler 1988; Kirchmair et al. 2013; Grimm et al. 2020]. Next, we turn off the memory and buffer drives while the two-photon pump remains active, thereby removing pairs of photons from $|\psi_2\rangle$. By parity conservation, $|+\rangle_{\alpha_1}$ is mapped to $|0\rangle$ and $|-\rangle_{\alpha_1}$ to $|1\rangle$ [Grimm et al. 2020]. When this mapping is adiabatic with respect to $\kappa_{\text{conf}}^{-1}$ and fast compared to κ_a^{-1}, quantum superpositions are preserved, yielding $|\psi_3\rangle = (|0\rangle + i\,|1\rangle)/\sqrt{2}$ [Albert et al. 2016]. In practice we use a square buffer drive of amplitude $\alpha_1 \approx 1.6$ in an attempt to maximize the fidelity of the $\pi/2$ gate around Z, while minimizing the loss of coherence during these mappings. Next, maintaining the two-photon pump, we activate a buffer drive $\varepsilon_{d,2} = -g_2^* \alpha_2^2$ where α_2 is real. The minus sign on the buffer drive translates into a pure imaginary amplitude for the stabilized coherent states. This maps $|0\rangle \to |+\rangle_{i\alpha_2}$, $|1\rangle \to -i\,|-\rangle_{i\alpha_2}$ [Grimm et al. 2020] and following the same reasoning: $|\psi_3\rangle \to |\psi_4\rangle = |i\alpha_2\rangle$ [Albert et al. 2016]. Conversely, an odd parity ρ_{odd} would be mapped to $|-i\alpha_2\rangle$. Information on the parity of $\rho_{\text{even/odd}}$ is now encoded on the amplitude of coherent states $|\pm i\alpha_2\rangle$. Finally, the two-photon pump is turned off, and the memory is displaced by $i\alpha_2$. The longitudinal pump ε_p^l is activated to distinguish between 0 and $4|\alpha_2|^2$ photons in the memory by heterodyne detection of the buffer. Note that the value of α_2 can be tuned to optimize the fidelity of the longitudinal readout. Preceding this entire sequence by a memory displacement of amplitude γ therefore measures $W(\gamma)$.

6.2.6 Reset of the memory

The SNR of the parity measurement is increased, leading to shorter data acquisition time. This is not to say we cannot compress it even further. The protocol for the generation and Wigner tomography of a memory state consists in a $\approx 15\mu$s pulse sequence followed by a long wait time for the memory to relax back to the vacuum. This wait time is typically $\approx 5/T_a$, where

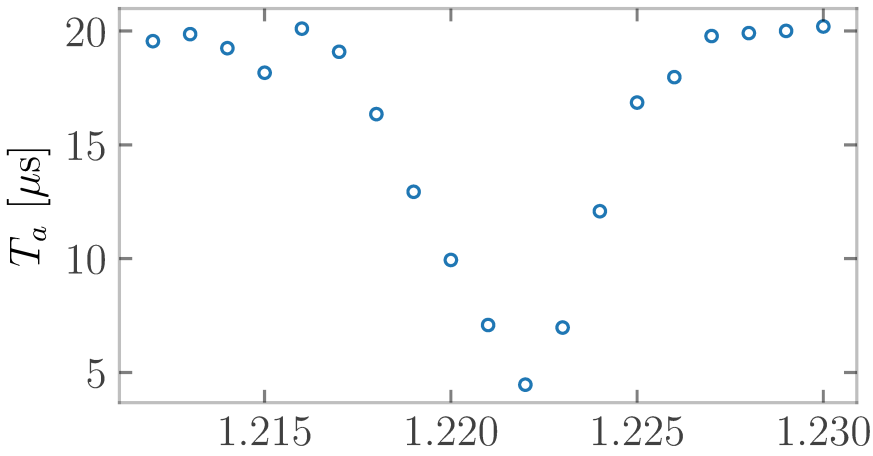

Figure 6.11 Calibration of the reset pump frequency. Memory lifetime T_a (y-axis) versus the reset pump frequency (x-axis). The data (open circles) are extracted from a fit of the form $\exp(-t/T_a)$ of the memory photon number decay versus time t in the presence of the reset pump.

$T_a = 1/\kappa_a$ is the memory lifetime. In order to reduce the acquisition time, we implement an active memory reset that reduces T_a, thereby accelerating the convergence of the memory to the vacuum. In this section we describe how we implement this reset mechanism.

Going back to Eq. 4.26, we consider the term in ΔE_J that was previously neglected. At the flux saddle point, this term contributes a non-linear interaction of the form $U^{\Delta}(\varphi) = 2\Delta E_J \cos(\varepsilon_p \cos(\omega_p t))\cos(\varphi)$. Assuming $\varepsilon_p, |\varphi| \ll 1$, we obtain:

$$U^{\Delta}(\varphi) \approx 2\Delta E_J(1 - \frac{1}{2}\varepsilon_p^2 \cos(\omega_p t)^2)(1 - \frac{1}{2}\varphi^2) \ . \tag{6.3}$$

The amplitude of the reset pump is denoted $\varepsilon_p^{\text{reset}}$ and its frequency is set to $\omega_p^{\text{reset}} = (\omega_b - \omega_a)/2$. Performing a rotating wave approximation on Eq. (6.3), we find:

$$U^{\Delta}_{\omega_p=\omega_p^{\text{reset}}}(\varphi)/\hbar = g_{\text{reset}}(\mathbf{a}\mathbf{b}^\dagger + \mathbf{a}^\dagger\mathbf{b}) \ , \tag{6.4}$$

where $g_{\text{reset}} = \Delta E_J {\varepsilon_p^{\text{reset}}}^2 \varphi_a \varphi_b / 8\hbar$. We place ourselves in the regime where $g_{\text{reset}} \ll \kappa_b$, which induces a dissipation rate on the memory of the form $\kappa_a^{\text{reset}} = 4g_{\text{reset}}^2/\kappa_b$.

In practice, we calibrate the reset pump by measuring the memory lifetime in the presence of the reset pump and vary its frequency (Fig. 6.11). When the reset pump frequency is far detuned from the frequency matching condition $(\omega_b - \omega_a)/2$, the memory lifetime is unaffected by the pump, remaining around 20 μs. However, when the pump frequency is set to $(\omega_b - \omega_a)/2$, the memory lifetime drops to ≈ 5 μs. Note that the buffer and memory frequencies that enter this frequency matching condition are those in the presence of the reset pump.

6.2.7 State preparation and measurement

Starting from the vacuum in the memory, we activate for a varying time ΔT both the two-photon pump and a buffer drive at various amplitudes ε_d. To assess the convergence of the state toward the cat qubit code space, we monitor the photon number in the memory with the longitudinal readout. The data are reported in Fig. 6.12. Three timescales are apparent. For $\Delta T \in [0, 250 \text{ ns}]$, we observe a rapid increase of photon number in the memory. Indeed, single photons from the buffer drive combine with the pump and pairs of photons are squeezed into the memory. For $\Delta T \in [250 \text{ ns}, 600 \text{ ns}]$, we observe damped oscillations in the memory photon number. Although we do not quantitatively understand the amplitude of these oscillations, they can be reproduced in simulation by adding a Kerr on the buffer, and a detuning of the pump and drive from the frequency matching condition. For $\Delta T \geqslant 600$ ns, the photon number reaches a plateau. This is expected when the memory has converged into the quantum manifold spanned by $|\pm\rangle_\alpha$. The specific state that is reached within this manifold depends on the initial state. Since the memory is initialized in the vacuum, the memory converges to the even cat $|+\rangle_\alpha$. However, due to inevitable losses in the memory, this state will evolve into a classical statistical mixture of $|\pm\rangle_\alpha$ at a rate $2|\alpha|^2\kappa_a$. The fidelity of preparation of $|+\rangle_\alpha$ is therefore maximized for the smallest time ΔT that insures the convergence of the memory in its steady state manifold. The cat states generated for Wigner tomography in Fig. 6.12 were prepared in $\Delta T = 1$ μs.

We can now demonstrate the effectiveness of the holonomic tomography protocol by measuring the Wigner functions of several states. In Fig. 6.13, we prepare and image states $|0\rangle, |+\rangle_\alpha, |-\rangle_\alpha$ and $|1\rangle$, labeled $|\psi_1\rangle$ to $|\psi_4\rangle$ for $\alpha^2 = 2.6$. (a) Pulse sequence: starting from the vacuum $|\psi_1\rangle$, the two-photon pump (blue) and buffer drive (orange) are activated. When no memory drive is played (dashed green), the even cat-state $|\psi_2\rangle$ is prepared. On the other hand, when the memory drive is played, it implements a $Z(\pi)$

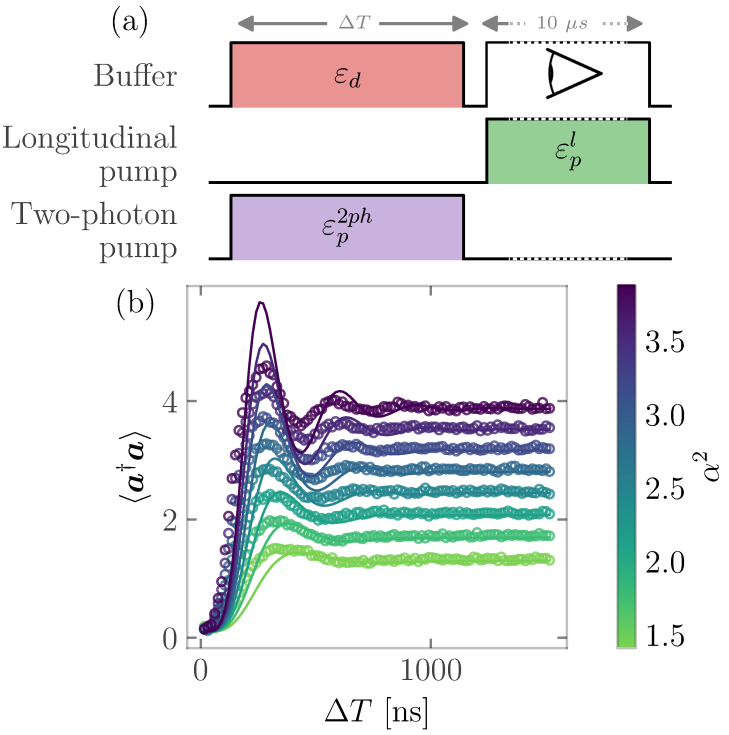

Figure 6.12 Cat inflation. Calibration of the convergence time towards the $|\pm\rangle_\alpha$ manifold, colloquially referred to as the inflation time. (a) Pulse sequence : starting from the vacuum in the memory, the two-photon pump and buffer drive is activated for variable time ΔT. The memory photon number is then measured by longitudinal readout through the buffer. (b) Photon number (y-axis) as a function of ΔT (x-axis) for drive amplitudes targeting various steady state memory photon numbers (color). The data (open circles) are overlaid with the results of a numerical simulation of Eq. 2.40 with $\rho(0) = |0\rangle\langle 0|$, $\varepsilon_d = \alpha^2 g_2$ and no fit parameters. The parameters entering this simulation are summarized in Table. 6.2 have all been independently calibrated.

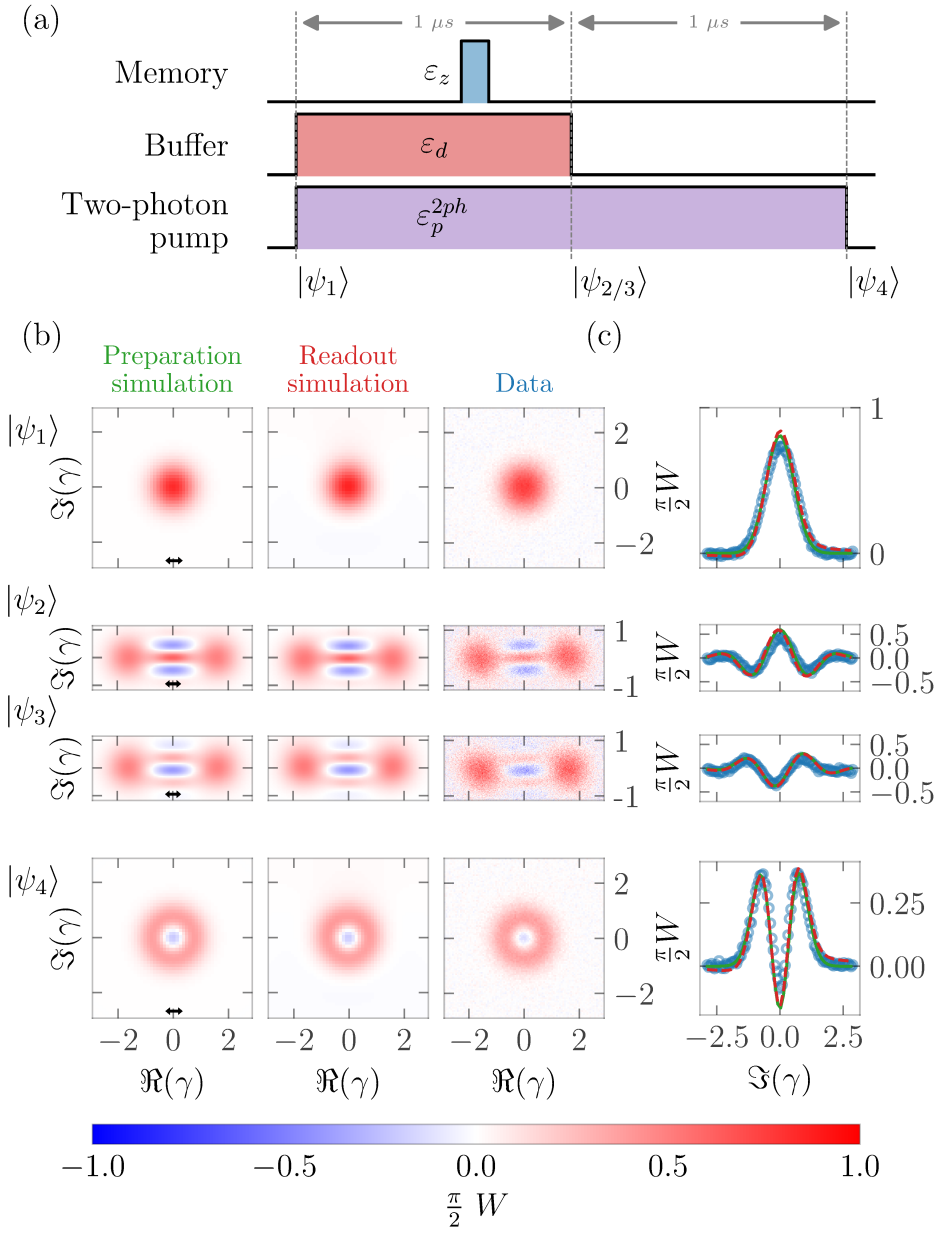

Figure 6.13 Quantum state tomography.

gate thereby preparing the odd cat state $|\psi_3\rangle$. Finally, in the latter case, the buffer drive is turned off to arrive at the first Fock state $|\psi_4\rangle$. At each step of this pulse sequence, the protocol of Fig. 2 is implemented to acquire the Wigner function of these states. (b) Wigner amplitude (color) over the complex plane (x-y axis) for each state $|\psi_{1,2,3,4}\rangle$. The data ("Measurement" column) are well reproduced by the Lindblad simulation of the entire protocol including the preparation only ("Preparation simulation" column) and both preparation and tomography pulse sequences ("Readout simulation" column). (c) Integration of the Wigners in (b) over the interval marked by the green double arrow.

