March 4, 2024

Notes for Phys 250: Quantum Theory of Measurement

Daniel Carney

Physics Division, Lawrence Berkeley National Laboratory
Berkeley, CA

Email: carney@lbl.gov

Abstract

These are collected notes for a course taught at UC Berkeley in Spring 2024, on the quantum theory of measurement.

Contents

0 Preliminaries **2**

1 Fundamentals: what is a measurement? **4**
1.1 The standard lore and its limitations 4
1.2 The modern lore: POVMs, etc. 10
1.3 Local and non-local measurements 13

2 Discrete-variable ("qubit") measurements **17**
2.1 Stern-Gerlach experiment 17
2.2 Hyperfine atomic qubit 22

3 Continuous-variable measurements **30**
3.1 Position measurements 30
3.2 Heisenberg's microscope and measurement back-action 32
3.3 Number and quadrature bases; homodyne measurements 34

4 Tomography of quantum states **43**
4.1 Examples on a single degree of freedom 44
4.1.1 Pauli tomography on one qubit 44
4.1.2 Wigner representations; homodyne tomography of an oscillator 48
4.2 Axiom of local tomography 53
4.2.1 Example: Pauli tomography of n qubits 54
4.3 Informationally complete measurements 55
4.3.1 Example: Random unitaries and SIC-POVMs 55

5 Time evolution and its measurement **56**
5.1 Quantum channels, operations, and instruments 56
5.2 Decoherence and dephasing 61
5.3 Local time evolution 67
5.3.1 The Lindblad equation 68
5.3.2 Markovian and non-Markovian noise 71
5.3.3 Quantum noise: spin-boson model 75
5.4 Thermalization: optical Bloch equations 79

A Quantum mechanics of the electromagnetic field **82**
A.1 Quantization of the field 82

0. Preliminaries

In this course we use "natural" units $\hbar = c = 1$. This means that all dimensionful quantities have units of either mass to some power or length to the inverse power, $M^n = 1/L^n$. In particular mass, energy, momentum, and frequency all have the same units. The reason you should get used to these units is not because they make formulas shorter (which is true) but because they make dimensional analysis extremely simple. A very useful tool for unit conversions is https://gutcalc.com/, which is also available as a phone app; it will let you input in an arbitrary mix of SI and natural units and output in whatever units you want.

I feel obliged to mention some cultural background before moving into the technical material. My training is in formal high energy theory, not quantum measurements, although that is now what I spend most of my time working on. These notes are my attempt to systematize a collection of things I've learned piecewise over the past ten years to solve certain problems I'm interested in. What this means for you is that

while I promise I am attempting to present the standard story, in various places I may say things that would be considered heterodox, or use language in a non-standard way. Just keep an eye out and feel free to ask questions at all times.

Finally, a brief comment on a key difference between the presentation here and the books available when I started to compile these notes. The main difficulty I had with the available materials is that very few of them really get into the details of physics calculations underlying the measurement process. This is a common issue in modern quantum information texts: often students are just handed equations in the language of oscillators and qubits, and at best some heuristic argument for the form of the equations, usually with arbitrary coefficients. This gives the mistaken impression that these kinds of things can't be derived from microscopic models. I have endeavored, especially in the early materials, to provide concrete examples that can be analyzed start-to-finish from detailed dynamical models. The particular goal is to let students see how the simple descriptions can be obtained systematically from exact (or at least, more exact) starting points.

Many people are owed thanks for helpful discussions while preparing these notes and learning the core of this material over the years, especially Jess Riedel and Jake Taylor. I would also like to thank the students at UC Berkeley for their patience and feedback, and particularly Jacob Beckey, who provided great advice and resources for many parts of this material.

	Pure states	General states
State	$\|\psi\rangle \in \mathcal{H}$	$\rho = \sum_i p_i \|i\rangle \langle i\|$
Measurement outcomes	Eigenvalues a of $A = A^\dagger$	same
Probabilities	$P(a) = \| \langle a\|\psi\rangle \|^2$	$P(a) = \mathrm{tr} \|a\rangle \langle a\| \rho$
State update	$\|\psi\rangle \to \|a\rangle$	$\rho \to \|a\rangle \langle a\|$

Table 1: The old school rules of quantum measurement.

1. Fundamentals: what is a measurement?

The majority of this course is going to be based on rather established, hard-nosed stuff like noise budgets, measurement efficiencies, and so forth, where like most of quantum mechanics we will largely follow a "shut up and calculate" approach. This is justified because the calculations all correspond to reality in a clear way.

In this first section, however, we are going to take a dive into one of the thorniest issues in modern physics: the actual definition of a measurement. The simple truth is that as of now we *do not have a rigorous definition of a measurement in quantum mechanics*! In quantum mechanics books this is usually mentioned and then ignored. In the handful of measurement theory books that exist to my knowledge, I was surprised to find that still not much more detail is usually given. Having now tried to produce a few lectures on this, I understand why: it's a brutal problem to treat in a way where you both reproduce the right heuristic rules, and also do not lie. What follows is my attempt.

1.1. The standard lore and its limitations

Here is what you probably learned in your quantum mechanics courses. The state of a system is a vector $|\psi\rangle \in \mathcal{H}$ in a Hilbert space. "Observables" are represented by Hermitian operators $A = A^\dagger$ on $\mathcal{H}$. The spectral theorem guarantees that these can be diagonalized, $A = \sum_a a \, |a\rangle \langle a|$. When we do a measurement of A, we obtain a one of the outcomes a, with probability distribution $P(a) = | \langle a|\psi\rangle |^2$. After the measurement, the state collapses to the corresponding eigenvector $|\psi\rangle \to |a\rangle$. This is summarized in Table 1.

You should notice that in the above, at no point did we actually specify how the

measurement is supposed to be enacted; this is just a list of rules about outcomes. This is a part of the so-called measurement problem. The reality is that we have a clear sense of certain circumstances where the above rules work and we know what we mean, but we also have a much more general set of circumstances we would also normally consider measurements, where these rules do not apply. We will get to these shortly. First, let's talk about what actually doing the measurement looks like in an example where the above story does apply: Schrödinger's cat.

The setup is famous: we have a box, within which there is an unstable uranium isotope, a vial of poison which breaks open when the uranium decays, and a cat. We are going to think of this whole setup as enacting a measurement *of the state of the uranium atom.* The basic idea is that the microscopic state of the atom gets "amplified" by coupling to the cat, which we can then measure according to the standard story given above.

First, let's review how the scene unfolds. Let $|e\rangle, |g\rangle$ denote the excited (unstable/parent) and ground (stable/daughter) state of the uranium atom. The initial state total state $|\Psi\rangle$ of the joint system evolves unitarily,

$$|\Psi(0)\rangle |e\rangle \otimes |\text{alive}\rangle \to |\psi(t)\rangle = A_e(t) |e\rangle \otimes |\text{alive}\rangle + A_g(t) |g\rangle \otimes |\text{dead}\rangle , \tag{1.1}$$

where $A_e(t) \sim e^{-t/2t_{1/2}}$ is the amplitude for the decay in terms of the half-life $t_{1/2}$. The cat's Hilbert space also includes the vial and poison. Now, at this stage we take for granted that we can "measure" the state of the cat, vial, and poison by simply opening the box and looking at it. There seems no doubt of this in real life. Moreover we will get one or the other outcome, dead or alive. One can write down a Hermitian operator

$$A = a_{\text{alive}} |\text{alive}\rangle \langle\text{alive}| + a_{\text{dead}} |\text{dead}\rangle \langle\text{dead}| = a_{\text{alive}}\Pi_{\text{alive}} + a_{\text{dead}}\Pi_{\text{dead}} \tag{1.2}$$

which is just the sum of two projectors $\Pi_{\text{alive,dead}}$ with some eigenvalues $a_{\text{alive,dead}}$, and claim we are measuring this observable. The probability distribution of the outcomes is given by the Born rule quoted above. To compute this we need

$$\langle\text{alive}|\Psi(t)\rangle = A_e(t) |e\rangle , \quad \langle\text{dead}|\Psi(t)\rangle = A_g(t) |g\rangle , \tag{1.3}$$

from which we obtain

$$P(\text{alive}) = |\langle\text{alive}|\Psi(t)\rangle|^2 = |A_e(t)|^2, \quad P(\text{dead}) = |\langle\text{alive}|\Psi(t)\rangle|^2 = |A_g(t)|^2. \tag{1.4}$$

The most subtle aspect is the wavefunction collapse rule. The formal way of writing is that given the outcome of the cat, we should project the whole system appropriately:

$|\psi\rangle \to \Pi_{\text{outcome}} |\psi\rangle / \sqrt{P(\text{outcome})}$, where the denominator normalizes the final state. Here, this means that we get

$$|\Psi(t)\rangle \to \begin{cases} |e\rangle \otimes |\text{alive}\rangle & \text{given outcome ``alive''} \\ |g\rangle \otimes |\text{dead}\rangle & \text{given outcome ``dead''} \end{cases}. \tag{1.5}$$

It should be emphasized that the use of kets $|e, g\rangle$ and $|\text{alive}, \text{dead}\rangle$ is a massive oversimplification of reality. There is not a single decay state but rather a whole spectrum of them—the various radiation products can be ejected with various energies and directions, for example—and similarly there are two huge subspaces of the cat+vial+poison Hilbert space corresponding to the two outcomes. This is more or less accurately captured in the general description using projectors like $\Pi_{\text{alive,dead}}$, where these now project not onto a certain specific vector but the high-dimensional alive/dead subspaces. We will return to this point later.