6.2.8 Rabi-Zeno gate calibration

One of the main results of this work is to perform coherent oscillations of superpositions of metastable states with macroscopic bit-flip times. These oscillations are induced by the interplay of two-photon dissipation and a coherent drive on the memory. While the drive attempts to translate the memory state in phase space, two-photon dissipation continuously stabilizes the manifold spanned by $|\pm\rangle_\alpha$. As a result, when the phase of memory drive is 90 degrees out of phase with α, we observe coherent oscillations between $|\pm\rangle_\alpha$. The coherent motion in a manifold stabilized by dissipation is known as quantum Zeno dynamics [Schäfer et al. 2014; Signoles et al. 2014; Bretheau et al. 2015; Touzard et al. 2018], and results in a rotation around the Z-axis in the cat-qubit code space.

In this section we describe the procedure we follow to calibrate the amplitude and phase of the memory drive (Fig. 6.14). Starting from a memory in vacuum, we implement the pulse sequence displayed in Fig.6.14a, where the phase θ_Z and duration ΔT of the memory drive is varied. We observe $\langle Y_\alpha\rangle$ oscillating versus ΔT, with an oscillation frequency that depends strongly on θ_Z (Fig. 6.14b). We identify the phase at witch the oscillation frequency is minimal with $\theta_Z = 0 \mod \pi$, that is, the memory is driven in phase with α. On the other hand, we identify the phase at which the oscillation frequency of $\langle Y_\alpha\rangle$ is maximal with $\theta_Z = \pi/2 \mod \pi$, that is, the memory drive is 90 degrees out of phase with α. To optimize the Z gate speed, we set $\theta_Z = \pi/2$ in the pulse sequence of Fig. 6.10a. Surprisingly, when we measure the evolution of Z_α in the presence of a memory drive with $\theta_Z = \pi/2$, we find that the Z_α population is skewed towards one of the two pointer states $|\pm\alpha\rangle$.

Having fixed the memory drive phase, we now vary its amplitude ε_Z. The data representing $\langle Y_\alpha\rangle$ as a function of drive amplitude and duration

are displayed in Fig. 6.14c. Two features are visible. First, as the drive amplitude is increased, the frequency of the oscillation increases proportionally (Fig. 6.14d). Second, these oscillations decay faster as ε_Z is increased, which is a manifestation of so-called non-adiabatic errors. We plot the number of oscillations per decay time in Fig. 6.14e, and find that there is an optimal gate amplitude that balances gate duration and induced dephasing.

As described in Eq. 2.58, in the presence of memory drive, the state of the buffer depends on the state of the memory, which induces decoherence. The buffer then continuously radiates a field proportional to these amplitudes at a rate $\kappa_b/2$.

Fig. 6.14 presents the calibration data for the Rabi-Zeno gate which allows use to rotate arround the Z axis in the cat-qubit code space. (a) Pulse sequence: a two-photon pump and buffer drive prepare $|+\rangle_\alpha$. A memory drive is then activated with a variable phase and for a variable time ΔT inducing coherent rotations in the steady state manifold. The rest of the pulse sequence measures $\mathbf{Y}_\alpha$ as described in Fig. 6.10. (b-left) Buffer output proportional to the expectation value of $\mathbf{Y}_\alpha$ (color) as a function of the memory drive angle (x-axis) and duration (y-axis). The memory drive angle is set to maximize the oscillation frequency (red dashed line). (b-right) With the angle now calibrated, the memory drive amplitude is varied (x-axis) and the expectation value of $\mathbf{Y}_\alpha$ (color) is measured versus pulse duration (y-axis). (c-left and c-right) Numerical simulation of the pulse sequence (a) with the circuit parameters given in table 6.2. (d-left) From (b-right) we extract the oscillation frequency (y-axis) versus memory drive amplitude (x-axis). The data (open circles) follow the expected linear trend (solid line). We calibrate ε_Z by fitting the data to the relation $\Omega_z = 4|\alpha|\varepsilon_Z$ [Guillaud and Mirrahimi 2019]. (d-right) Number of coherent oscillations in units of oscillation frequency times fitted dephasing time (y-axis) versus memory drive amplitude ε_Z (x-axis). The data (open circles) reach a maximum when ε_Z balances rotation speed and induced dephasing (red star and dashed red line). Surprisingly, the data are well explained by an analytical formula (Eq. (2.58)) valid in the adiabatic regime $8g_2\alpha/\kappa_b \ll 1$ (solid blue line), although in our experiment, $8g_2\alpha/\kappa_b \sim 3.6$. On the other hand, the full master equation simulation underestimates the dephasing time (solid purple line).

An important consequence of Eq. 2.57 is that it becomes possible to perform a Z measurement of the cat-qubit by applying a memory drive and looking at the field radiated by the buffer. Note that driving the system harder (i.e increasing ε_Z) would result in a higher readout contrast, but would also activate higher-order effects such as the ones neglected in this linearized dynamics that could eventually induce bit-flips by displacing the

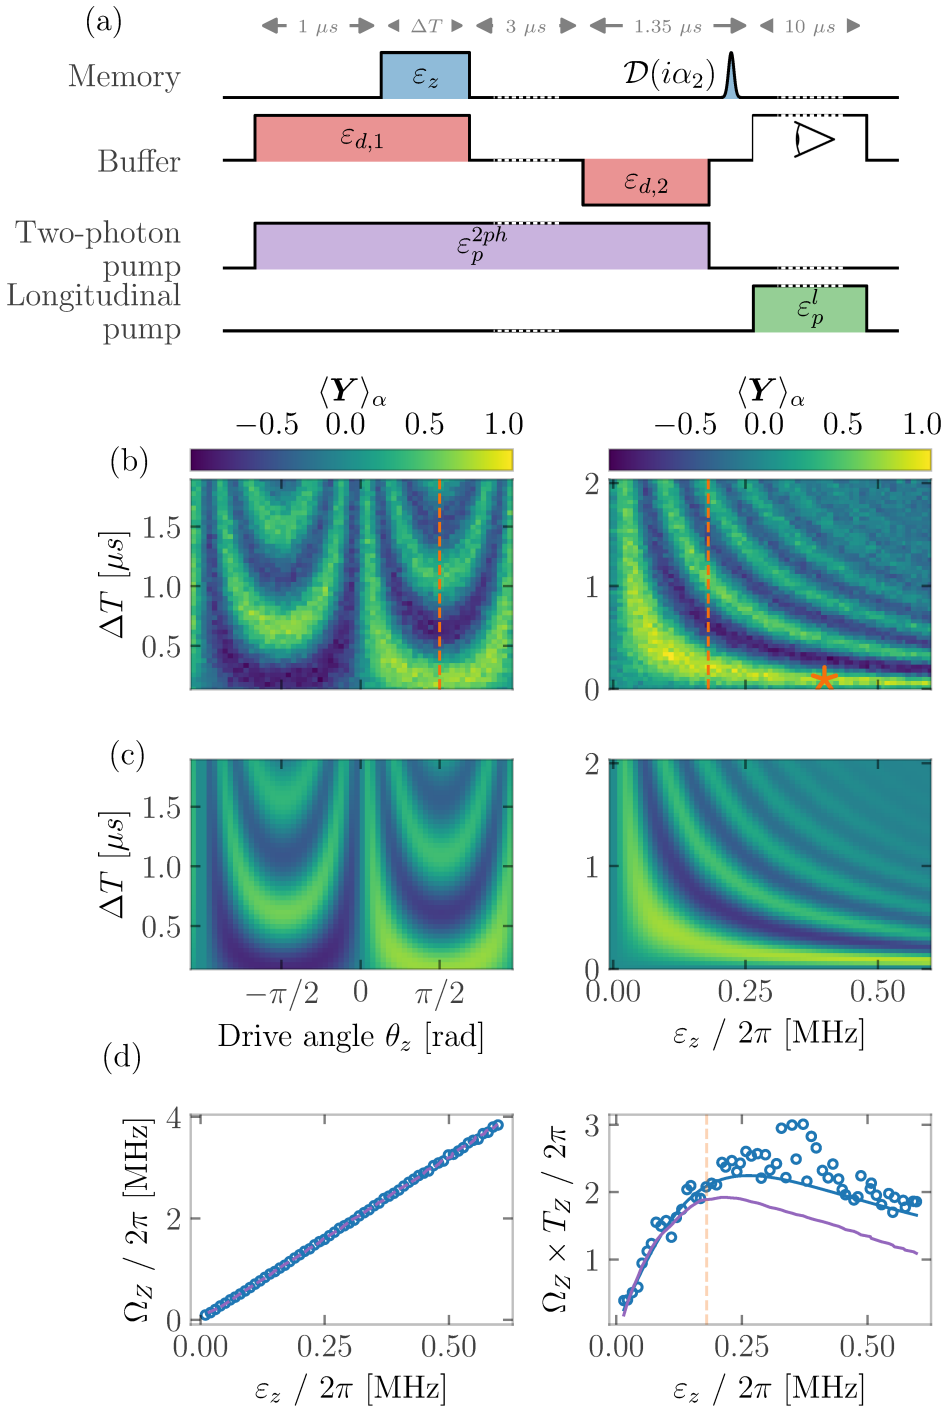

Figure 6.14 Rabi-Zeno one a dissipative cat qubit.

memory too far from its steady state. In the next section, we employ this technique to measure time trajectories.

6.3 Performances of the transmon free dissipative cat qubit

The measurements of phase-flip and bit-flip times of our cat-qubit are displayed in Fig. 6.16. We prepare $|+\rangle_\alpha$ for various average photon numbers $|\alpha|^2$ by starting from a memory mode in the vacuum, and activating the corresponding buffer drive amplitude and two-photon pump. The preparation duration (400 ns for Fig. 6.16 or 1 μs for Fig. 6.15), is chosen to be longer than $1/\kappa_{\text{conf}} \approx 120$ ns, ensuring sufficient time to reach the steady state manifold, and on par with $T_Z \approx 490$ ns for the largest states at 11.3 photons, ensuring the preservation of measurable quantum coherence. Using our novel tomography tool, we image the Wigner functions of these states and observe interference fringes that take negative values. While the contrast of these fringes reduces with $|\alpha|^2$, they remain visible up to $|\alpha|^2 \approx 11.3$ photons (Fig. 6.15a,b). Note that in the cat-qubit code-space $\langle \boldsymbol{X}_\alpha \rangle = \langle \boldsymbol{P} \rangle$, and hence we extract the phase-flip time by monitoring the photon number parity decay over time. We measure phase-flip times ranging from $T_Z = 2.7$ μs for $|\alpha|^2 = 2.5$ to $T_Z = 490$ ns for $|\alpha|^2 = 11.3$ (Fig. 6.16a). Finally, we monitor the switching between $|\pm\alpha\rangle$ over time (Fig. 6.16b). To this end we prepare $|+\alpha\rangle$ by displacing the memory from the vacuum, before applying the two-photon pump and a buffer drive whose amplitude is adjusted to stabilize $|\pm\alpha\rangle$ for a variable time t. During this time, the state may switch to $|-\alpha\rangle$, causing a bit-flip. We detect the population of $|\alpha\rangle$ at time t by setting the buffer drive to map $|\pm\alpha\rangle \rightarrow |\pm\alpha'\rangle$ where $\alpha' \approx 2.1$, then interrupting the pump and buffer drive, and finally displacing the memory by α'. This maps $|-\alpha'\rangle \rightarrow |0\rangle$ and $|\alpha'\rangle \rightarrow |2\alpha'\rangle$. Next, we activate the longitudinal pump to distinguish between these two states. For bit-flip times exceeding ≈ 100 ms, this method leads to impractically long acquisition times. Instead, for these long bit-flip times that occur at $|\alpha|^2 \gtrsim 7$, we sample the real-time trajectory of the memory field. After initializing the memory in $|\alpha\rangle$ and activating the two-photon exchange, we apply a weak drive of amplitude ε_Z on the memory for 250 μs every millisecond. This slightly displaces the state out of the steady state manifold. In response to this perturbation, the buffer develops an average field amplitude $\langle \mathbf{b} \rangle = \mp \frac{\varepsilon_Z}{2\alpha^* g_2}$ depending on the state $|\pm\alpha\rangle$ in the memory [Gautier et al. 2023]. This field is then integrated by heterodyne detection (Fig. 6.16b, bottom panel) for the pulse duration of $T_{\text{int}} = 250$ μs. For $\alpha^2 \gtrsim 7$, we have $T_Z \ll T_{\text{int}} \ll T_X$ and hence we observe bit-flip events

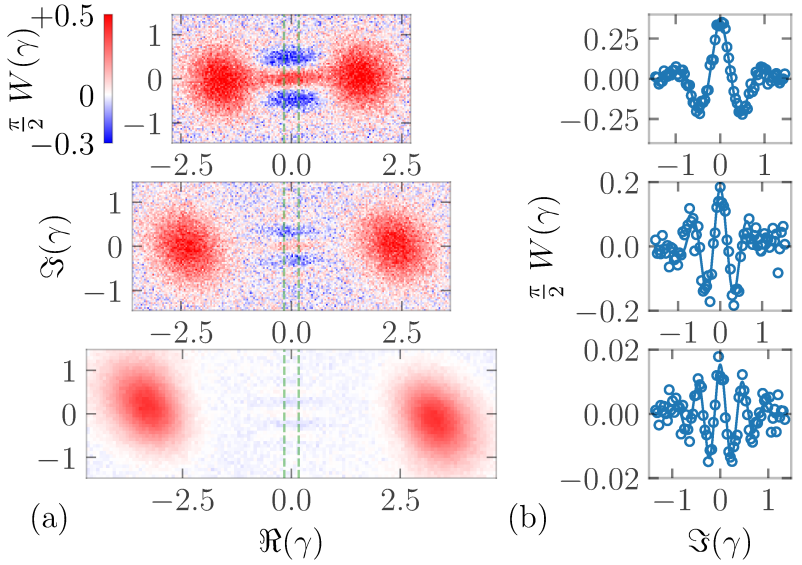

Figure 6.15 Cat states with increasing sizes. Each row of figures represents a photon number $|\alpha|^2 = 2.5, 5.6, 11.3$ (top, middle, bottom). (a) Measured Wigner functions of the memory state $|+\rangle_\alpha$ prepared in 1 μs. Constraining these maps to sum to one sets an absolute scale for our parity measurement. (b) Integration of the measured Wigner functions (y-axis) over the window delimited by the green dashed lines in (a) versus the imaginary axis (x-axis). We fit the analytical formula of these oscillations [Haroche and Raimond 2006a] (solid lines) to the data (open circles).

in real time, so T_X is well estimated from a single trace lasting $\approx 100\ T_X$. Using these methods, we measure T_X for $|\alpha|^2 = 2.5, 5.6, 11.3$, and observe a spectacular increase from 313 μs, to 56 ms, to 15 seconds.

We demonstrate quantum control of our cat-qubit and its impact on bit-flip protection in Fig. 6.17. After preparing the cat-qubit in $|+\rangle_\alpha$, we add a drive of amplitude ε_Z on the memory mode. The interplay of this coherent drive and two-photon dissipation induces Zeno-blocked oscillations [Touzard et al. 2018] at angular-frequency $\Omega_Z = 4|\alpha|\varepsilon_Z$ [Mirrahimi et al. 2014], together with a drive-induced dephasing that increases with ε_Z (see Eq. 2.58). For $|\alpha^2| = 11.3$, we observe a π rotation in 235 ns. Since the benefit of the cat-qubit is to lighten the hardware overhead for error correction by sparing the need for active bit-flip correction, it is crucial to verify that our quantum manipulations do not break bit-flip protection. We measure the

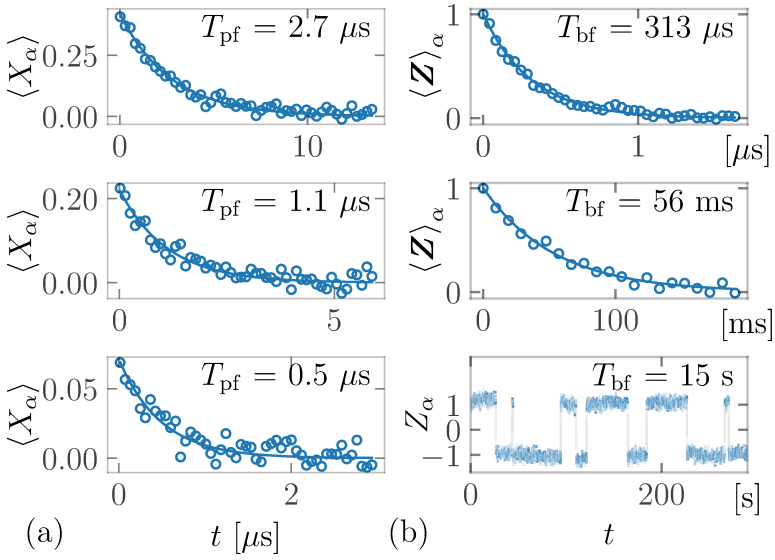

Figure 6.16 Cat-qubit phase-flip and bit-flip time measurements. Each row of figures represents a photon number $|\alpha|^2 = 2.5, 5.6, 11.3$ (top, middle, bottom). (a) Evolution of the expectation value of $\boldsymbol{X}_\alpha$ (y-axis) versus time (x-axis) for a memory state $|+\rangle_\alpha$ prepared in 400 ns. The data (open circles) follow an exponential decay (blue solid line) from which we extract the phase-flip time T_Z. (b-top,middle) Expectation value of $\mathbf{Z}_\alpha$ (y-axis) versus time (x-axis). The data (open circles) follow an exponential decay (solid line) from which we extract the bit-flip time T_X. (b-bottom) Real-time trajectory cropped from the full data-set of the memory switching between $Z_\alpha = \pm 1$ (y-axis) versus time (x-axis). It is acquired by applying a drive on the memory and collecting the buffer fluorescence (solid line).

scaling of errors for $|\alpha|^2 \in [1.4, 11.3]$, a range on which we can measure both the phase-flip and bit-flip rates (Fig. 6.16b,c). We observe the bit-flip time multiply by 4.2 for every added photon, culminating at 15 seconds. Importantly, in the presence of the continuous memory drive, the bit-flip only slightly reduces, remaining above 10 seconds for $|\alpha|^2 = 11.3$. On the other hand, the measured dephasing rate ${T_Z}^{-1}$ increases linearly with $|\alpha|^2$, closely following the theoretical prediction $\Gamma_Z^{\kappa_a} = 2\kappa_a|\alpha|^2$. Notably, the oscillator lifetime extracted from a linear fit to the data is 25% larger than the one obtained from spectroscopy, possibly due to the interplay of the strong two-photon pump and uncontrolled parametric processes.

In conclusion, this experiment demonstrates quantum tomography and coherent control of a cat-qubit without breaking bit-flip protection up to bit-flip times of 10 seconds. This constitutes a 10^4 improvement over other cat-qubit implementations, and a 10^6 enhancement over the oscillator lifetime. We measure a phase-flip time of 490 ns, and perform a π rotation around the Z axis in 235 ns. Although we achieved $g_2/\kappa_a \approx 80 \gg 1$, this ratio needs to be further increased to improve measurement fidelities and reduce state preparation and gate errors below the error correction threshold [Guillaud and Mirrahimi 2019; Gautier et al. 2022]. Possible directions for progress are circuit engineering to increase g_2 [Aiello et al. 2022; Marquet et al. 2023], optimized gate design [Eickbusch et al. 2022; Gautier et al. 2023] and the integration of recent advances in nanofabrication [Place et al. 2021; Wang et al. 2022b; Kono et al. 2023] in order to improve our oscillator lifetime by at least one order of magnitude. With these improvements in hand, we envision assembling multiple cat-qubits in hardware-efficient error correcting architectures [Guillaud and Mirrahimi 2019; Puri et al. 2020; Darmawan et al. 2021; Chamberland et al. 2022b], and operating them to correct phase-flips without breaking bit-flip protection.

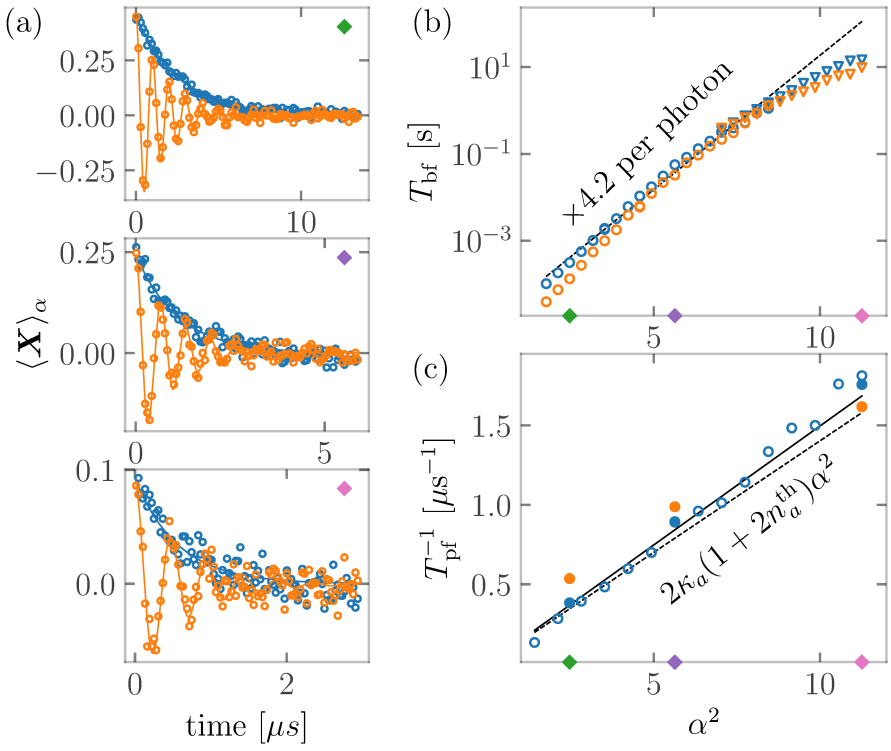

Figure 6.17 Quantum control that preserves bit-flip protection. (a) The cat-qubit is initialized in $|+\rangle_\alpha$ for $|\alpha^2| = 2.5, 5.6, 11.3$ (top, middle, bottom) with a preparation duration of 400 ns, and the expectation value of $\boldsymbol{X}_\alpha$ (y-axis) is measured (open circles) versus time (x-axis). In the presence of a memory drive (orange), the cat-qubit undergoes coherent Zeno-blocked oscillations around its Z axis. These oscillations decay exponentially with time. The decay in absence of oscillations (blue) are superimposed for reference. (b) The bit-flip time (y-axis, log scale) increases exponentially with photon number (x-axis), multiplying by 4.2 for each added photon, up to about 7 photons. It is extracted from the measured (open circles) or simulated (dashed line) average population transfer over time between $|\pm\alpha\rangle$, or from measured single real-time trajectories (triangles). (c) Dephasing rate (y-axis) as a function of photon number (x-axis). It follows a linear trend (solid line) with a slope that is 25% larger than the one given by memory dissipation rate κ_a (dashed line). The data points (circles) corresponding to panels (a) are marked in full dots. The data in (b,c) is taken in the absence (blue) or presence (orange) of a memory drive of same amplitude as in (a).

7

APPENDIX A: PRECISIONS ON THE TWO-MODE MODEL

7.1 Two mode model with detunings

In this model, we introduce a detuning on the buffer Δ_b, representative of a detuning of the buffer drive, and a detuning on the memory Δ_b, which can be caused by a detuning of the drive of the buffer drive of the parametric pump. This approach is reflective of real-life scenarios, where frequencies are not perfectly calibrated. Accounting for detunings allows us to analyze the bandwidth of the processes under study. We modify accordingly the two-model model of Eq. 2.40:

$$\begin{aligned} \mathbf{H}/\hbar &= -\Delta_a \mathbf{a}^\dagger \mathbf{a} - \Delta_b \mathbf{b}^\dagger \mathbf{b} + (g_2 {\mathbf{a}^\dagger}^2 \mathbf{b} - \varepsilon_d \mathbf{b}^\dagger + \text{h.c.}) \\ \boldsymbol{L}_b &= \sqrt{\kappa_b}\mathbf{b} \\ \boldsymbol{L}_a &= \sqrt{\kappa_a}\mathbf{a} \end{aligned} \tag{7.1}$$

We can write the equation of motion seen by $\mathbf{a}$ and $\mathbf{b}$. From there, we can gain further insight by performing the adiabatic elimination of the buffer. Following [Azouit et al. 2016; Berdou et al. 2023], we get the following model.

$$\begin{aligned} \mathbf{H}/\hbar &= -\Delta_a \mathbf{a}^\dagger \mathbf{a} + \frac{\Delta_b g_2^2}{\Delta_b^2 + \kappa_b^2/4} {\mathbf{a}^\dagger}^2 \mathbf{a}^2 - \left(\frac{\varepsilon_d g_2}{\Delta_b + i\kappa_b/2} {\mathbf{a}^\dagger}^2 + \text{h.c.} \right) \\ \boldsymbol{L}_{a^2} &= \sqrt{\frac{\kappa_b g_2^2}{\Delta_b^2 + \kappa_b^2/4}} \mathbf{a}^2 \\ \boldsymbol{L}_a &= \sqrt{\kappa_a}\mathbf{a} \end{aligned} \tag{7.2}$$

If the deutnings are zero, we find Eq. 2.40. After adiabatic elimination, Δ_b gets converted to an effective self-Kerr of the memory.