So in what sense did we measure the quantum state of the atom itself? From Eq. (1.5), we know that with $P(e) = |A_e|^2$ we found the atom in the unstable excited state and the cat alive, and similarly for ground state and dead cat. The cat serves as a measurement "record": classical information, stable under reasonable perturbations to the environment, which tells us the outcome. What about the atom state by itself? To calculate the state of just the atom, we use the usual partial trace rules

$$\rho_{\text{atom}} = \text{tr}_{\text{cat}} |\Psi\rangle \langle\Psi| = \begin{cases} |e\rangle \langle e| & \text{given outcome ``alive''} \\ |g\rangle \langle g| & \text{given outcome ``dead''} \end{cases}. \tag{1.6}$$

In other words, the atom state is similarly projected onto either $|e\rangle$ or $|g\rangle$. So, this whole process acts to make a *projective measurement* on the atom, obeying the textbook rules above: we have an operator $A_{\text{atom}} = a_e |e\rangle \langle e| + a_g |g\rangle \langle g|$, and then when we measure A_{atom} we get the outcomes e, g with the probabilities $P(e), P(g)$. This is the way that measurement of a quantum system is often presented: the measurement apparatus (the cat) is not mentioned explicitly and we just work with operators on the system of interest (the atom). It is taken for granted that some measurement system like the cat can be found. Much of this course will be dedicated to understanding how these cats actually work.

As an aside, notice that the nominal eigenvalues $a_{\text{alive,dead}}$ or $a_{e,g}$ never affect anything. This reflects a basic fact about measurement that is usually glossed over: the choice of a measurement is not so much about choosing some Hermitian operator A, but rather about choosing a *basis*. This may sound esoteric, but for example,

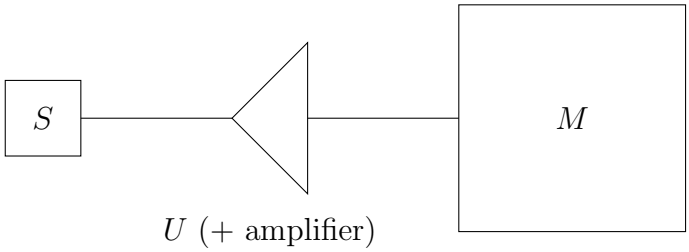

Figure 1: Abstract definition of a measurement on a system S. First, S is coupled to a larger measurement system M, and in particular there is some amplification of the S information. Then, a von Neumann (repeatable, projective) measurement is made on M, which then enacts a generalized measurement on S following the rules in Table 2.

notice that a measurement of either the position x or its cube x^3 obviously can't give any different information about where a particle is. The real point is just that you measure *in the position basis*.

The above example highlights the central elements of what is, at present, the closest we have to a definition of a measurement in quantum mechanics. We have a system S (e.g., the uranium atom), whose state we want to measure, and a measurement apparatus M (e.g., the poison and cat). Now we make a circular definition: to do a measurement on S, we first entangle S with M via some kind of dynamics, and then we measure M. This "measurement" on M is never really precisely defined; we just assume that M is large, and the entangling dynamics can "amplify" the information about S into a stable set of states of M, that we can fall back on a heuristic definition of measurement on M. Then, we can use the rules of partial traces to figure out what happened to the original system S. See Fig. 2.

The amplification step is critically important. As a toy model to illustrate what it means, consider a qubit $S = \text{span}\{|0\rangle, |1\rangle\}$ and some large number $n \gg 1$ of readout qubits $M = \text{span}\{|000\cdots\rangle, |100\cdots\rangle, \ldots\}$, with 2^n total states. The amplification step could look something like a dynamical, unitary evolution U that takes

$$|\Psi\rangle = (a|0\rangle + b|1\rangle) \otimes |000\cdots\rangle \to U|\Psi\rangle = a|0\rangle \otimes |000\cdots\rangle + b|1\rangle \otimes |111\cdots\rangle . \quad (1.7)$$

The idea is that we can safely assume that we can measure this large qubit system. For example if these are electron spins, the M states on the right-hand side will have a large magnetic moment and could be measured with a simple magnetometer. If we measure in the total-M-spin-z basis, the same logic used in the cat example shows that this will produce a projective spin-z measurement on the S qubit.

This amplifier is just a CNOT gate, where the qubit S is the control and the states $|\tilde{0}\rangle = |000\cdots\rangle$ and $|\tilde{1}\rangle = |111\cdots\rangle$ of the amplifier form a "logical subspace", in the language of error correction. Then the amplifier modeled here acts as as,

$$U = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} \tag{1.8}$$

where the matrix elements are in the $|0\rangle|\tilde{0}\rangle, |0\rangle|\tilde{1}\rangle, |1\rangle|\tilde{0}\rangle, |1\rangle|\tilde{1}\rangle$ basis; this is the usual CNOT gate in quantum computing. In this example, the amplification is "noiseless": the two S basis states are mapped perfectly onto an orthogonal pair of pure M states. In a real amplifier, the final M states will be more complicated, for example superpositions of many different bit strings, or perhaps the final states will not even be pure. In this case what a measurement should do is actually pick out some pair of subspaces of the M Hilbert space, where all the states in the subspaces still accurately correspond to the right system states. This is a simple form of error correction. We will discuss this kind of readout noise in great detail throughout this course.

Before moving on to the general formalism sketched above, let's motivate it with another example, one in which the measurement on the microscopic system does *not* obey the standard framework of Table 1: photodetection. Suppose we have a semiconductor or some kind of system that has a photoelectric effect; we can model this as a bunch of electrons, each of which has a Hamiltonian with a bound ground state $|g\rangle$ with energy $E = -\Delta$, and then some band of excited (conducting/free) states $|E\rangle$ with energies $E > 0$. The basic operating principle is that an incident photon of frequency $\omega \geq \Delta$ can be absorbed and excite one of the electrons to some conducting state $|g\rangle \to |E\rangle$, which is then amplified and detected. The amplification step can be done in various ways; for example one can imagine having a large voltage bias, so that once the electron is excited, its energy gets amplified to some large value, and is then smashed into a metal plate or something such that an easily measurable shower of particles is produced. The whole chain of systems including the electrons, bias, and final detector constitute the measurement apparatus M.

Now for the system S, suppose we prepare a mode of the electromagnetic field in some superposition of the vacuum $|0\rangle$ and a single photon $|1\rangle$. We will discuss the

details of this later, but for now let's model the interaction U as

$$\begin{aligned} U\left|0\right\rangle\left|g\right\rangle &= \left|0\right\rangle\left|g\right\rangle, \\ U\left|1\right\rangle\left|g\right\rangle &= \left|0\right\rangle\left|e\right\rangle. \end{aligned} \tag{1.9}$$

Here $|e\rangle$ is some particular excited state in the $\{|E\rangle\}_{E>0}$ subspace; for example, one where the gap $\omega_{eg} = \omega_{\text{light}}$ is exactly on resonance with the input light. The unitary in Eq. (1.9) is a very simplified picture, in which the photodetector is a perfect absorber (in optics language, it has "quantum efficiency" $\eta = 1$). In homework, you will explore what happens with a finite detection efficiency. Following the same logic as in the uranium and cat case, we will now assume that we can projectively measure the electron; we will leave the amplifier, final detector and so forth implicit, and just act with projectors $\Pi_e = |e\rangle\langle e|$ and $\Pi_g = |g\rangle\langle g|$. With an arbitrary initial state of the electromagnetic mode, the interaction produces the final joint state

$$\left|\Psi\right\rangle\left(a\left|0\right\rangle + b\left|1\right\rangle\right)\otimes\left|g\right\rangle \rightarrow U\left|\Psi\right\rangle = a\left|0\right\rangle\otimes\left|g\right\rangle + b\left|0\right\rangle\otimes\left|e\right\rangle. \tag{1.10}$$

Measuring the electron (plus amplifier, ...) we now obtain the outcome e, i.e., the detector "clicks" as we have absorbed a photon, with probability $|a|^2$. We hear no click since no photon was absorbed, with probability $|b|^2$. This is obviously the right answer, given the initial state.