It is also very insightful to compute the number of photons in the steady state as a function of the detunings, to this end, we write the semi-classical model derived from Eq. 7.1:

$$\begin{aligned}\frac{da}{dt} &= \left(i\Delta_a - \frac{\kappa_a}{2}\right) a - 2ig_2 a^* b \\ \frac{db}{dt} &= \left(i\Delta_b - \frac{\kappa_b}{2}\right) b - ig_2^* a^2 + i\varepsilon_d\end{aligned} \tag{7.3}$$

In the steady state, this system may admit 1, 3, or 5 solutions denoted α and β. Performing a comprehensive stability analysis of these solutions in all cases is challenging, especially since the system may exhibit metastable non-zero solutions, making the conversation complex. In practice, it appears that assuming the solution with the highest photon number in the memory is the most stable is a reasonable empirical approximation.

Now, when it comes to solving this system, it all begins with a crucial observation: there is always a solution where the memory is in the vacuum state. On the other hand, assuming that a is not equal to zero, we can divide the first equation by $|a|$, and inject it into the second equation, leading to:

$$\begin{aligned}|\alpha|^2 &= \frac{\varepsilon_d}{g_2^*} e^{-2i\arg(\alpha)} + z \\ z &= \frac{(i\kappa_a/2 + \Delta_a)(i\kappa_b/2 + \Delta_b)}{2|g_2|^2}\end{aligned} \tag{7.4}$$

The separation of $|\alpha|$ and $\arg(\alpha)$ on opposite sides of the equation allows for an elegant geometric interpretation of the system in the complex plane. The right-hand side describes a circle with a radius of $|\varepsilon_d/g_2|$ centered on z, while the left-hand side represents the positive real axis. In this representation, there can be 0, 1, or 2 intersections between this circle and the positive real axis, resulting in 1, 3, or 5 solutions for the system. This corresponds to one solution for the vacuum state and two solutions for each intersection, since both $\pm\alpha$ are valid solutions. A comprehensive overview of the system solutions is provided in Fig. 7.1 [Leghtas et al. 2015].

$$\begin{aligned}|\alpha|^2 &= \frac{\varepsilon_d}{g_2^*} e^{-2i\arg(\alpha)} + z \\ z &= \frac{(i\kappa_a/2 + \Delta_a)(i\kappa_b/2 + \Delta_b)}{2|g_2|^2}\end{aligned} \tag{7.5}$$

Examining the photon number plotted against detuning reveals a distinctive diagram, often referred to as the diamond due to its rhombus shape. Rich physics is intertwined on its borders. We've previously encountered

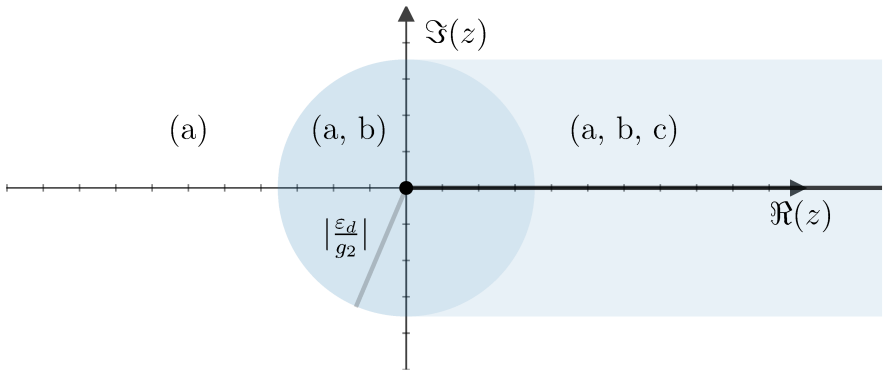

(a)	(b)	(c)
$\alpha_\infty = 0$	$\lvert\alpha_\infty\rvert^2 = \Re(z) + u$	$\lvert\alpha_\infty\rvert^2 = \Re(z) - u$
	$u = \sqrt{\lvert\varepsilon_d/g_2\rvert^2 - \Im(z)}$	
$\beta_\infty = \frac{-\varepsilon_d}{i\Delta_b + i\kappa_b/2}$	$\arg(\alpha_\infty) = \frac{\pi}{2}(1 \pm 1) - \frac{1}{2}\arg(\lvert\alpha_\infty\rvert^2 - z)$	
	$\beta_\infty = \frac{\Delta_a + i\kappa_a/2}{2g_2} e^{2i\arg(\alpha_\infty)}$	

Figure 7.1 Geometrical resolution of the two-mode model steady state. Depending on the value of the complex parameter z, the system can have a unique solution where the memory is in the vacuum state (white area). If the circle centered on z with a radius of ε_d/g_2 intersects the positive real axis once, then the system acquires two more admissible solutions (dark blue area). Finally, if the circle crosses the positive real axis twice, the system has a total of five solutions (light blue area).

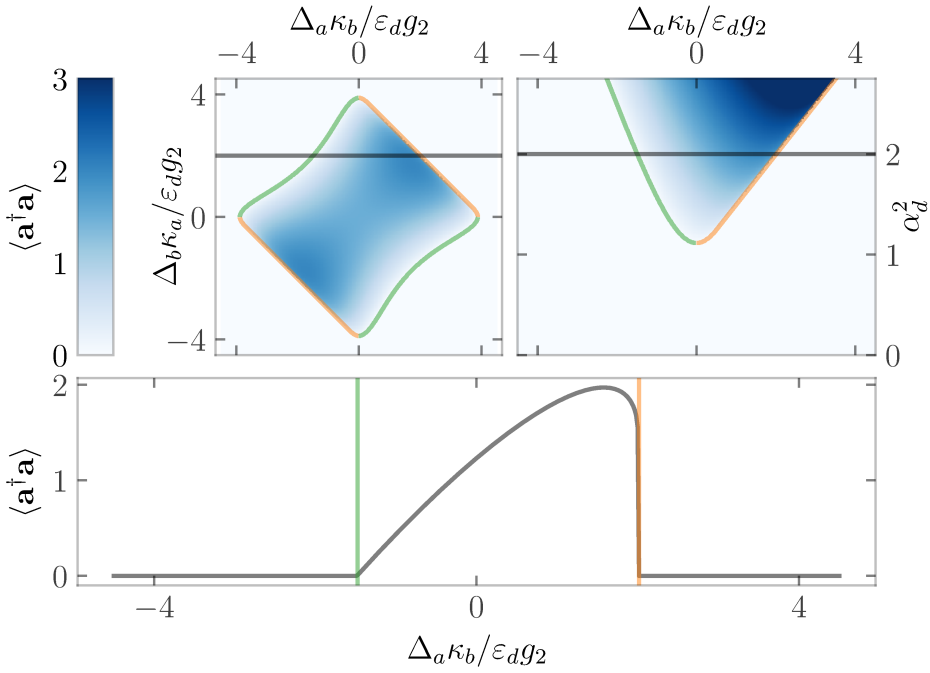

Figure 7.2 Semi-classical response of the two-mode model in the regime $\kappa_1/\kappa_2 \approx 300$. The memory photon number in the steady state is displayed against different experimental parameters. In the top-left panel it is plotted against detuning on the buffer (y-axis) and on the memory (x-axis). The top-right panel shows it against the memory detuning (x-axis) and the two-photon drive amplitude (y-axis). The bottom panel focusses on the black solid line visible on both top panels. Green lines delineate the second order dissipative phase transition, while the orange lines indicate the first order phase transition. In normalized unites, the slope of this line against memory and buffer detunings is -1.

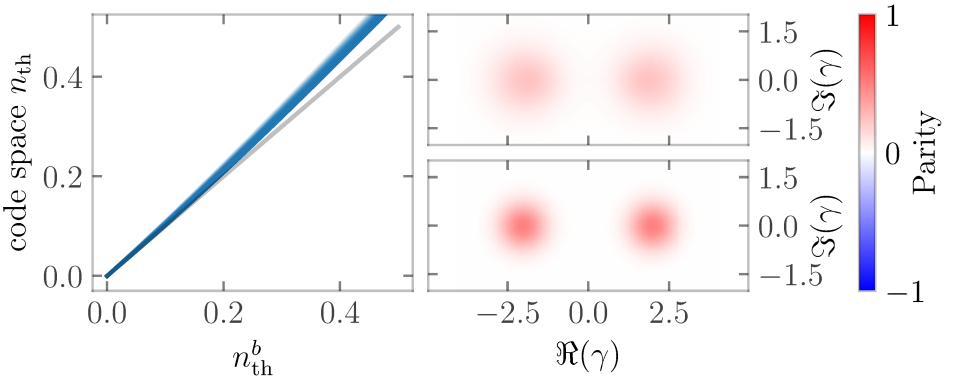

Figure 7.3 Impact of the buffer thermall occupation.

an example of a second-order dissipative phase transition, and this diamond plot showcases instances of both first and second-order transitions. From a practical perspective, there is particular interest in analytically determining the shape of the region bounded by the first-order transition.

The upper right and bottom left edges of the diamond are two lines of same slope κ_a/κ_b, a handy rule of thumb to bear in mind.

$$\kappa_a \Delta_b + \kappa_b \Delta_a = \pm 4 |\varepsilon_d g_2| \tag{7.6}$$

7.2 Impact of thermal population

In this section, we aim to provide some insights and present numerical results regarding the influence of thermal populations in both the memory and the buffer on the cat qubit code space.

Incorporating finite temperatures into both the memory and the buffer amounts to replace of the one-photon loss operators $\boldsymbol{L}_a$ and $\boldsymbol{L}_b$ in Eq. 2.40 with:

$$\begin{aligned}
\boldsymbol{L}_{a\uparrow} &= \sqrt{n_{\text{th}}^a \kappa_a}\,\mathbf{a}^\dagger \\
\boldsymbol{L}_{a\downarrow} &= \sqrt{(1+n_{\text{th}}^a)\kappa_a}\,\mathbf{a} \\
\boldsymbol{L}_{b\uparrow} &= \sqrt{n_{\text{th}}^b \kappa_b}\,\mathbf{b}^\dagger \\
\boldsymbol{L}_{b\downarrow} &= \sqrt{(1+n_{\text{th}}^b)\kappa_b}\,\mathbf{b}
\end{aligned} \tag{7.7}$$

Considering two-photon stabilization, thermal populations in the buffer and memory modes assume asymmetrical roles. In Fig. 7.3, we fit a displaced Gaussian to the logical zero state $|0\rangle_\alpha$ for various thermal occupations in both memory and buffer modes, plotted against the two-photon drive α_d^2. In the limit of a large drive, the thermal population of the memory becomes negligible as the confinement rate increases, suppressing the effect of the local operator $\boldsymbol{L}_{a\uparrow}$. However, in the displaced frame around α and after adiabatic elimination, the thermal occupation of the buffer converts to thermal occupation of the displaced operator $\tilde{\mathbf{a}}$. Consequently, the logical zero and one states of the code space are now displaced thermal states with a thermal occupation of n_{th}^b. A convenient parameterization for the code space, spanned by these states, is as follows:

$$\begin{aligned}\boldsymbol{\rho} &= \frac{1}{\mathcal{N}}\begin{bmatrix}\mathbf{D}(\alpha) & \mathbf{D}^\dagger(\alpha)\end{bmatrix}\begin{bmatrix}1+Z & X-iY\\ X+iY & 1-Z\end{bmatrix}\begin{bmatrix}\boldsymbol{\rho}_{\text{th}}\mathbf{D}^\dagger(\alpha)\\ \boldsymbol{\rho}_{\text{th}}\mathbf{D}(\alpha)\end{bmatrix}\\ \mathcal{N} &= 2 + 2X\,\text{Tr}\left(\mathbf{D}(\alpha)\boldsymbol{\rho}_{\text{th}}\mathbf{D}^\dagger(\alpha)\boldsymbol{\rho}_{\text{th}}\right)\end{aligned} \tag{7.8}$$

Where $\boldsymbol{\rho}_{\text{th}}$ is a thermal state described in Eq. 5.14, with thermal occupation n_{th}^b.

7.3 Flux map

The dependence between the frequency of the buffer and the applied flux is easy to observe and provides valuable information. In fact, this measurement is an essential step in the device calibration procedure. To model this situation, we consider an ATS with a parallel capacitor of energy E_C. In this simplified picture, the mode frequency $\omega_b/2\pi$ can be computed as follows:

$$\omega_b = \frac{1}{\hbar}\sqrt{8E_C E_{L,\text{eq}}} \quad \propto \quad \sqrt{E_{L,\text{eq}}} \tag{7.9}$$

In this expression, $E_{L,\text{eq}}$ is the equivalent inductance of the ATS for a given flux point. In a classical picture, this energy is determined by the curvature of the potential of the ATS U_{ATS} (see Eq. 4.22) around its global minimum:

$$E_{L,\text{eq}} = \frac{\partial^2 U_{\text{ATS}}}{\partial\varphi^2}\Big|_{\arg\min U_{\text{ATS}}(\varphi)} \tag{7.10}$$

To simplify the analysis, we non-dimensionalize the problem by scaling it to dimensionless units.

$$
\begin{aligned}
E_J^* &= 2\left((E_J\cos(\varphi_\Sigma))^2 + (\Delta E_J\sin(\varphi_\Sigma))^2\right)^2/E_L \\
\varphi^* &= \arg(E_J\cos(\varphi_\Sigma)\cos(\varphi_\Delta) - \Delta E_J\sin(\varphi_\Sigma)\sin(\varphi_\Delta) \\
&\quad + i(E_J\cos(\varphi_\Sigma)\sin(\varphi_\Delta) + \Delta E_J\sin(\varphi_\Sigma)\cos(\varphi_\Delta)))
\end{aligned}
\tag{7.11}
$$

In terms of these variables, we can define the dimensionless potential of the ATS:

$$
U^*(\varphi) = \frac{1}{2}\varphi^2 - E_J^*\cos(\varphi+\varphi^*) \tag{7.12}
$$

Finding the global minima of this potential amounts to finding the smallest solution to an equation involving a sine and a linear term:

$$
\frac{\partial U^*}{\partial\varphi} = \varphi + E_J^*\sin(\varphi+\varphi^*) = 0 \tag{7.13}
$$

For efficient numerical resolution, we solve this system once and for all, varying E_J^* and φ^*, we record the values of the second derivative of $U^*(\varphi)$ in Fig. 7.4. From there, computing the flux map of the ATS consists merely in recasting the variables of the problem in the appropriate units.

The flux map exhibits several distinctive features specific to the ATS. When plotted against common and differential flux, it remains invariant under translations of $(\pi, \pm\pi)$. The frequency is maximized at zero flux. The key feature of the ATS is that it exhibits points which frequencies are at first order insensitive to flux variations: saddle points. Remarkably, these points are also optimal for implementing the four-wave mixing responsible for the two-photon exchange mechanism.

Assuming no asymmetry ($\Delta E_J = 0$), all the saddle points lie at the same frequency and are also free Kerr. Based on this observation and considering the desirability of these two properties, a natural question arises: What is the simplest Josephson-based circuit that features operating points that are both Kerr-free and frequency saddle points? Specifically, how few flux loops can this circuit have? The writer's intuition regarding this question is that the answer is two, and somehow, the ATS may represent the simplest solution in a yet-to-be-defined sense.

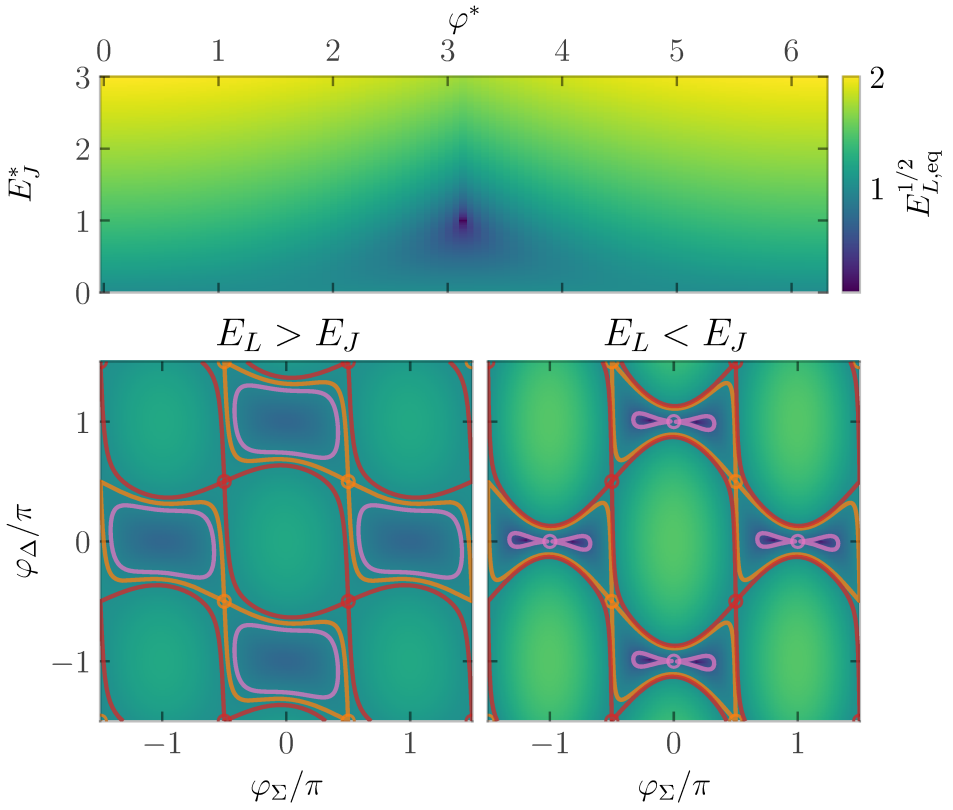

Figure 7.4 ATS frequency vs. common and differential flux. The equivalent frequency $E_{L,\mathrm{eq}}^{1/2}$ is displayed (top) as a function of the dimensionless energy E_J^* and phase φ^*. The ATS flux maps in the bottom panels are interpolated from this data for various parameter sets. These maps are 2π-periodic with respect to the common flux φ_Σ (x-axis) and the differential flux φ_Δ (y-axis). When the junction energy E_J is smaller than the inductance energy E_L (bottom-left) the flux map displays two classes of saddle points (orange and red solid lines) with a frequency difference that depends on the ATS asymmetry ΔE_J. The isofrequency curve just below the saddle points (pink solid line) has the shape of a peanut, with its main axis coinciding with the common flux direction. For junction energy greater than that of the inductance (bottom-right) the peanut-shaped curve is pinched, and a third class of saddle point appears. However, these points are somewhat artificial, as they result from the kissing of two parabolas and correspond to a regime where the ATS is bistable.

8

APPENDIX B: EFFICIENT RESONATOR FIT TECHNIQUE

This appendix outlines an efficient method for analyzing signal outputs from resonator spectroscopy across various geometries. Fitting resonators constitutes the bread and butter of the CQED engineer. Therefore, having fast and reliable procedures for this task is essential.

To the best of my knowledge, the literature currently lacks a satisfactory method for fitting the scattering parameter of a resonator efficiently, except for the ”circle fit” technique described in [Probst et al. 2015] from 2014. However, this method is imperfect and suffers from two main drawbacks:

First, it merely attempts to fit a circle to the data on the complex plane. While it is true that the data points generally form a circle, this approach neglects the dynamic information encoded as the frequency ω sweeps through the resonance. We will demonstrate that the gradient of the data provides significant insights. Spectroscopy should be viewed as a 1D manifold existing within a 3D space defined by ω, $\Re(S)$, and $\Im(S)$, and the fitting should be executed directly in this space.

Secondly, if the fit for the electrical delay is incorrect, the data will appear to lie on a circle with radius α centered at 0 on the complex plane. Consequently, while the fitting process may technically succeed, the parameters it yields will be entirely inaccurate.

In this text, we introduce an analytical method designed to address this fitting challenge more robustly. This method has been implemented in a public GitHub repository `https://github.com/UlysseREGLADE/abcd_rf_fit` and is already in use by several groups. It is colloquially known as abcd-rf-fit.

8.1 Scattering Parameter Formula

We adopt the convention described by [Gardiner and Collett 1985] for the input/output relationship:

$$a^i_{out} - a^i_{in} = \sqrt{\kappa_c} a \tag{8.1}$$

$$\frac{\partial a}{\partial t} = -\frac{i}{\hbar}\frac{\partial H}{\partial a^\dagger} - \frac{\kappa_i}{2}a - \sum_i \left(\frac{\kappa^i_c}{2} a + \sqrt{\kappa^i_c} a^i_{in} \right) \tag{8.2}$$

We are going to consider a single resonator with internal losses κ_i coupled to one or two lines at rates κ^1_c and κ^2_c. In the rotating frame at frequency ω, we consider the Hamiltonian of a bare linear resonator:

$$\mathbf{H}/\hbar = (\omega - \omega_0) a^\dagger a \tag{8.3}$$

Note that in this expression, the resonator frequency appears negative. This choice of convention, consistent with [Gardiner and Collett 1985], affects how the scattering parameter rotates in the complex plane—either clockwise or anticlockwise. It is important to remember that this rotation direction might vary depending on the VNA brand used to collect the data.

Now, the scattering parameter $S_{ij}(\omega)$ is defined by:

$$S_{ij}(\omega) = \frac{a^i_{out}}{a^j_{in}} \tag{8.4}$$

8.1.1 Transmission

We consider coupled to two antennas:

$$S_{21}(\omega) = S_T(\omega) = \frac{a^2_{out}}{a^1_{in}} \tag{8.5}$$

Which gives us:

$$S_T(\omega) = \frac{2\sqrt{\kappa^1_c \kappa^2_c}}{2i(\omega - \omega_0) + \kappa_i + \kappa^1_c + \kappa^2_c} \tag{8.6}$$

The only fittable parameters in this formula are ω_0 and $\kappa = \kappa_i + \kappa^1_c + \kappa^2_c$. Since the numerator is constant, it will be entirely eaten by the amplification coefficient, hence we write:

$$S_T(\omega) = \frac{1}{2i(\omega - \omega_0) + \kappa} \tag{8.7}$$

8.1.2 Reflection

We consider a cavity coupled to a single antenna.

$$S_{11}(\omega) = S_R(\omega) = \frac{a^1_{out}}{a^1_{in}} \tag{8.8}$$

We find:

$$S_R(\omega) = \frac{i(\omega - \omega_0) + (\kappa_i - \kappa_c)/2}{i(\omega - \omega_0) + (\kappa_i + \kappa_c)/2} \tag{8.9}$$

In this case, both ω_0, κ_i and κ_c are fittable, which makes a reflection measurement much more valuable than a transmission one. Both the transmission and reflection formula are derived in [Campagne-Ibarcq 2015].

8.1.3 Hanger

To get the hanger scattering parameter, one should modify the input/output relation as follows:

$$a^i_{out} - a^i_{in} = \frac{\sqrt{\kappa^i_c}}{2} a \tag{8.10}$$

In this context, we can now compute:

$$S_{11}(\omega) = S_H(\omega) = \frac{a^1_{out}}{a^1_{in}} \tag{8.11}$$

We find:

$$S_H(\omega) = \frac{i(\omega - \omega_0) + \kappa_i/2}{i(\omega - \omega_0) + (\kappa_i + \kappa_c)/2} \tag{8.12}$$

Again, ω_0, κ_i and κ_c are fittable.

8.1.4 Hanger with impedance mismatch

In [Probst et al. 2015] is described how to take into account the effect of an impedance mismatch in the context of the hanger geometry. In this case, κ_c can be viewed as a complex number, in this situation, we write:

$$\kappa_c = |\kappa_c| e^{i\phi_0} \tag{8.13}$$

In this case, only the real part of the complex coupling rates contribute to the total loss-rate of the cavity κ.

$$\kappa = \kappa_i + \Re(\kappa_c) \tag{8.14}$$

One should introduce the parameter ϕ_0 in the scattering parameter of the hanger as follow:

$$S_{HM}(\omega) = \frac{2i(\omega - \omega_0) + \kappa - \Re(\kappa_c)(1 + i\tan(\phi_0))}{2i(\omega - \omega_0) + \kappa} \tag{8.15}$$

From a geometrical point of view in the complex plan, it allows the rotation of the circle described by $S_H(\omega)$ around 1. Empirically, one can also see it as a modification of the input/output relation:

$$a^i_{out} - a^i_{in} = \frac{\sqrt{|\kappa^i_c|e^{i\phi_0}}}{2}a \tag{8.16}$$

In practise, adding this degree of freedom to the formula loosens the precision of the fit of κ_i and κ_c. Hence, one should be cautious when allowing for this offset.

8.1.5 Reflection with impedance mismatch

Taking inspiration for what was done for the hanger, we allow the rotation around 1 in the complex plan by writing the input/output relation as follows. However, this is more a mathematical trick than real physics:

$$a^i_{out} - a^i_{in} = \sqrt{|\kappa^i_c|e^{i\phi_0}}a \tag{8.17}$$

The modified scattering parameter for the reflection geometry is as follow:

$$S_{RM}(\omega) = \frac{2i(\omega - \omega_0) + \kappa - 2\Re(\kappa_c)(1 + i\tan(\phi_0))}{2i(\omega - \omega_0) + \kappa} \tag{8.18}$$

Again, the precision of the fits of κ_i and κ_c are loosened, and this formula is to be taken with the same caveat than the previous one. abcd-rf-fit will return a warning when the value of ϕ_0 in greater than 0.25 .

8.1.6 Effect of electrical delay and amplification chain

In a real life scenario of circuit QED, the scattering parameter of a resonator is always dressed.

First, since these resonators are meant to operate at very low energy scales, one always uses an extensive attenuation chain to send signal in, and a powerful amplification chain to retrieve the outputted signal. This results in an arbitrary complex multiplicative prefactor $\alpha \in \mathbb{C}$.

Second, the resonator is always at a finite distance from the instrument used to measure it. As a consequence, the phase of the outputted signal will vary as a function of its frequency depending on this distance and the speed of light in the medium that carries the signal. This has the effect to multiply the scattering parameter by a factor $e^{2i\pi\lambda\omega}$, where $\lambda \in \mathbb{R}$ is a typical time called the electrical delay that encodes for cables length and the speed of light among them.

Hence, what is actually to be fitted in most cases is:

$$S(\omega) = \alpha \times S_X(\omega) \times e^{2i\pi\lambda\omega} \tag{8.19}$$

$$X = T, R, H, RM, HM \tag{8.20}$$

8.2 Fit of a rationnal function of degree one

Ingnoring the electrical for now, one can observe that all the scattering parameters we described can be written in the form:

$$S_X(\omega) = \frac{a + b\omega}{c + d\omega}, (a, b, c, d) \in \mathbb{C} \tag{8.21}$$

We propose an efficient procedure to extract these coefficients from a noised signal.