What is the resulting final state of the photon field? Again using the partial trace as described above, we find that the state of the photon, conditioned on the different measurement outcomes, is

$$\rho_{\text{photon}} = \text{tr}_{\text{electron}}\left|\Psi\right\rangle\left\langle\Psi\right| = \begin{cases} \left|0\right\rangle\left\langle 0\right| & \text{given outcome "e}'' \\ \left|0\right\rangle\left\langle 0\right| & \text{given outcome "g}'' \end{cases}. \tag{1.11}$$

Oh no! In either case, the resulting photon state is just the vacuum. That this was going to be the answer is physically obvious: either there wasn't a photon in the first place, or there was and then it got absorbed. But this is naively at odds with the usual "wavefunction collapse" rule, which one might have said would project the photon into two different states. Indeed, if we thought of this as "measuring the number operator", the standard update postulate would have said that if we got 1 that the resulting photon state was $|1\rangle$. On the other hand, the measurement outcomes *do* correspond to perfect ("projective") information about the photon.

What this example illustrates is that there are operations we can do, like absorptive photodetection, which any sane person would call a measurement, but which

	Pure states	General states
State	$\|\psi\rangle \in \mathcal{H}$	$\rho = \sum_i p_i \|i\rangle \langle i\|$
Measurement outcomes	Labels a of POVM $E_a = K_a^\dagger K_a$	same
Probabilities	$P(a) = \langle\psi\|E_a\|\psi\rangle$	$P(a) = \operatorname{tr} E_a\rho$
State update	$\|\psi\rangle \to \frac{K_a\|\psi\rangle}{\sqrt{P(a)}}$	$\rho \to \frac{K_a\rho K_a^\dagger}{P(a)}$

Table 2: The modern rules of quantum measurement.

do not obey the simple rules of Table 1. In this case, the state update rule is not correct. More generally, we can modify some of the other naive measurement rules as well. In the next section, we turn to this generalized, modern picture of quantum measurements.

1.2. The modern lore: POVMs, etc.

Let's spell out the above measurement procedure in detail, which will lead us immediately into some examples of measurements more general than projective ones. Fix a system S in an arbitrary initial state ρ_S, and measurement apparatus M in some "register" state $|0\rangle$. We couple them through unitary dynamics:

$$\rho = \rho_S \otimes |0\rangle\langle 0| \to \rho' = U\left(\rho_S \otimes |0\rangle\langle 0|\right) U^\dagger \tag{1.12}$$

and then projectively measure the apparatus in some basis $|a\rangle$. We get outcome a with probability

$$\begin{aligned} P(a) &= \operatorname{tr} \Pi_a \rho' \\ &= \operatorname{tr}_S \langle a|U\left(\rho_S \otimes |0\rangle\langle 0|\right) U^\dagger|a\rangle \\ &= \operatorname{tr}_S K_a \rho_S K_a^\dagger, \end{aligned} \tag{1.13}$$

where we defined the *Kraus operators*

$$K_a := \langle a|U|0\rangle\,. \tag{1.14}$$

Notice that these are operators on the system Hilbert space $\mathcal{H}_S$; they are not necessarily Hermitian, but satisfy a completeness relation that is important to make sense

of the out probabilities:

$$\sum_a K_a^\dagger K_a = \sum_a \langle 0|U^\dagger|a\rangle \langle a|U|0\rangle = 1_S \tag{1.15}$$

Following the discussion in the previous section, we emphasize again that there is an amplification step, where the state of S is highly amplified into a much larger measurement apparatus M. See Fig. 1.1. However, we will often just model M as a simple quantum system and leave the amplification and large measurement apparatus implicit. This may sound esoteric in the abstract; we will do a number of examples shortly.

Thus, in the end, those are the rules: a measurement is defined by a set of "outcomes" $\{a\}$, each of which comes with a Kraus operator, and the experiment produces outcome a with probability given by Eq. (1.13). Projective measurements are a special case, when $K_a = K_a^2 = K_a^\dagger$ are projection operators and Hermitian. Note the physical meaning of the "outcome" a: it is some specific state $|a\rangle$ recorded on the measurement apparatus, for example $|\text{alive}\rangle$ in the Schrödinger's cat example.

Looking at the state of the joint system and measurement apparatus after the measurement is very instructive. Assuming that we get outcome a, the usual von Neumann rules say that we should project the M state down to $|a\rangle$, or more precisely we should act with the projection operator $1_S \otimes |a\rangle\langle a|$ on the joint system. Doing so, we find that the state is updated to

$$\begin{aligned} \rho' &\to \frac{1}{P(a)} \left(1_S \otimes |a\rangle\langle a|\right) \rho' \left(1_S \otimes |a\rangle\langle a|\right)^\dagger \\ &= \frac{1}{P(a)} |a\rangle\langle a| U \left(\rho_S \otimes |0\rangle\langle 0|\right) U^\dagger |a\rangle\langle a| \\ &= \frac{K_a \rho_S K_a^\dagger}{P(a)} \otimes |a\rangle\langle a| . \end{aligned} \tag{1.16}$$

The factor $P(a)$ is needed to ensure the final state is normalized correctly. The interpretation of the last line is very clear: given the outcome a, the measurement apparatus M is projected into the appropriate "pointer" state $|a\rangle$, and the system state is updated as

$$\rho_S \to \frac{K_a^\dagger \rho_S K_a}{P(a)}. \tag{1.17}$$

Notice that K_a, in general, is not a projection operator, which means that the updated system state is not necessarily a particular pure state. In other words, a general measurement does not completely "collapse" the system state.

As we will emphasize in the next sections with examples, the state update rule can be quite tricky. Even in simple scenarios one has to be careful that the projective measurement acting on the apparatus is described correctly in order to get the right Kraus operators. We will illustrate this in the next section with a simple example, the Stern-Gerlach experiment.

Sometimes this story is presented slightly differently. Notice that in the final expression for the outcome probability distribution, Eq. (1.13), we could choose a different set of Kraus operators

$$K'_a = UK_a \tag{1.18}$$

and still get the same set of probability outcomes, where U can be an arbitrary unitary on $\mathcal{H}_S$. This corresponds, physically, to the fact that we can extract the same system information using a variety of different apparatuses. A more invariant object that we can form is

$$E_a := K_a^\dagger K_a. \tag{1.19}$$

These make up what is called a Projective Operator-Valued Measure (POVM), which is a set of positive operators $E_a > 0$ on $\mathcal{H}_S$ satisfying $\sum_a E_a = 1_S$. The name is because for any set of outcomes $\mathcal{A} = \{a_1, a_2, \ldots\}$, we can form a positive operator $\mathcal{O}_\mathcal{A} = \sum_{a\in\mathcal{A}} E_a$, and use these to define a probability distribution

$$P(a \in \mathcal{A}) = \mathrm{tr}_S\, \mathcal{O}_\mathcal{A} \rho_S. \tag{1.20}$$

In other words, this produces a measure (in the sense of measure spaces) on the space of measurement outcomes. Obviously, this distribution is invariant under (1.18). Note that the special case of projective measurements, where the E_a are all projection operators, similarly define a measure; these are sometimes referred to as Projective Operator Measures (PVMs). Sometimes people talk about measurements being defined by a choice of POVM. The issue with this that it does not uniquely specify the update to the system state after the measurement: clearly Eq. (1.17) is not invariant under the choice of different Kraus operators.

These generalized rules are summarized in Table 2. As we have emphasized, these rules encapsulate just about anything one would like to think of as a measurement in quantum mechanics. Of course, they still rely on a kind of circular definition, where we take for granted some notion of projective measurement on the apparatus M with the usual Born rules. Figuring out if there is a more general definition which does not enforce this step ad hoc is the essence of the so-called "measurement problem",

hello

Figure 2: **Local vs non-local measurements.** Left: A local measurement on $\mathcal{H}_S = \mathcal{H}_{S_1} \otimes \mathcal{H}_{S_2} \otimes \cdots$. Each factor $\mathcal{H}_{S_i}$ is coupled to a single measurement apparatus M_i through a unitary U_i, so that the whole system is coupled to the whole measurement apparatus via $U = U_1 \otimes U_2 \otimes \cdots$. Each apparatus is then measured in some basis $|a_i\rangle_{M_i}$. Right: An example non-local measurement on a bipartite system $\mathcal{H}_S = \mathcal{H}_{S_1} \otimes \mathcal{H}_{S_2}$. Here there is only one measurement apparatus M, and the unitary U couples both $S_{1,2}$ to M.

and thinking about it too hard leads one deep into the weeds of things like the Many-Worlds Interpretation, QBism, and so forth. While there may be light at the end of that very dark tunnel, in this course we will instead follow the pragmatic route, shutting up and calculating based on these generalized measurement rules.

1.3. Local and non-local measurements

(Under construction: need to fix the definition here to allow for classical post-processing, as in Pauli tomography. –dc)

In quantum measurements, the distinction between local and non-local measurements is very important. The language is slightly different from the notion of locality used in other contexts, so here we briefly spell out the definitions. The use of local measurements and non-local measurements will play a crucial role in our study of state reconstruction ("tomography"), in Ch. 4.