8.2.1 Side note: least square regression of an n-degree polynomial

As described in the [Bach 2020], given a set of points $(x_i, y_i), i \in [\![1, N]\!]$, one can write the empirical risk $R(w)$ associated with the least square regression of an n-degree polynomial with coefficients $w_0, ..., w_n$ as follow:

$$R_{LS}(w) = \frac{1}{N}||Xw - y||_2^2 \tag{8.22}$$

with:

$$X = \begin{bmatrix} 1, x_1, \dots, x_1^n \\ \vdots \\ 1, x_N, \dots, x_N^n \end{bmatrix}, y = \begin{bmatrix} y_1 \\ \vdots \\ y_N \end{bmatrix} w = \begin{bmatrix} w_0 \\ \vdots \\ w_N \end{bmatrix} \tag{8.23}$$

We can now easily solve for $\nabla_w R(w) = 0$ to find the global minimum of this convex optimization problem:

$$w = (X^T X)^{-1} X y \tag{8.24}$$

Not only is this formulation extremely elegant, it allows for very fast computation. It is actually the one implemented in the Python library numpy in the function 'polyfit'.

We take inspiration from this approach in the next section to find an efficient fitting procedure in our case.

8.2.2 Least square regression of a degree one rationnal function

We are given a set $(\omega_i, s_i), i \in [\![1, N]\!]$, where ω_i is a frequency and s_i the signal measured by the instrument, for instance a Vectorial Network Analyser (VNA).

We would like the find the global minimum of the least square risk. In this context it reads:

$$R_{LS}(a, b, c, d) = \frac{1}{N} \sum_i \left| s_i - \frac{a + b\omega_i}{c + d\omega_i} \right|^2 \tag{8.25}$$

It is not obvious how to minimize this risk analytically. To this end, we introduce the following matrices:

$$X = \begin{bmatrix} 1, \omega_1 \\ \vdots \\ 1, \omega_N \end{bmatrix}, y = \begin{bmatrix} s_1 \\ \vdots \\ s_N \end{bmatrix} w = \begin{bmatrix} a \\ b \end{bmatrix}, m = \begin{bmatrix} c \\ d \end{bmatrix} \tag{8.26}$$

We now define the following empirical risk:

$$R(w, m) = ||\mathcal{D}(y) X m - X w||_2^2 \tag{8.27}$$

Where $\mathcal{D}$ is a shorthand for the diagonal matrix with the elements of y on its diagonal:

$$\mathcal{D}(y) = \begin{bmatrix} y_1 & & \\ & \ddots & \\ & & y_N \end{bmatrix} \tag{8.28}$$

Note that this risk is not yet the one associated with the least square regression, however its zeros include the one of the least square risk. Indeed, if $R_{LS}(w,m) = 0$ then we have $R(w,m) = 0$.

We now solve for $\nabla_w R(w,m) = \nabla_m R(w,m) = 0$:

$$X^\dagger X w - X^\dagger \mathcal{D}(y) X m = 0 \tag{8.29}$$

$$X^\dagger \mathcal{D}(|y|^2) X w - X^\dagger \mathcal{D}(y^*) X m = 0 \tag{8.30}$$

To compress this expression, we write:

$$A = (X^\dagger X)^{-1} X^\dagger \mathcal{D}(y) X \tag{8.31}$$

$$B = (X^\dagger \mathcal{D}(|y|^2) X)^{-1} X^\dagger \mathcal{D}(y^*) X \tag{8.32}$$

$$C = \begin{bmatrix} 0_2, A \\ B, 0_2 \end{bmatrix} \tag{8.33}$$

Leading to:

$$\begin{bmatrix} w \\ m \end{bmatrix} = C \begin{bmatrix} w \\ m \end{bmatrix} \tag{8.34}$$

Diagonalizing this 4×4 matrix and looking for its eigenvector associated with the eigenvalue 1 allows to efficiently solve for a, b, c, d.

This is very nice and efficient to compute, though in the case of noised signals the fit while eventually fail since this risk doesn't coincide with $R_{LS}(w,m)$.

To solve this issue, one can observe that the $R(w,m)$ we introduce corresponds to $R_{LS}(w,m)$ where each data point has been weighted by $|c+d\omega|$ which has the bad taste of being maximal away from the resonance. Fortunately for us, it happens that the empirical gradient of the signal encodes for this quantity:

$$\sqrt{|\nabla_\omega S_X(\omega)|} = \sqrt{\left|\frac{b}{d}\left(\frac{c}{d} - \frac{a}{b}\right)\right|} \frac{1}{|c + d\omega|} \propto \frac{1}{|c + d\omega|} \tag{8.35}$$

You can now pick your preferred empirical estimator of the gradient $\tilde{\nabla}_X(y)$ such as the convolution with the derivative of a gaussian kernel for instance. In the limit where this estimation of the gradient is accurate, we have $R_{LS}(w,m) = R^*(w,m)$, where:

$$R^*(w,m) = \frac{1}{N}||\mathcal{D}\left(\sqrt{|\tilde{\nabla}_X(y)|}\right)(\mathcal{D}(y)Xm - Xw)||_2^2 \tag{8.36}$$

The corrected values of the matrices A and B are:

$$A = (X^\dagger \mathcal{D}(|\tilde{\nabla}_X(y)|)X)^{-1}X^\dagger \mathcal{D}(y|\tilde{\nabla}_X(y)|)X \tag{8.37}$$
$$B = (X^\dagger \mathcal{D}(|\tilde{\nabla}_X(y)||y|^2)X)^{-1}X^\dagger \mathcal{D}(y^*|\tilde{\nabla}_X(y)|)X \tag{8.38}$$

Performing the conversion from the coefficients a, b, c, d to the relevant quantities depending on the resonator geometry is now only a matter of simple algebra.

8.2.3 Estimation of the electrical delay

Assuming a flat frequencial landscape, the best estimator of the electrical delay one can write is the following:

$$\lambda_{\text{flat}} = \frac{1}{N}\sum_{i=0}^{N-1}\frac{\arg(s_{i+1}/s_i)}{\omega_{i+1}-\omega_i} \tag{8.39}$$

Please observe that this estimator should be much more robust to noised signals than performing an unwrapping of the signal phase such as the one implemented in the Python library Numpy in the function ‘unwrap‘.

In the case of a simple resonator, the signal performs at worst a full revolution around the origin of the complex plane. Since on modern hardware the diagonalisation of a four by four matrix such as C is fast, one can brute force all the values of lambda in the range:

$$\lambda \in \left[\lambda_{\text{flat}} - \frac{1.5}{\max_i(\omega_i)-\min_i(\omega_i)}, \lambda_{\text{flat}} + \frac{1.5}{\max_i(\omega_i)-\min_i(\omega_i)}\right] \tag{8.40}$$

This simple method proves to be really efficient.

BIBLIOGRAPHY

Gidney, Craig and Martin Ekerå (1905). "How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits (2019)". In: *arXiv preprint arXiv:1905.09749* (cit. on p. 37).

Ambegaokar, Vinay and Alexis Baratoff (1963). "Tunneling between superconductors". In: *Physical review letters* 10.11, p. 486 (cit. on p. 86).

Matthaei, GL, Leo Young, and EM Jones (1963). "Design of microwave filters impedance-matching networks and coupling structures volume 2". In: *vol* 1, p. 526 (cit. on p. 97).

Halbritter, J (1974). "On surface resistance of superconductors". In: *Zeitschrift für Physik* 266.3, pp. 209–217 (cit. on p. 92).

Guckenheimer, John and Philip Holmes (1983). "Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields". In: New York, NY: Springer New York, pp. 1–65 (cit. on p. iii).

Gardiner, Crispin W and Matthew J Collett (1985). "Input and output in damped quantum systems: Quantum stochastic differential equations and the master equation". In: *Physical Review A* 31.6, p. 3761 (cit. on pp. 45, 93, 178, 179).

Wolinsky, M. and H. J. Carmichael (1988). "Quantum noise in the parametric oscillator: From squeezed states to coherent-state superpositions". In: *Phys. Rev. Lett.* 60 (18), pp. 1836–1839 (cit. on p. iv).

Yurke, B. and D. Stoler (1988). "The dynamic generation of Schrödinger cats and their detection". In: *Physica B+C* 151.1, pp. 298–301 (cit. on p. 157).

Shor, Peter W (1994). "Algorithms for quantum computation: discrete logarithms and factoring". In: *Proceedings 35th annual symposium on foundations of computer science*. Ieee, pp. 124–134 (cit. on p. 16).

Devoret, Michel H et al. (1995). "Quantum fluctuations in electrical circuits". In: *Les Houches, Session LXIII* 7.8, pp. 133–135 (cit. on p. 81).

Kitaev, A Yu (1995). "Quantum measurements and the Abelian stabilizer problem". In: *arXiv preprint quant-ph/9511026* (cit. on p. 17).

Shor, Peter W (1995). "Scheme for reducing decoherence in quantum computer memory". In: *Physical review A* 52.4, R2493 (cit. on pp. 32, 33).

Bennett, Charles H., David P. DiVincenzo, John A. Smolin, and William K. Wootters (Nov. 1996). "Mixed-state entanglement and quantum error correction". In: *Physical Review A* 54.5, pp. 3824–3851 (cit. on p. 32).

Ekert, A. and C. Macchiavello (1996). *Error Correction in Quantum Communication*. arXiv: quant-ph/9602022 [quant-ph] (cit. on p. 32).

Lutterbach, LG and L Davidovich (1997). "Method for direct measurement of the Wigner function in cavity QED and ion traps". In: *Physical review letters* 78.13, p. 2547 (cit. on p. 121).

Shor, Peter W. (1997). *Fault-tolerant quantum computation*. arXiv: quant-ph/9605011 [quant-ph] (cit. on p. 32).

Fefferman, Charles L (2000). "Existence and smoothness of the Navier-Stokes equation". In: *The millennium prize problems* 57, p. 67 (cit. on p. 17).

Bertet, Patrice, Alexia Auffeves, Paolo Maioli, Stefano Osnaghi, Tristan Meunier, Michel Brune, Jean-Michel Raimond, and Serge Haroche (2002). "Direct measurement of the Wigner function of a one-photon Fock state in a cavity". In: *Physical Review Letters* 89.20, p. 200402 (cit. on p. 121).

Breuer, Heinz-Peter and Francesco Petruccione (2002). *The theory of open quantum systems*. OUP Oxford (cit. on p. 45).

Dennis, Eric, Alexei Kitaev, Andrew Landahl, and John Preskill (2002). "Topological quantum memory". In: *Journal of Mathematical Physics* 43.9, pp. 4452–4505 (cit. on p. 36).

Goldstein, Herbert, Charles Poole, and John Safko (2002). *Classical mechanics* (cit. on p. 75).

Zurek, Wojciech Hubert (2003). "Decoherence, einselection, and the quantum origins of the classical". In: *Rev. Mod. Phys.* 75 (3), pp. 715–775 (cit. on p. iv).

Hong, Jia-Shen G and Michael J Lancaster (2004). *Microstrip filters for RF/microwave applications*. John Wiley & Sons (cit. on pp. 97, 140).

Gambetta, Jay, Alexandre Blais, D. I. Schuster, A. Wallraff, L. Frunzio, J. Majer, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf (Oct. 2006). "Qubit-photon interactions in a cavity: Measurement-induced dephasing and number splitting". In: *Physical Review A* 74.4 (cit. on p. 70).

Haroche, Serge and J-M Raimond (2006a). *Exploring the quantum: atoms, cavities, and photons*. Oxford university press (cit. on pp. 41, 166).

Haroche, Serge and Jean-Michel Raimond (2006b). *Exploring the Quantum: Atoms, Cavities, and Photons*. Oxford: Oxford University Press (cit. on p. 155).

Abbas, Norliana Mohd, Darius G Solomon, and Md Fuad Bahari (2007). "A review on current research trends in electrical discharge machining

(EDM)". In: *International Journal of machine tools and Manufacture* 47.7-8, pp. 1214–1228 (cit. on pp. 97, 98).

Koch, Jens, M Yu Terri, Jay Gambetta, Andrew A Houck, David I Schuster, Johannes Majer, Alexandre Blais, Michel H Devoret, Steven M Girvin, and Robert J Schoelkopf (2007). "Charge-insensitive qubit design derived from the Cooper pair box". In: *Physical Review A* 76.4, p. 042319 (cit. on p. 106).

Schuster, D. I., A. A. Houck, J. A. Schreier, A. Wallraff, J. M. Gambetta, A. Blais, L. Frunzio, J. Majer, B. Johnson, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf (Feb. 2007). "Resolving photon number states in a superconducting circuit". In: *Nature* 445.7127, pp. 515–518 (cit. on p. 115).

Aliferis, Panos and John Preskill (2008). "Fault-tolerant quantum computation against biased noise". In: *Physical Review A* 78.5, p. 052331 (cit. on pp. iv, 39).

Schreier, Joseph A, Andrew A Houck, Jens Koch, David I Schuster, Bradley R Johnson, Jerry M Chow, Jay M Gambetta, J Majer, Luigi Frunzio, Michel H Devoret, et al. (2008). "Suppressing charge noise decoherence in superconducting charge qubits". In: *Physical Review B* 77.18, p. 180502 (cit. on p. 107).

Barchielli, Alberto and Matteo Gregoratti (2009). *Quantum trajectories and measurements in continuous time: the diffusive case*. Vol. 782. Springer Science & Business Media (cit. on p. 135).

Gottesman, Daniel (2009). *An Introduction to Quantum Error Correction and Fault-Tolerant Quantum Computation*. arXiv: 0904.2557 [quant-ph] (cit. on p. 36).

Huttel, Andreas K, Gary A Steele, Benoit Witkamp, Menno Poot, Leo P Kouwenhoven, and Herre SJ van der Zant (2009). "Carbon nanotubes as ultrahigh quality factor mechanical resonators". In: *Nano letters* 9.7, pp. 2547–2552 (cit. on p. 92).

Petit, Nicolas and Pierre Rouchon (2009). "Automatique: Dynamique et contrôle des systèmes". In: (cit. on p. 62).

Macha, Pascal, SHW van Der Ploeg, G Oelsner, E Il'Ichev, H-G Meyer, S Wünsch, and M Siegel (2010). "Losses in coplanar waveguide resonators at millikelvin temperatures". In: *Applied Physics Letters* 96.6 (cit. on p. 95).

Nielsen, Michael A. and Isaac L. Chuang (2010). *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press (cit. on p. iv).

Mabuchi, Hideo (Oct. 2011). "Nonlinear interferometry approach to photonic sequential logic". In: *Applied Physics Letters* 99.15, p. 153103 (cit. on p. iii).

Pozar, David M (2011). *Microwave engineering*. John wiley & sons (cit. on p. 98).

Fowler, Austin G., Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland (2012). "Surface codes: Towards practical large-scale quantum computation". In: *Phys. Rev. A* 86 (3), p. 032324 (cit. on p. iv).

Girvin (2012). "Superconducting qubits coupled to microwave photons". In: *Quantum Machines: Measurement and Control of Engineered Quantum Systems (Oxford University Press 2011)* (cit. on pp. 80, 106).

Curtright, Thomas L, David B Fairlie, and Cosmas K Zachos (2013). *A concise treatise on quantum mechanics in phase space*. World Scientific Publishing Company (cit. on p. 121).

Kirchmair, Gerhard, Brian Vlastakis, Zaki Leghtas, Simon E. Nigg, Hanhee Paik, Eran Ginossar, Mazyar Mirrahimi, Luigi Frunzio, S. M. Girvin, and R. J. Schoelkopf (Mar. 14, 2013). "Observation of quantum state collapse and revival due to the single-photon Kerr effect". In: *Nature* 495.7440, pp. 205–209 (cit. on p. 157).

Vlastakis, Brian, Gerhard Kirchmair, Zaki Leghtas, Simon E Nigg, Luigi Frunzio, Steven M Girvin, Mazyar Mirrahimi, Michel H Devoret, and Robert J Schoelkopf (2013). "Deterministically encoding quantum information using 100-photon Schrödinger cat states". In: *Science* 342.6158, pp. 607–610 (cit. on p. 121).

Barends, R, J Kelly, A Veitia, A Megrant, AG Fowler, B Campbell, Y Chen, Z Chen, B Chiaro, A Dunsworth, et al. (2014). "Rolling quantum dice with a superconducting qubit". In: *Physical Review A* 90.3, p. 030303 (cit. on p. 111).

Deb, Manan and Dibakar Sen (2014). "Design of double toggle switching mechanisms". In: *Mechanism and Machine Theory* 71, pp. 163–190 (cit. on p. 27).

Flurin, Emmanuel (2014). "The Josephson mixer : a swiss army knife for microwave quantum optics". Ph.D. thesis. ENS Paris (cit. on p. 145).

Girvin, Steven M (2014). *Circuit QED: superconducting qubits coupled to microwave photons* (cit. on p. 81).

Locatelli, Nicolas, Vincent Cros, and Julie Grollier (2014). "Spin-torque building blocks". In: *Nature materials* 13.1, pp. 11–20 (cit. on p. 23).

Mirrahimi, Mazyar, Zaki Leghtas, Victor V. Albert, Steven Touzard, Robert J. Schoelkopf, Liang Jiang, and Michel H. Devoret (2014). "Dynami-

cally protected cat-qubits: A new paradigm for universal quantum computation". In: *New J. Phys.* 16.4, p. 045014 (cit. on pp. iv, v, 166).

Schäfer, F., I. Herrera, S. Cherukattil, C. Lovecchio, F.S. Cataliotti, F. Caruso, and A. Smerzi (Jan. 2014). "Experimental realization of quantum zeno dynamics". In: *Nature Communications* 5.1 (cit. on p. 162).

Signoles, Adrien, Adrien Facon, Dorian Grosso, Igor Dotsenko, Serge Haroche, Jean-Michel Raimond, Michel Brune, and Sébastien Gleyzes (Aug. 2014). "Confined quantum Zeno dynamics of a watched atomic arrow". In: *Nature Physics* 10.10, pp. 715–719 (cit. on p. 162).

Singh, V, SJ Bosman, BH Schneider, Ya M Blanter, A Castellanos-Gomez, and GA Steele (2014). "Optomechanical coupling between a multilayer graphene mechanical resonator and a superconducting microwave cavity". In: *Nature nanotechnology* 9.10, pp. 820–824 (cit. on pp. 72, 92).

Bretheau, L., P. Campagne-Ibarcq, E. Flurin, F. Mallet, and B. Huard (2015). "Quantum dynamics of an electromagnetic mode that cannot contain N photons". In: *Science* 348.6236, pp. 776–779. eprint: `https://www.science.org/doi/pdf/10.1126/science.1259345` (cit. on p. 162).

Campagne-Ibarcq, Philippe (2015). "Quantum backaction and feedback in superconducting circuits". Thèse de doctorat dirigée par Huard, Benjamin et Devoret, Michel Physique Paris, Ecole normale supérieure 2015. PhD thesis (cit. on p. 180).

Kienzler, Daniel, H-Y Lo, B Keitch, L De Clercq, F Leupold, F Lindenfelser, M Marinelli, Vlad Negnevitsky, and JP Home (2015). "Quantum harmonic oscillator state synthesis by reservoir engineering". In: *Science* 347.6217, pp. 53–56 (cit. on p. 129).

Leghtas, Zaki, Steven Touzard, Ioan M Pop, Angela Kou, Brian Vlastakis, Andrei Petrenko, Katrina M Sliwa, Anirudh Narla, Shyam Shankar, Michael J Hatridge, et al. (2015). "Confining the state of light to a quantum manifold by engineered two-photon loss". In: *Science* 347.6224, pp. 853–857 (cit. on pp. iv, 125, 171).

Micron (2015). DDR4 SDRAM MT40A2G4. Micron (cit. on p. 22).

Probst, S., F. B. Song, P. A. Bushev, A. V. Ustinov, and M. Weides (Feb. 2015). "Efficient and robust analysis of complex scattering data under noise in microwave resonators". In: *Review of Scientific Instruments* 86.2 (cit. on pp. 178, 180).

Albert, Victor V, Chi Shu, Stefan Krastanov, Chao Shen, Ren-Bao Liu, Zhen-Biao Yang, Robert J Schoelkopf, Mazyar Mirrahimi, Michel H Devoret, and Liang Jiang (2016). "Holonomic quantum control with continuous variable systems". In: *Physical review letters* 116.14, p. 140502 (cit. on pp. 142, 155–157).

Azouit, Remi, Alain Sarlette, and Pierre Rouchon (2016). *Adiabatic elimination for open quantum systems with effective Lindblad master equations*. arXiv: 1603.04630 [quant-ph] (cit. on pp. 69, 170).

Chen, Lily, Lily Chen, Stephen Jordan, Yi-Kai Liu, Dustin Moody, Rene Peralta, Ray A Perlner, and Daniel Smith-Tone (2016). *Report on post-quantum cryptography*. Vol. 12. US Department of Commerce, National Institute of Standards and Technology … (cit. on p. 16).

Raimond, Jean Michel (2016). "Atoms and photons". In: *M2 ICFP* (cit. on p. 109).

Bernstein, Daniel J and Tanja Lange (2017). "Post-quantum cryptography". In: *Nature* 549.7671, pp. 188–194 (cit. on p. 16).

Burnett, Jonathan, James Sagar, Oscar W Kennedy, Paul A Warburton, and Jonathan C Fenton (2017). "Low-loss superconducting nanowire circuits using a neon focused ion beam". In: *Physical Review Applied* 8.1, p. 014039 (cit. on p. 95).

Torrejon, Jacob, Mathieu Riou, Flavio Abreu Araujo, Sumito Tsunegi, Guru Khalsa, Damien Querlioz, Paolo Bortolotti, Vincent Cros, Kay Yakushiji, Akio Fukushima, et al. (2017). "Neuromorphic computing with nanoscale spintronic oscillators". In: *Nature* 547.7664, pp. 428–431 (cit. on p. 23).

Walter, Theodore, Philipp Kurpiers, Simone Gasparinetti, Paul Magnard, Anton Potočnik, Yves Salathé, Marek Pechal, Mintu Mondal, Markus Oppliger, Christopher Eichler, et al. (2017). "Rapid high-fidelity single-shot dispersive readout of superconducting qubits". In: *Physical Review Applied* 7.5, p. 054020 (cit. on p. 107).

Abobeih, Mohamed H, Julia Cramer, Michiel A Bakker, Norbert Kalb, Matthew Markham, Daniel J Twitchen, and Tim H Taminiau (2018). "One-second coherence for a single electron spin coupled to a multi-qubit nuclear-spin environment". In: *Nature communications* 9.1, p. 2552 (cit. on p. 38).

Minganti, Fabrizio (Oct. 2018). "Out-of-Equilibrium Phase Transitions in Nonlinear Optical Systems". Theses. Université Sorbonne Paris Cité (cit. on pp. 53, 54).

Tilloy, Antoine (2018). "Exact signal correlators in continuous quantum measurements". In: *Physical Review A* 98.1, p. 010104 (cit. on p. 135).

Touzard, S., A. Grimm, Z. Leghtas, S. O. Mundhada, P. Reinhold, C. Axline, M. Reagor, K. Chou, J. Blumoff, K. M. Sliwa, S. Shankar, L. Frunzio, R. J. Schoelkopf, M. Mirrahimi, and M. H. Devoret (2018). "Coherent Oscillations inside a Quantum Manifold Stabilized by Dissipation". In: *Phys. Rev. X* 8 (2), p. 021005 (cit. on pp. vii, 142, 157, 162, 166).

Tuckett, David K, Stephen D Bartlett, and Steven T Flammia (2018). "Ultrahigh error threshold for surface codes with biased noise". In: *Physical review letters* 120.5, p. 050505 (cit. on pp. iv, 39).

Arute, Frank, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando GSL Brandao, David A Buell, et al. (2019). "Quantum supremacy using a programmable superconducting processor". In: *Nature* 574.7779, pp. 505–510 (cit. on p. 106).

Cao, Yudong, Jonathan Romero, Jonathan P Olson, Matthias Degroote, Peter D Johnson, Mária Kieferová, Ian D Kivlichan, Tim Menke, Borja Peropadre, Nicolas PD Sawaya, et al. (2019). "Quantum chemistry in the age of quantum computing". In: *Chemical reviews* 119.19, pp. 10856–10915 (cit. on p. 17).

Flühmann, Christa, Thanh Long Nguyen, Matteo Marinelli, Vlad Negnevitsky, Karan Mehta, and JP Home (2019). "Encoding a qubit in a trapped-ion mechanical oscillator". In: *Nature* 566.7745, pp. 513–517 (cit. on p. 47).

Guillaud, Jérémie and Mazyar Mirrahimi (Dec. 2019). "Repetition Cat Qubits for Fault-Tolerant Quantum Computation". In: *Physical Review X* 9.4 (cit. on pp. vii, 39, 40, 142, 163, 168).

Hu, Ling, Yuwei Ma, Weizhou Cai, Xianghao Mu, Yuan Xu, Weiting Wang, Yukai Wu, Haiyan Wang, YP Song, C-L Zou, et al. (2019). "Quantum error correction and universal gate set operation on a binomial bosonic logical qubit". In: *Nature Physics* 15.5, pp. 503–508 (cit. on p. 23).

Touzard, S., A. Kou, N. E. Frattini, V. V. Sivak, S. Puri, A. Grimm, L. Frunzio, S. Shankar, and M. H. Devoret (2019). "Gated Conditional Displacement Readout of Superconducting Qubits". In: *Phys. Rev. Lett.* 122 (8), p. 080502 (cit. on p. 146).

You, Xinyuan, James A Sauls, and Jens Koch (2019). "Circuit quantization in the presence of time-dependent external flux". In: *Physical Review B* 99.17, p. 174512 (cit. on p. 88).

Bach, Francis (2020). "Machine learning-Master ICFP 2019-2020 Statistical Learning Theory". In: (cit. on p. 182).

Grimm, A., N. E. Frattini, S. Puri, S. O. Mundhada, S. Touzard, M. Mirrahimi, S. M. Girvin, S. Shankar, and M. H. Devoret (Aug. 2020). "Stabilization and operation of a Kerr-cat qubit". In: *Nature* 584.7820, pp. 205–209 (cit. on pp. vi, 48, 157).

Henriet, Loïc, Lucas Beguin, Adrien Signoles, Thierry Lahaye, Antoine Browaeys, Georges-Olivier Reymond, and Christophe Jurczak (2020).

"Quantum computing with neutral atoms". In: *Quantum* 4, p. 327 (cit. on p. 80).

Kudra, M, J Biznárová, A Fadavi Roudsari, JJ Burnett, D Niepce, S Gasparinetti, B Wickman, and P Delsing (2020). "High quality three-dimensional aluminum microwave cavities". In: *Applied Physics Letters* 117.7 (cit. on pp. 93, 96).

Lei, Chan U, Lev Krayzman, Suhas Ganjam, Luigi Frunzio, and Robert J Schoelkopf (2020). "High coherence superconducting microwave cavities with indium bump bonding". In: *Applied Physics Letters* 116.15 (cit. on p. 97).