Locality in quantum mechanics is always a statement about Hilbert space factors; it has no a priori connection to spacetime locality. Suppose we have a system with Hilbert space $\mathcal{H}$ and we have some way to decompose it into factors

$$\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2 \otimes \cdots . \tag{1.21}$$

For example, this could be a bunch of qubits, or it could be the spatial degree of freedom and spin degree of freedom of a single spin-1/2 particle $\mathcal{H} = \text{span}\,\{|\mathbf{x}\rangle\} \otimes \text{span}\,\{|\uparrow\rangle, |\downarrow\rangle\}$. In the first example, one would usually arrange the qubits at different locations in space, so this decomposition does reflect spacetime locality. In the second example, however, the two degrees of freedom literally live *at the same spacetime point*, so the decomposition has nothing to do with spacetime locality.

One often talks about "local interactions". This means that we have a Hamiltonian that couples, for example, sets of k factors. Here we are going to be concerned with

a different issue: local vs. non-local measurements. The formal definition of a local measurement is that it is described by POVMs that factor:

$$E_{\mathbf{a}} = E_{a_1} \otimes E_{a_2} \otimes \cdots \qquad \text{(local measurement)} \tag{1.22}$$

where each E_{a_i} acts on $\mathcal{H}_i$, the ith tensor factor, and $\mathbf{a} = (a_1, a_2, \ldots)$ is some vector of outcomes. To understand this definition, it is helpful to go back to our system-apparatus picture. Suppose that to each system site S_i with Hilbert space $\mathcal{H}_i$, we attach a measurement system M_i, and a (2-local) coupling U_i that couples S_i to M_i. The total coupling then factors $U = U_1 \otimes U_2 \otimes \cdots$. See Fig. 2. We can then find a basis $|a_i\rangle_i$ on each of these M_i, and consider the total Kraus operator for measuring all of these at once:

$$K_{\mathbf{a}} = \langle a_1 a_2 \cdots | U | 00\cdots \rangle_M = \prod_i \langle a_i | U_i | 0 \rangle_{M_i} = \bigotimes_i K_{a_i} \tag{1.23}$$

which in turn gives a POVM

$$E_{\mathbf{a}} = K_a^\dagger K_a = \left(\bigotimes_i K_{a_i}^\dagger \right) \left(\bigotimes_j K_{a_j} \right) = \bigotimes_i E_{a_i} \tag{1.24}$$

with $E_{a_i} = K_{a_i}^\dagger K_{a_i}$. This last equality follows from associativity of the tensor product over normal matrix multiplication: $(A \otimes B)(A' \otimes B') = AA' \otimes BB'$. This overly formal discussion encodes a very simple idea: a local measurement is just one that can be described by locally coupling a single measurement apparatus to each factor $\mathcal{H}_i$ of the system Hilbert space and measuring those apparati in parallel.

This can be compared with non-local measurements. These involve coupling multiple system factors to a single measurement apparatus; see Fig. 2 for an example. This produces a POVM which does not factor as in Eq. (1.22). As an instructive example, consider the case of two qubits $\mathcal{H}_S = \mathcal{H}_A \otimes \mathcal{H}_B$ being measured by a third qubit $\mathcal{H}_M$, in a protocol where we couple the qubits via a pair of CNOT gates [see Eq. (1.8)]

$$U_{SM} = U_{AM} \otimes U_{BM} = \mathrm{CNOT}_{AM} \otimes \mathrm{CNOT}_{BM}, \tag{1.25}$$

and then measure M in the computational basis $\{|a\rangle_M\} = \{|0\rangle_M, |1\rangle_M\}$. It is a nice

exercise to work out the Kraus operators:

$$K_0 = \langle 0|U_{SM}|0\rangle_M = |0\rangle\langle 0|_A \otimes |0\rangle\langle 0|_B + |1\rangle\langle 1|_A \otimes |1\rangle\langle 1|_B = \begin{pmatrix} 1 & & & \\ & 0 & & \\ & & 0 & \\ & & & 1 \end{pmatrix}$$
$$K_1 = \langle 1|U_{SM}|0\rangle_M = |0\rangle\langle 0|_A \otimes |1\rangle\langle 1|_B + |1\rangle\langle 1|_A \otimes |0\rangle\langle 0|_B = \begin{pmatrix} 0 & & & \\ & 1 & & \\ & & 1 & \\ & & & 0 \end{pmatrix} \quad (1.26)$$

where the matrices are in the $|00\rangle, |01\rangle, |10\rangle, |11\rangle$ basis on $\mathcal{H}_S = \mathcal{H}_A \otimes \mathcal{H}_B$. These give a simple POVM $E_a = K_a^\dagger K_a = K_a$ since the K_a just square back to themselves, and it is easy to see that the POVM is complete $E_0 + E_1 = 1_S$. It is, however, non-local: clearly the K_a are not products $K_0 \neq K_{A_0} \otimes K_{B_0}$, and similarly for the other Kraus operator and POVMs.

Further reading

- Peres, Chapter 12
- Preskill, Chapter 3
- Jordan and Siddiqui, Chapters 2-4

Problems

1. **A slightly more realistic photodetector model**, which includes non-unit absorption efficiency.

   (a) Define η as the fractional power/energy of incident light that is absorbed, i.e., $E_{\text{absorbed}} = \eta E_{\text{incident}}$. Following the photodetector model in the text above, let $\mathcal{H}_S = \text{span}\{|0\rangle, |1\rangle\}$ denote the space of either 0 or 1 incident photons, and $\mathcal{H}_M = \text{span}\{|g\rangle, |e\rangle\}$ denote the relevant electron ground and excited states. Write down a unitary U on $\mathcal{H}_S \otimes \mathcal{H}_M$ that describes the interaction, now including the possibility that an incident photon might not be absorbed. Show that U is unitary and explain how it correctly encodes the definition of η.

(b) Assume the photodetector is prepared in its ground state and compute the Kraus operators and POVM elements.

(c) Consider an equal superposition of the incident photon $\sim |0\rangle + |1\rangle$. Calculate the probability that the detector clicks. Give an interpretation of this result, i.e., compare the answer when $\eta = 1$ to the finite efficiency case $\eta \neq 1$; what is the physics of the difference?

2. **A Hamiltonian model for the amplifier**. As discussed above, our toy amplifier model Eq. (1.7) is essentially the CNOT gate. Here we will explore how this gate can be generated through Hamiltonian evolution, and how a "weak" version can be generated.

(a) Let $H = g(1_S - Z) \otimes (1_M - \tilde{X})$ be the "logical" Hamiltonian on the qubit + amplifier, where as above we're in the "logical" space $\tilde{\mathcal{H}}_M = \text{span}\{|\tilde{0}\rangle, |\tilde{1}\rangle\}$ for the amplifier, and the $Z, \tilde{X}$ are Pauli operators. The coupling constant g has units of a frequency. Compute $U(t) = e^{iHt}$ explicitly.

(b) Calculate how long the interaction should be turned on, $t = t_*$, to produce an exact CNOT gate.

(c) Consider the measurement discussed above, where the amplifier is initialized in $|\tilde{0}\rangle$, coupled to the qubit with $U(t)$, and then the amplifier is measured in the $|\tilde{0}, \tilde{1}\rangle$ basis. Compute the Kraus operators and POVM elements as a function of $\theta = gt$.

(d) Find the probabilities $P(\tilde{0})$ and $P(\tilde{1})$ with an arbitrary initial qubit state $|\psi\rangle = a|0\rangle + b|1\rangle$. For small $t \approx 0$, give an interpretation in terms of information gained about the qubit state.

(e) Calculate the post-measurement states of the qubit, given the outcome $\tilde{0}$ or $\tilde{1}$. Compare the limits $t \approx 0$ and $t = t_*$. Give an interpretation in terms of information gained about the qubit state, and the resulting "collapse".

2. Discrete-variable ("qubit") measurements

In this section, we will illustrate the basic formalism with a few choice examples where the system S has a finite dimensional Hilbert space. First we analyze the Stern-Gerlach experiment in substantially more detail than one usually encounters. We then move to an actual qubit architecture, a trapped ion or neutral atom, where the states $|0\rangle, |1\rangle$ refer to electron levels split by a hyperfine interaction.

2.1. Stern-Gerlach experiment

Let's begin with the Stern-Gerlach experiment, not because it is important (which it is) but because it is so simple that we can treat every aspect of it analytically.

Everyone knows the setup. We prepare a beam of electrically neutral, spin-1/2 particles.[1] This is sent through a localized magnetic field with a gradient $\partial B_z/\partial z \neq 0$. The magnetic dipole interaction can then be Taylor expanded, leading to a term

$$V = \boldsymbol{\mu} \cdot \mathbf{B}(\mathbf{x}) \supset \mu_0 \frac{\partial B_z}{\partial z} \sigma_z z \tag{2.1}$$

which leads to a force along the z-axis

$$F_z = \mu_0 \frac{\partial B_z}{\partial z} \sigma_z \tag{2.2}$$

which depends on the spin-z state of the atom. Here $\boldsymbol{\mu} = \mu_0 \boldsymbol{\sigma}$ is the magnetic moment operator on the spin. Thus, an atom prepared in $|0\rangle = |\downarrow\rangle$ will be pushed downwards and an atom prepared in $|1\rangle = |\uparrow\rangle$ will fly upwards. Thus, one can measure the spin state by seeing where the atom is in space after the interaction.

Now let us analyze this in detail, and in particular see what happens when we take into account the finite width of the beam and finite time of the free flight after the magnetic field interaction. The Hilbert space $\mathcal{H}_S$ represents the spin state (we will continue to use qubit notation $|0, 1\rangle$ as defined in the previous paragraph). The apparatus $\mathcal{H}_M$ will be represented by a one-dimensional particle, $\mathcal{H}_M = \text{span}\{|z\rangle \mid z \in \mathbf{R}\}$, i.e., as usual we suppress the actual amplifier and readout and just refer to a single degree of freedom into which we will encode the spin state. This choice of $\mathcal{H}_M$ is actually not quite right in a pretty interesting way, which we will discuss later in this section.[2]

[1] Historically, the experiment was done with neutral silver atoms, because they have vanishing nuclear spin and a single valence electron, so the whole atom has magnetic momentum given entirely by the electron spin.

[2] Warning: the discussion on this point will be heterodox; all the books I found just take for granted that this is the right description.

First we need to specify an initial state of the measurement apparatus M. Here this is an arbitrary position-space wavefunction $\psi_M(z)$; we will ignore the transverse directions. (Implicitly we are assuming that the beam travels with some fixed momentum along the x-axis). Thus let's specify a quasi-realistic initial condition, a Gaussian beam

$$|0\rangle_M = \int dz\, \psi_M(z)\, |z\rangle_M\,, \quad \psi_M(z) = N \exp\left\{-\frac{z^2}{4\Delta z_0^2}\right\}, \tag{2.3}$$

where $N = (2\pi\Delta z_0^2)^{1/4}$ is the usual Gaussian normalization. The initial uncertainty Δz_0 will play an important role in what follows. We are thinking of this as the state of the beam right when it hits the magnetic field, which we will think of as $t = 0$. This itself is an approximation since really there is also a wavepacket on the transverse axes and thus a finite spread and interaction time, but let's not get carried away.

Next, we need to figure out the unitary U that couples S to M. Let v be the beam velocity along the beam axis (say, x), and assume that the field gradient is localized to a length δx. If we assume that v is sufficiently large and δx is sufficiently small (you will make this precise in a homework exercise), we can approximate the time evolution as an instantaneous pulse

$$U_{\rm int} = \exp\left\{-i\Delta p z \sigma_z\right\}, \tag{2.4}$$

where the impulse delivered along the z-axis is

$$\int dt\, F_z \approx \frac{\delta x}{v}\mu_0 \frac{\partial B_z}{\partial z}\sigma_z =: \Delta p \sigma_z. \tag{2.5}$$

A common terminology is that we have a "state-dependent force". Notice that in $U_{\rm int}$, the quantity Δp is a number, while z and σ_z are operators. Specifically, since z is the generator of translations in momentum space, and σ_z is diagonal in the spin-z basis, what this operator does is to take

$$\begin{aligned} U_{\rm int}\, |0\rangle_S\, |p\rangle_M &= |0\rangle_S\, |p + \Delta p\rangle_M \\ U_{\rm int}\, |1\rangle_S\, |p\rangle_M &= |1\rangle_S\, |p - \Delta p\rangle_M\,, \end{aligned} \tag{2.6}$$

where here and after $p = p_z$ is the momentum in the z-axis, so $[z, p] = i$. The initial state [Eq. (2.3)] written in momentum space is

$$|0\rangle_M = \int dp\, \psi_M(p)\, |p\rangle\,, \quad \psi_M(p) = \frac{N}{\sqrt{2\pi}} \exp\left\{-\frac{p^2}{4\Delta p_0^2}\right\}, \quad \Delta p_0^2 = 1/4\Delta z_0^2, \tag{2.7}$$

another Gaussian centered around $p = 0$. The joint state thus evolves as

$$\begin{aligned} U_{\text{int}} \left|0\right\rangle_S \left|0\right\rangle_M &= \left|0\right\rangle_S \int dp\, \psi_M(p) \left|p + \Delta p\right\rangle_M \\ U_{\text{int}} \left|1\right\rangle_S \left|0\right\rangle_M &= \left|1\right\rangle_S \int dp\, \psi_M(p) \left|p - \Delta p\right\rangle_M . \end{aligned} \tag{2.8}$$

Finally, we have to deal with the fact that the particle will freely evolve under the pure kinetic term

$$H_{\text{free}} = \frac{\mathbf{p}^2}{2m} \tag{2.9}$$

between the interaction point $t = 0$ and when it hits the screen at $t > 0$.

The free evolution of a Gaussian wavepacket causes two effects: it moves the average location in position space $\overline{z}(t) = pt/m$, and it also causes wavepacket spreading $\Delta z(t) \sim t/(m\Delta z_0)$. Dealing with the wavepacket spreading is annoying. Let's assume that the beam freely evolves for t sufficiently short (again made precise in the homework) so that we can ignore it. To be more specific, let's re-write the post-interaction M states in position space

$$\begin{aligned} \int dp\, \psi_M(p) \left|p \pm \Delta p\right\rangle &= \int \frac{dz dp}{\sqrt{2\pi}}\, \psi_M(p) e^{i(p \pm \Delta p)z} \left|z\right\rangle \\ &= \int dz\, \psi_M(z) e^{\pm i \Delta p z} \left|z\right\rangle , \end{aligned} \tag{2.10}$$

using the fact that the state is Gaussian. This last expression is intuitive: the M states just after the interaction are Gaussians centered around $z = 0$ but now with some non-trivial momentum $\pm\Delta p$. Now let's approximate the free evolution as simply moving the centers of these packets by the appropriate amount:

$$U_{\text{free}}(t) \int dz\, \psi_M(z) e^{\pm i \Delta p z} \left|z\right\rangle \approx \int dz\, \psi_M \left(z \pm \frac{\Delta p}{m} t\right) \left|z\right\rangle . \tag{2.11}$$

This can be more formally justified using a stationary phase approximation.

The total evolution from the interaction is $U = U_{\text{free}}(t) U_{\text{int}}$. From the above results we have

$$\begin{aligned} U \left|0\right\rangle_S \left|0\right\rangle_M &= \left|0\right\rangle_S \left|0(t)\right\rangle_M \\ U_{\text{int}} \left|1\right\rangle_S \left|0\right\rangle_M &= \left|1\right\rangle_S \left|1(t)\right\rangle_M , \end{aligned} \tag{2.12}$$

where the conditional states of the "measurement apparatus", i.e., the states of the motional degree of freedom of the atom, are

$$\begin{aligned} \left|0(t)\right\rangle_M &= \int dz\, \psi_M \left(z - \frac{\Delta p}{m} t\right) \left|z\right\rangle \\ \left|1(t)\right\rangle_M &= \int dz\, \psi_M \left(z + \frac{\Delta p}{m} t\right) \left|z\right\rangle , \end{aligned} \tag{2.13}$$

that is, position-space Gaussians shifted by a relative amount $2\Delta pt/m$.

Finally, we can consider what any of the above means in terms of the actual measurement. Let's continue to follow the general formalism given in the first chapter. We assume that we can do a von Neumann measurement on the position variable, i.e., projectively measure it. Notice that, in general for finite $t > 0$, the two M states are not orthogonal. Their overlap is just the overlap of two shifted Gaussians:

$$\langle 0(t)|1(t)\rangle_M = \exp\left\{-\frac{(2\Delta pt/m)^2}{\Delta z_0^2}\right\} \neq 0. \tag{2.14}$$

which means that we *cannot do a projective measurement in the* $|0(t)\rangle, |1(t)\rangle$ *basis*. Of course, as $t \to \infty$, the overlap goes to zero and then the projective measurement makes sense; this is the limit in which the experiment is usually implicitly described. However, we can be much more general.

Since we can't do a projective measurement on the $|0(t)\rangle, |1(t)\rangle$ states, let's resort to what we *can* do: projectively measure the position z. This means that the measurement outcomes $a = z$ are drawn from a continuous spectrum, i.e., the atom in principle can land at any value of $-\infty < z < \infty$. Using Eqs. (2.12) and (2.13), the Kraus operators are

$$K_z = \langle z|U|0\rangle_M = \psi_M\left(z - \frac{\Delta p}{m}t\right)|0\rangle\langle 0|_S + \psi_M\left(z + \frac{\Delta p}{m}t\right)|1\rangle\langle 1|_S, \tag{2.15}$$

which in turn gives the POVM elements

$$E_z = \left|\psi_M\left(z - \frac{\Delta p}{m}t\right)\right|^2 |0\rangle\langle 0|_S + \left|\psi_M\left(z + \frac{\Delta p}{m}t\right)\right|^2 |1\rangle\langle 1|_S. \tag{2.16}$$

It is instructive to check that these satisfy the completeness relation $\int dz E_z = 1_S$.