Lescanne, Raphaël (Feb. 2020a). "Engineering multi-photon dissipation in superconducting circuits for quantum error correction". Theses. Université Paris sciences et lettres (cit. on pp. v, 73).

Lescanne, Raphaël, Marius Villiers, Théau Peronnin, Alain Sarlette, Matthieu Delbecq, Benjamin Huard, Takis Kontos, Mazyar Mirrahimi, and Zaki Leghtas (Mar. 2020b). "Exponential suppression of bit-flips in a qubit encoded in an oscillator". In: *Nature Physics* 16.5, pp. 509–513 (cit. on pp. vi, 86, 128).

Puri, Shruti, Lucas St-Jean, Jonathan A. Gross, Alexander Grimm, Nicholas E. Frattini, Pavithran S. Iyer, Anirudh Krishna, Steven Touzard, Liang Jiang, Alexandre Blais, Steven T. Flammia, and S. M. Girvin (Aug. 2020). "Bias-preserving gates with stabilized cat qubits". In: *Science Advances* 6.34 (cit. on pp. iv, vii, 39, 142, 168).

Sandlin, Destin (2020). *The Computer that Controlled the Saturn V (Behind the Scenes ft Linus Tech Tips) - Smarter Every Day*. Youtube. URL: https://www.youtube.com/watch?v=6mMK6iSZsAs (cit. on p. 21).

Darmawan, Andrew S., Benjamin J. Brown, Arne L. Grimsmo, David K. Tuckett, and Shruti Puri (2021). "Practical Quantum Error Correction with the XZZX Code and Kerr-Cat Qubits". In: *PRX Quantum* 2 (3), p. 030345 (cit. on pp. iv, vii, 142, 168).

Grimsmo, Arne L and Shruti Puri (2021). "Quantum error correction with the Gottesman-Kitaev-Preskill code". In: *PRX Quantum* 2.2, p. 020101 (cit. on p. 47).

Place, Alexander P. M., 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 (Mar. 2021). "New material platform for supercon-

ducting transmon qubits with coherence times exceeding 0.3 milliseconds". In: *Nature Communications* 12.1 (cit. on p. 168).

Robert, Anton, Panagiotis Kl Barkoutsos, Stefan Woerner, and Ivano Tavernelli (2021). "Resource-efficient quantum algorithm for protein folding". In: *npj Quantum Information* 7.1, p. 38 (cit. on p. 17).

Scholl, Pascal, Michael Schuler, Hannah J Williams, Alexander A Eberharter, Daniel Barredo, Kai-Niklas Schymik, Vincent Lienhard, Louis-Paul Henry, Thomas C Lang, Thierry Lahaye, et al. (2021). "Quantum simulation of 2D antiferromagnets with hundreds of Rydberg atoms". In: *Nature* 595.7866, pp. 233–238 (cit. on p. 30).

Aiello, Gianluca, Mathieu Féchant, Alexis Morvan, Julien Basset, Marco Aprili, Julien Gabelli, and Jérôme Estève (Nov. 2022). "Quantum bath engineering of a high impedance microwave mode through quasiparticle tunneling". In: *Nature Communications* 13.1 (cit. on p. 168).

Beverland, Michael E., Prakash Murali, Matthias Troyer, Krysta M. Svore, Torsten Hoefler, Vadym Kliuchnikov, Guang Hao Low, Mathias Soeken, Aarthi Sundaram, and Alexander Vaschillo (2022). *Assessing requirements to scale to practical quantum advantage*. arXiv: 2211.07629 [quant-ph] (cit. on p. 37).

Burgelman, Michiel, Pierre Rouchon, Alain Sarlette, and Mazyar Mirrahimi (2022). "Structurally Stable Subharmonic Regime of a Driven Quantum Josephson Circuit". In: *Phys. Rev. Appl.* 18 (6), p. 064044 (cit. on p. 132).

Chamberland, Christopher, Kyungjoo Noh, Patricio Arrangoiz-Arriola, Earl T Campbell, Connor T Hann, Joseph Iverson, Harald Putterman, Thomas C Bohdanowicz, Steven T Flammia, Andrew Keller, et al. (2022a). "Building a fault-tolerant quantum computer using concatenated cat codes". In: *PRX Quantum* 3.1, p. 010329 (cit. on p. 92).

Chamberland, Christopher, Kyungjoo Noh, Patricio Arrangoiz-Arriola, Earl T. Campbell, Connor T. Hann, Joseph Iverson, Harald Putterman, Thomas C. Bohdanowicz, Steven T. Flammia, Andrew Keller, Gil Refael, John Preskill, Liang Jiang, Amir H. Safavi-Naeini, Oskar Painter, and Fernando G.S.L. Brandão (Feb. 2022b). "Building a Fault-Tolerant Quantum Computer Using Concatenated Cat Codes". In: *PRX Quantum* 3.1 (cit. on pp. iv, vii, 142, 168).

Daley, Andrew J, Immanuel Bloch, Christian Kokail, Stuart Flannigan, Natalie Pearson, Matthias Troyer, and Peter Zoller (2022). "Practical quantum advantage in quantum simulation". In: *Nature* 607.7920, pp. 667–676 (cit. on p. 40).

Eickbusch, Alec, Volodymyr Sivak, Andy Z. Ding, Salvatore S. Elder, Shantanu R. Jha, Jayameenakshi Venkatraman, Baptiste Royer, S. M. Girvin, Robert J. Schoelkopf, and Michel H. Devoret (Oct. 2022). "Fast universal control of an oscillator with weak dispersive coupling to a qubit". In: *Nature Physics* 18.12, pp. 1464–1469 (cit. on p. 168).

Frattini, Nicholas E., Rodrigo G. Cortiñas, Jayameenakshi Venkatraman, Xu Xiao, Qile Su, Chan U Lei, Benjamin J. Chapman, Vidul R. Joshi, S. M. Girvin, Robert J. Schoelkopf, Shruti Puri, and Michel H. Devoret (2022). *The squeezed Kerr oscillator: spectral kissing and phase-flip robustness*. eprint: arXiv:2209.03934 (cit. on pp. vi, 48).

Gautier, Ronan, Alain Sarlette, and Mazyar Mirrahimi (2022). "Combined Dissipative and Hamiltonian Confinement of Cat Qubits". In: *PRX Quantum* 3 (2), p. 020339 (cit. on p. 168).

McEwen, Matt, Lara Faoro, Kunal Arya, Andrew Dunsworth, Trent Huang, Seon Kim, Brian Burkett, Austin Fowler, Frank Arute, Joseph C Bardin, et al. (2022). "Resolving catastrophic error bursts from cosmic rays in large arrays of superconducting qubits". In: *Nature Physics* 18.1, pp. 107–111 (cit. on p. 138).

Minganti, Fabrizio and Dolf Huybrechts (2022). "Arnoldi-Lindblad time evolution: Faster-than-the-clock algorithm for the spectrum of time-independent and Floquet open quantum systems". In: *Quantum* 6, p. 649 (cit. on p. 65).

Mylnikov, Valentin Yu, Sergey O Potashin, Grigorii S Sokolovskii, and Nikita S Averkiev (2022). "Dissipative phase transition in systems with two-photon drive and nonlinear dissipation near the critical point". In: *Nanomaterials* 12.15, p. 2543 (cit. on pp. 53, 54).

Paquette, Alexandre, Joël Griesmar, Gabriel Lavoie, Romain Albert, Florian Blanchet, Alexander Grimm, Ulrich Martel, and Max Hofheinz (2022). "Absorptive filters for quantum circuits: Efficient fabrication and cryogenic power handling". In: *Applied Physics Letters* 121.12 (cit. on p. 111).

Wang, Chenlu, Xuegang Li, Huikai Xu, Zhiyuan Li, Junhua Wang, Zhen Yang, Zhenyu Mi, Xuehui Liang, Tang Su, Chuhong Yang, et al. (2022a). "Towards practical quantum computers: Transmon qubit with a lifetime approaching 0.5 milliseconds". In: *npj Quantum Information* 8.1, p. 3 (cit. on pp. 106, 111).

Wang, Chenlu, 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 (Jan. 2022b). "Towards practical quantum computers: transmon qubit with a lifetime approaching 0.5 milliseconds". In: *npj Quantum Information* 8.1 (cit. on p. 168).

Beaulieu, Guillaume, Fabrizio Minganti, Simone Frasca, Vincenzo Savona, Simone Felicetti, Roberto Di Candia, and Pasquale Scarlino (2023). "Observation of first-and second-order dissipative phase transitions in a two-photon driven Kerr resonator". In: *arXiv preprint arXiv:2310.13636* (cit. on p. 54).

Berdou, Camille, Anil Murani, Ulysse Reglade, William C Smith, Marius Villiers, José Palomo, Michael Rosticher, A Denis, Pascal Morfin, Matthieu Delbecq, et al. (2023). "One hundred second bit-flip time in a two-photon dissipative oscillator". In: *PRX Quantum* 4.2, p. 020350 (cit. on pp. vi, vii, 56, 131, 132, 136, 142, 170).

Bluvstein, Dolev, Simon J. Evered, Alexandra A. Geim, Sophie H. Li, Hengyun Zhou, Tom Manovitz, Sepehr Ebadi, Madelyn Cain, Marcin Kalinowski, Dominik Hangleiter, J. Pablo Bonilla Ataides, Nishad Maskara, Iris Cong, Xun Gao, Pedro Sales Rodriguez, Thomas Karolyshyn, Giulia Semeghini, Michael J. Gullans, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin (Dec. 2023). "Logical quantum processor based on reconfigurable atom arrays". In: *Nature* 626.7997, pp. 58–65 (cit. on p. 80).

Gautier, Ronan, Mazyar Mirrahimi, and Alain Sarlette (2023). *Designing High-Fidelity Gates for Dissipative Cat Qubits*. arXiv: arXiv:2303.00760 [quant-ph] (cit. on pp. 165, 168).

Gouzien, Élie, Diego Ruiz, Francois-Marie Le Régent, Jérémie Guillaud, and Nicolas Sangouard (2023). "Performance analysis of a repetition cat code architecture: Computing 256-bit elliptic curve logarithm in 9 hours with 126 133 cat qubits". In: *Physical Review Letters* 131.4, p. 040602 (cit. on p. 40).

Guillaud, Jérémie, Joachim Cohen, and Mazyar Mirrahimi (2023). "Quantum computation with cat qubits". In: *SciPost Physics Lecture Notes*, p. 072 (cit. on pp. 58, 69, 72).

Kono, Shingo, Jiahe Pan, Mahdi Chegnizadeh, Xuxin Wang, Amir Youssefi, Marco Scigliuzzo, and Tobias J. Kippenberg (2023). *Mechanically Induced Correlated Errors on Superconducting Qubits with Relaxation Times Exceeding 0.4 Milliseconds*. eprint: arXiv:2305.02591 (cit. on p. 168).

Lei, Chan U, Suhas Ganjam, Lev Krayzman, Archan Banerjee, Kim Kisslinger, Sooyeon Hwang, Luigi Frunzio, and Robert J Schoelkopf (2023). "Mi-

crowave loss characterization using multi-mode superconducting resonators". In: *arXiv preprint arXiv:2305.01872* (cit. on p. 93).

Lu, Yao, Aniket Maiti, John WO Garmon, Suhas Ganjam, Yaxing Zhang, Jahan Claes, Luigi Frunzio, Steven M Girvin, and Robert J Schoelkopf (2023). "High-fidelity parametric beamsplitting with a parity-protected converter". In: *nature communications* 14.1, p. 5767 (cit. on p. 140).

Marquet, A, A Essig, J Cohen, N Cottet, A Murani, E Abertinale, S Dupouy, A Bienfait, T Peronnin, S Jezouin, R Lescanne, and B Huard (2023). *Cat qubit with 0.3 s bit-flip time owing to passive two photon dissipation* (cit. on p. 168).

"Suppressing quantum errors by scaling a surface code logical qubit" (2023). In: *Nature* 614.7949, pp. 676–681 (cit. on pp. iv, 37).

Ulysse Réglade and Adrien Bocquet, Ronan Gautier, Antoine Marquet, Emanuele Albertinale, Natalia Pankratova, Mattis Hallén, Felix Rautschke, Lev-Arcady Sellem, Pierre Rouchon, Alain Sarlette, Mazyar Mirrahimi, Philippe Campagne-Ibarcq, Raphaël Lescanne, Sébastien Jezouin, and Zaki Leghtas (2023). *Quantum control of a cat-qubit with bit-flip times exceeding ten seconds*. arXiv: 2307.06617 [quant-ph] (cit. on p. 141).

Dubovitskii, Kirill, Denis M. Basko, Julia S. Meyer, and Manuel Houzet (2024). *Theory of quasiparticle-induced errors in driven-dissipative Schrödinger cat qubits*. arXiv: 2403.15310 [cond-mat.mes-hall] (cit. on p. 138).

Ruiz, Diego, Jérémie Guillaud, Anthony Leverrier, Mazyar Mirrahimi, and Christophe Vuillot (2024a). "LDPC-cat codes for low-overhead quantum computing in 2D". In: *arXiv preprint arXiv:2401.09541* (cit. on p. 37).

Ruiz, Diego, Jérémie Guillaud, Anthony Leverrier, Mazyar Mirrahimi, and Christophe Vuillot (2024b). *LDPC-cat codes for low-overhead quantum computing in 2D*. arXiv: 2401.09541 [quant-ph] (cit. on p. 40).

Ken Shirrif's blog (n.d.). http://www.righto.com/2020/03/the-core-memory-inside-saturn-v-rockets.html. Accessed: 2024-04-22 (cit. on p. 21).