To get some intuition for what this POVM is telling us, consider the simple case where we prepare the beam in the equally weighted superposition state

$$|\psi\rangle_S = (|0\rangle_S + |1\rangle_S)/\sqrt{2}. \tag{2.17}$$

The probability distribution of hits on the screen is then given by

$$P(z) = \langle\psi|E_z|\psi\rangle = \frac{1}{2}\left|\psi_M\left(z - \frac{\Delta p}{m}t\right)\right|^2 + \frac{1}{2}\left|\psi_M\left(z + \frac{\Delta p}{m}t\right)\right|^2. \tag{2.18}$$

This is just a sum of two Gaussians whose peaks are shifted by $2\Delta pt/m$. In the limit $t \to \infty$ these peaks are cleanly separated and we recover the usual Stern-Gerlach measurement, where the two spin states are sent to two distinct locations

on the screen. In the limit $t \approx 0$, on the other hand, the two Gaussians are nearly overlapping, and we get very little information.

Consider, instead of the pure state (2.17), preparing the totally depolarized state

$$\rho_S = \begin{pmatrix} 1/2 & 0 \\ 0 & 1/2 \end{pmatrix} \tag{2.19}$$

which is a classical statistical mixture of spin up and down. The distribution of hits on the screen is

$$P(z) = \mathrm{tr}_S(\rho_S E_z) = \frac{1}{2}\left|\psi_M\left(z - \frac{\Delta p}{m}t\right)\right|^2 + \frac{1}{2}\left|\psi_M\left(z + \frac{\Delta p}{m}t\right)\right|^2, \tag{2.20}$$

identical to (2.18)! The key point is that the POVM and Kraus operators are *diagonal* in the spin-z basis. Physically, this reflects the fact that the measurement is insensitive to any "coherence" in the quantum state of the spin, and it ultimately traces back to the interaction (2.6) being diagonal in the spin-z basis.

Finally, let's introduce a little information theory. We can quantify the information "gain" in this measurement in terms of entropy. Consider preparing the mixed state Eq. (2.19), which has von Neumann entropy

$$S(0) = -\,\mathrm{tr}_S(\rho_S \ln \rho_S) = -\left[\frac{1}{2}\ln\frac{1}{2} + \frac{1}{2}\ln\frac{1}{2}\right] = \ln 2. \tag{2.21}$$

The second equality is supposed to remind you how to compute a von Neumann entropy: diagonalize ρ_S and just sum over the eigenvalues. This state has the most possible entropy for a two-dimensional system. Now, suppose we prepare the mixed state (2.19) and do a measurement, obtaining a hit on the screen at position z. The post-measurement state is

$$\rho_S \to \rho_S(z) = \frac{K_z \rho_S K_z^\dagger}{P(z)} = N(t)\begin{pmatrix} e^{-\lambda_-(t)} & 0 \\ 0 & e^{-\lambda_+(t)} \end{pmatrix}, \tag{2.22}$$

where

$$N(t) = \frac{1}{e^{-\lambda_-(t)} + e^{-\lambda_+(t)}}, \quad \lambda_\pm(t) = \frac{(z - \Delta p t/m)^2}{2\Delta z_0^2}. \tag{2.23}$$

Consider the case of a weak measurement, which here means small t, i.e., we don't let the particle propagate long enough to clearly distinguish the two Gaussians on the screen. In this case, we can approximate

$$\rho_S(z) \approx \begin{pmatrix} \frac{1}{2} - \frac{z\Delta p t}{2\Delta z_0^2} & 0 \\ 0 & \frac{1}{2} + \frac{z\Delta p t}{2\Delta z_0^2} \end{pmatrix} \tag{2.24}$$

which is just the original, totally mixed state, plus a small correction. The entropy works out to

$$S(z) \approx \ln 2 - \frac{z^2 \Delta p^2 t^2}{2\Delta z_0^4}. \tag{2.25}$$

What this says is that our post-measurement state, conditioned on the outcome that we got a hit at location z, has slightly lower entropy than the original, as it should. You can check various other intuitive facts: for example, if the hit is far from the origin (larger $|z|$), the entropy decreases more; if the initial beam is more collimated (smaller Δz_0), the entropy decreases more; etc.

Finally, a comment on the state update rule in this system. The preceding discussion about entropy used the state update $\rho \to K_z \rho K_z^\dagger / P(z)$ to calculate the state of the spin S after the measurement. The answer is that we have a spin which is slightly less polarized, which sounds entirely reasonable. In the case of the initial pure state (2.17) and a projective measurement ($t \to \infty$, for example), the state update would say that the spin is projected into either the up or down state. But is it? What is the actual, physical condition of the spin after these measurements? It's a mess: the atom gets absorbed by the screen, winds up in some complicated bound state with the other screen atoms, etc. There's absolutely no way that the spin is just sitting there in some pure spin up state. It is instructive to decide for yourself whether this is actually the right description, or if we implicitly mis-applied one of the measurement rules outlined in Table 2. I'll leave this exercise to you in the homework. In the next section, we will instead give another kind of qubit mesurement where the physical picture of the post-mesurement state is much more clear.

2.2. Hyperfine atomic qubit

Now we move on to a qubit architecture that is actually used in practice. Consider some atom with a single valence electron, for example neutral hydrogen H or an ion like Be^+. The electrons form a discrete set of bound states $|n\ell m\rangle$, with discrete energy levels. A beautiful way to encode a qubit is to find two such levels, call them $|0\rangle$ and $|1\rangle$, split by an energy $\Delta E_{01} = \omega$ lying in either the microwave (100 MHz $\lesssim \omega \lesssim$ 10 GHz) or optical (100 GHz $\lesssim \omega \lesssim$ 1000 THz) range [1]. Typically these splittings come from the perturbation due to the hyperfine coupling of the nucleus to the electron, thus the name. Applying drives at this frequency for different durations (pulse lengths) can then enact arbitrary single-qubit gates, using standard Rabi oscillation physics. Coupling these atoms together to form n-body gates can be

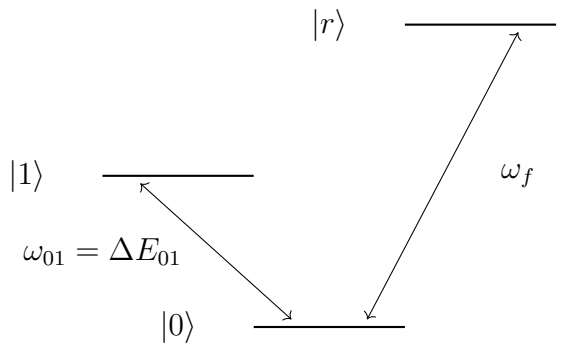

Figure 3: **Level diagram for Dehmelt's electron shelving method.** Here $|0\rangle, |1\rangle$ represent two electron energy levels in the atom, typically split by a hyperfine interaction. The level $|r\rangle$ is connected to the $|0\rangle$ state by a fluorescence transition but not to the state $|1\rangle$, at least to leading order. We assume that the spontaneous emission timescales $\tau_{10} \gg \tau_{r0}$. Adapted from Liebfried, Blatt, Monroe, and Wineland, *Quantum Dynamics of Single Trapped Ions*, Rev. Mod. Phys. 75 (2003).

done in a variety of ways, for example forming ionic crystals and using the phonons, as in the proposal of Cirac and Zøller.

Here, we will focus on how the readout of a single such qubit can be performed, using a method first suggested by Dehmelt in 1975 known as "electron shelving". This is a nice example which shows how a projective measurement arises from a sequence of non-projective measurements.

The basic goal is to engineer a measurement in the $|0\rangle, |1\rangle$ basis. In electron shelving, we also make use of some third level $|r\rangle$ in the atomic spectrum, with the following properties:

- There is a transition $|r\rangle \to |0\rangle$, in which a photon with frequency ω_f is spontaneously emitted. This transition is "strong" in the sense that it happens with very high probability in a very short time τ_f. The subscript f stands for fluorescence. Obviously there is also a transition $|0\rangle \to |r\rangle$ by unitarity.

- There is *not* a transition $|r\rangle \leftrightarrow |1\rangle$, or at least this transition is very suppressed compared to the $|r\rangle \leftrightarrow |0\rangle$ transition.

These energy levels and transitions are summarized in Fig. 3. So although we fundamentally want to use $|0\rangle$ and $|1\rangle$ as the qubit, the actual system Hilbert space is really three-dimensional: $\mathcal{H}_S = \text{span}\{|0\rangle, |1\rangle, |r\rangle\}$.

For the measurement apparatus $\mathcal{H}_M$ we will use photons of frequency ω_f. To highlight the basic quantum mechanics, we will think of the measurement apparatus M as a single photon (or more accurately, a sequence of single photons, discussed at the end). To describe its Hilbert space we start with a simple picture of the essential physics.

Consider what happens when we send a single photon, say in a collimated laser beam so it has fixed momentum $|\mathbf{k}\rangle_M$ with frequency $\omega_f = |\mathbf{k}|$, at the atom. If the atom is initialized in $|0\rangle$, then

$$|0\rangle_S \otimes |\mathbf{k}\rangle_M \to A_e \, |r\rangle_S \otimes |\text{vac}\rangle_M + \sqrt{1 - |A_e|^2} \, |0\rangle_S \otimes |\mathbf{k}\rangle_M \,, \tag{2.26}$$

where A_e is the absorption amplitude for the $|0\rangle \to |r\rangle$ transition, and $|\text{vac}\rangle_M$ is the state of the electromagnetic field once the photon is absorbed. On the other hand, if the atom is in $|1\rangle$ or $|r\rangle$, then to leading order nothing happens,

$$\begin{aligned} |1\rangle_S \otimes |\mathbf{k}\rangle_M &\to |1\rangle_S \otimes |\mathbf{k}\rangle_M \\ |r\rangle_S \otimes |\mathbf{k}\rangle_M &\to |r\rangle_S \otimes |\mathbf{k}\rangle_M \,. \end{aligned} \tag{2.27}$$

Now, after a time of order τ_f, if the photon was absorbed and the atom excited to $|r\rangle$, the fluorescent decay will occur, producing an emitted photon state $|\gamma\rangle_M$ which is no longer collimated but has some kind of angular distribution (e.g., a dipole pattern if $|r\rangle \to |0\rangle$ is an allowed dipole transition),

$$|r\rangle_S \otimes |\text{vac}\rangle_M \to |0\rangle_S \otimes |\gamma\rangle_M \,. \tag{2.28}$$

Putting all of this together, we can model the Hilbert space of the light as $\mathcal{H}_M = \text{span}\{|\text{vac}\rangle, |\mathbf{k}\rangle, |\gamma\rangle\}$, and the interaction as a unitary whose important matrix elements are

$$\begin{aligned} U \, |0\rangle_S \, |\mathbf{k}\rangle_M &= A_e \, |0\rangle_S \, |\gamma\rangle_M + \sqrt{1 - |A_e|^2} \, |0\rangle_S \, |\mathbf{k}\rangle_M \\ U \, |1\rangle_S \, |\mathbf{k}\rangle_M &= |1\rangle_S \, |\mathbf{k}\rangle_M \\ U \, |r\rangle_S \, |\mathbf{k}\rangle_M &= |r\rangle_S \, |\mathbf{k}\rangle_M \,. \end{aligned} \tag{2.29}$$

At this stage, we can forget about the $|r\rangle$ state, since an initial state in the $\text{span}\{|0\rangle, |1\rangle\}$ subspace will remain in that subspace under (2.29). We can also forget about the state $|\text{vac}\rangle$ of the electromagnetic field, for the same reason.

Now, let's use the interaction (2.29) to enact a measurement of the qubit state. We prepare the initial state of the measurement apparatus, the single photon:

$$|0\rangle_M = |\mathbf{k}\rangle_M \tag{2.30}$$

Figure 4: **Trapped ion qubit readout with Dehmelt shelving.** Each row corresponds to a different measurement of 53 trapped ion qubits, showing either a bright spot ($|0\rangle$) or dark spot ($|1\rangle$), in our notation. Figure from Chris Monroe's group at Maryland/IonQ. See https://iontrap.umd.edu/2017/12/01/quantum-simulation-with-individual-control-of-53-qubits/.

and then assume that we can projectively measure photon in the basis $|\mathbf{k}\rangle, |\gamma\rangle$. This is a big simplification: to measure the $|\gamma\rangle$ state projectively would physically require having a perfect photodetector with $4\pi - \epsilon$ coverage, where the ϵ is a second photodetector aligned right on the beam spot defined by $\mathbf{k}$. See Fig. **??**. In reality the photodetector will have some less-than-full coverage and finite efficiency, but as we will see shortly this isn't very important because we can repeat this measurement with many photons.

With this setup, the Kraus operators are easy to calculate:

$$\begin{aligned} K_0 &= \langle 0|U|0\rangle_M = \sqrt{1-|A_e|^2}\,|0\rangle\langle 0|_S + |1\rangle\langle 1|_S \\ K_\gamma &= \langle \gamma|U|0\rangle_M = A_e\,|0\rangle_S\langle 0|_S\,. \end{aligned} \tag{2.31}$$

The corresponding POVM is

$$\begin{aligned} E_0 &= \left(1-|A_e|^2\right)|0\rangle\langle 0|_S + |1\rangle\langle 1|_S \\ E_\gamma &= |A_e|^2\,|0\rangle_S\langle 0|_S\,. \end{aligned} \tag{2.32}$$

These clearly satisfy the completeness relation $E_0 + E_\gamma = 1_S$ (where we've dropped the $|r\rangle$ state from the S basis).

Ok, now let's see how this produces a qubit measurement. Prepare the qubit in

an arbitrary superposition

$$|\psi\rangle_S = a\,|0\rangle_S + b\,|1\rangle_S\,. \tag{2.33}$$

The outcome of a measurement is that we we get either an unscattered photon $|\mathbf{k}\rangle_M = |0\rangle_M$ or a scattered one $|\gamma\rangle_M$. These occur with probabilities

$$\begin{aligned} P(\text{unscat}) = P(0) &= \langle\psi|E_0|\psi\rangle_S = \left(1-|A_e|^2\right)|a|^2 + |b|^2 \\ P(\text{scat}) = P(\gamma) &= \langle\psi|E_\gamma|\psi\rangle_S = |A_e|^2|a|^2. \end{aligned} \tag{2.34}$$

Notice in particular that the photon will scatter only if the qubit was in $|0\rangle$. In contrast, the photon can go through unscattered whether the qubit was in $|1\rangle$ or $|0\rangle$, although the latter is not likely, occurring with a probability proportional to $\epsilon^2 := 1 - |A_e|^2 < 1$. This factor ϵ may be much less than one $\epsilon \ll 1$ but could be almost 1, but it is definitely less than 1, because the absorption process is unitary. Thus, if we get a scattered photon we immediately know the qubit state, but if we get an unscattered photon we have some remaining uncertainty.

To make this more concrete, consider the post-measurement state, using the usual rule $|\psi\rangle \to K_{\text{outcome}}\,|\psi\rangle\,/\sqrt{P(\text{outcome})}$. Here this is

$$|\psi\rangle_S \to \begin{cases} \dfrac{\sqrt{1-|A_e|^2}a\,|0\rangle_S + b\,|1\rangle_S}{\sqrt{(1-|A_e|^2)|a|^2+|b|^2}} & \text{given outcome ``unscattered'' } (0) \\ \dfrac{A_e a}{\sqrt{|A_e|^2|a|^2}}\,|0\rangle_S & \text{given outcome ``scattered'' } (\gamma) \end{cases}. \tag{2.35}$$

In the case that we see the photon get scattered, the resulting state is simply $|0\rangle_S$, up to some phase. That is, the qubit is projected to the $|0\rangle_S$ state. On the other hand, if the photon goes through unscattered, the resulting state

$$|\psi_0\rangle_S = \frac{\sqrt{1-|A_e|^2}a\,|0\rangle_S + b\,|1\rangle_S}{\sqrt{(1-|A_e|^2)|a|^2+|b|^2}} \approx \frac{a}{|b|}\epsilon\,|0\rangle_S + \frac{b}{|b|}\,|1\rangle_S \tag{2.36}$$

is still a coherent superposition of $|0\rangle_S$ and $|1\rangle_S$. However, the coefficient in front of $|1\rangle_S$ is order one, while the coefficient in front of $|0\rangle_S$ is order $\epsilon < 1$, so the state is "more" $|1\rangle_S$ than initially. Unlike the Stern-Gerlach example, here it is physically very clear that the qubit really is in this state after the measurement. In words, the first photon measurement has partially collapsed the qubit state into either $|0\rangle$ or $|1\rangle$.

Since the qubit state really physically changes like this, we can do the measurement again! What happens? If the photon was scattered initially, then the qubit is in $|0\rangle_S$ with unit probability, and therefore a subsequent measurement (another photon) will

leave the qubit in $|0\rangle_S$ [see Eq. (2.29)]. On the other hand, if the initial photon went unscattered, the probability of getting *another* unscattered photon is

$$\begin{aligned} P(0|0) &= \langle\psi_0|E_0|\psi_0\rangle_S \\ &= \frac{(1-|A_e|^2)^2|a|^2+|b|^2}{(1-|A_e|^2)|a|^2+|b|^2} \\ &\approx 1-\frac{|a|^2}{|b|^2}\epsilon^2. \end{aligned} \tag{2.37}$$

What this means is that if the first photon was unscattered, then with probability nearly one, the second photon will also be unscattered, which gives us more confidence that the qubit is in the state $|1\rangle_S$. Notice that the "error" is now down by another power of ϵ: the qubit state becomes

$$\begin{aligned} |\psi\rangle \to |\psi_0\rangle &= \frac{K_0\,|\psi\rangle}{\sqrt{P(0)}} \\ \to |\psi_{0,0}\rangle &= \frac{K_0K_0\,|\psi\rangle}{\sqrt{P(0)P(0|0)}} \\ &\approx \frac{a}{|b|^2}\epsilon^2\,|0\rangle+\frac{b}{|b|}\,|1\rangle\,, \end{aligned} \tag{2.38}$$

where you can check the last equation easily. Again, we see that the state is "more" projected onto the $|1\rangle_S$ state.

The upshot is that if we send in a bunch of photons, the first few act to collapse the qubit state down to either $|0\rangle_S$ or $|1\rangle_S$, and then all the subsequent ones will simply scatter or not, accordingly. The actual data is simple: you set up some CCDs or other photon collectors in a solid angle around the qubit, and either you see a bright spot (for a qubit in $|0\rangle_S$, which causes all the light to scatter) or you see a dark spot (for a qubit in $|1\rangle_S$, which causes almost all the light to pass through unscattered). See Fig. 4. In this way we turn a partial/weak measurement into an effectively projective measurement, including the amplification step. This is, in practice, how almost all projective measurements arise.

Further reading

- Jordan and Siddiqui, ch 2-4
- Steve Girvin's lecture notes on circuit QED, ch 7.
  Available at https://girvin.sites.yale.edu/lectures.

Problems

1. **Stern-Gerlach approximations and interpretation.**

   (a) Write an expression for the exact time evolution of the beam just before it crosses the magnetic field until just after. Do not forget the free evolution from the kinetic energy. Find quantitative conditions on the beam velocity v and width of the field δx that justify our approximation of an instantaneous interaction.

   (b) Suggest a method to non-destructively detect the position of the silver atom, rather than the destructive measurement made by the screen.

2. **Non-orthogonal states cannot be distinguished in a single measurement.** Let $|\psi\rangle$ and $|\phi\rangle$ be two states of some system S, and assume that they are not orthogonal $\langle\psi|\phi\rangle \neq 0$. Show that there is no measurement that can determine with complete certainty which state the system was prepared in. Hint: use the fact that unitary evolution preserves inner products. [The purpose of this problem is to justify the discussion around Eq. (2.14), although there we applied the argument to M, not S. This is problem 1.17 in K. Jacobs, *Quantum Measurement Theory and its Applications*].

3. **Entropy and the approach to a projective measurement**. Consider our hyperfine measurement model. For simplicity assume A_e is real, and assume the non-absorption probability is small: $1 - |A_e|^2 = \epsilon^2 \ll 1$. The answer to all of these problems can be expressed just to leading order in ϵ.

   (a) Prepare the qubit in the totally mixed state of equal weights. Find the states ρ_0, ρ_γ of the qubit after a single photon measurement, given outcome 0 or γ.

   (b) Calculate the von Neumann entropy of the state ρ_0 and compare it to the initial entropy.

   (c) Now suppose we do another measurement and get 0 again. Find the conditional state $\rho_{0,0}$, and calculate its entropy.

   (d) Conjecture or prove a formula for the entropy of the state assuming N consecutive outcomes of an unscattered photon $0, 0, 0, \ldots$. Interpret this in terms of purity of the resulting state.

References

[1] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland, "Quantum dynamics of single trapped ions," *Reviews of Modern Physics* **75** no. 1, (2003) 281.

3. Continuous-variable measurements

We now turn to measurements on continuous variable systems, where $\dim \mathcal{H} = \infty$. We will be non-rigorous in the usual ways; in particular we assume that we can make sense of the Hilbert space in the first place and don't have to worry about things like the von Neumann algebra type. In practice we are just going to talk about things isomorphic to a non-relativistic particle moving in one spatial dimension.

The main difference between this and discrete-variable systems is that we have to deal with continuous bases like position $|x\rangle$ or momentum $|p\rangle$. One should be skeptical of the operational meaning of these bases. Does it really make sense to make projective measurements in the $|x\rangle$ basis? If we did that, the post-measurement state would be a position eigenstate. Put another way, its momentum state would be completely flat in momentum, i.e., would have components at arbitrarily high energy. Such states would be pathological for many reasons! In reality, measurements of continuous variables always come with some level of "error", or in our language, are not completely strong/projective in the basis.

We begin with a simple model of a quasi-realistic position measurement. This will set the stage for a discussion of what Heisenberg uncertainty means and what actually happens if we sequentially measure a pair of non-commuting observables. We then move on to the most important cases of real measurements in the continuous domain: measurements of electromagnetic modes.

3.1. Position measurements

Consider making a measurement of the location of a particle S prepared in some state

$$|\psi\rangle_S = \int_{-\infty}^{\infty} dx\, \psi_S(x)\, |x\rangle_S \,. \tag{3.1}$$

We we could do this, for example, with a time-of-flight measurement: we prepare some other particle M (say, a photon) which we model as starting in some Gaussian wavepacket

$$|0\rangle_M = \int_{-\infty}^{\infty} dx\, \frac{e^{-x^2/4\Delta x_M^2}}{(2\pi\Delta x_M^2)^{1/4}}\, |x\rangle_M \,. \tag{3.2}$$

We then shoot M at S, assume it reflects, and measure where M is after some time t (or really, we set up some M detector and ask for the time t at which M is reflected back to the detector). Crucially, the detector particle M has its own position-space uncertainty Δx_M, which means that it will only give us imperfect information about

the location of S. Given the discussion in the previous sections, it seems natural to model this whole process as a set of Kraus operators

$$K_x = \int_{-\infty}^{\infty} dy \, \frac{e^{-(x-y)^2/4\Delta x_M^2}}{(2\pi\Delta x_M^2)^{1/4}} \, |y\rangle \langle y|_S \tag{3.3}$$

on the system. Here x means the outcome of the measurement of the position of S, as usual. It is instructive to demonstrate that the POVM

$$E_x = K_x^\dagger K_x = \int_{-\infty}^{\infty} dy \, \frac{e^{-(x-y)^2/2\Delta x_M^2}}{(2\pi\Delta x_M^2)^{1/2}} \, |y\rangle \langle y|_S \tag{3.4}$$

satisfies the usual completeness relation $\int dx E_x = 1_S$ (notice that the width in the Gaussian changed by a factor of 2, as did the normalization power). One could derive these Kraus operators by explicitly modeling the interaction U between M and S as we did above, but let's just start directly at the Kraus operator level. What this POVM represents is a measurement of the S position, but the outcomes have a measurement uncertainty Δx_M due to the "imperfect" apparatus.

What happens when we measure the position of S? To get some intuition, consider preparing S in a state which is localized in space, centered at $x = x_0$ with some spread Δx_S^2, say the Gaussian state:

$$|\psi_{x_0}\rangle = \int_{-\infty}^{\infty} dx \, \frac{e^{-(x-x_0)^2/4\Delta x_S^2}}{(2\pi\Delta x_S^2)^{1/4}} \, |x\rangle \, . \tag{3.5}$$

Here and for the rest of this section, we are only dealing with the system S states, so we suppress the S subscripts. The probability of "finding the particle at x" (i.e., obtaining outcome x), is

$$\begin{aligned} P(x) &= \langle\psi_{x_0}|E_x|\psi_{x_0}\rangle \\ &= \left(\int_{-\infty}^{\infty} dy \, \frac{e^{-(y-x_0)^2/4\Delta x_S^2}}{(2\pi\Delta x_S^2)^{1/4}} \, \langle y|\right) \times \left(\int_{-\infty}^{\infty} dz \, \frac{e^{-(x-z)^2/2\Delta x_M^2}}{(2\pi\Delta x_M^2)^{1/4}} \, |z\rangle \langle z|\right) \\ &\times \left(\int_{-\infty}^{\infty} dy' \, \frac{e^{-(y'-x_0)^2/4\Delta x_S^2}}{(2\pi\Delta x_S^2)^{1/4}} \, |y'\rangle\right) \\ &= \frac{e^{-(x-x_0)^2/4\Delta x_{\text{tot}}^2}}{(2\pi\Delta x_{\text{tot}}^2)^{1/2}}, \end{aligned} \tag{3.6}$$

where

$$\Delta x_{\text{tot}}^2 = \Delta x_S^2 + \Delta x_M^2 \tag{3.7}$$

represents the *total* uncertainty in the measurement. The first term comes from the uncertainty in the system state itself, while the second comes from the imperfect

nature of the measurement. This is, of course, a normalized Gaussian distribution centered at the (most likely value of the) position of the system S, namely $x = x_0$. The uncertainties add in quadature, as one would expect.

The probability distribution of outcomes, Eq. (3.6), is wider than both the initial system state and measurement uncertainty. What about the post-measurement state? One might expect that the width narrows, if we do a good measurement. In general, given that we find the particle at location x, the state update rule gives

$$|\psi_{x_0}\rangle \to \frac{K_x |\psi_{x_0}\rangle}{\sqrt{P(x)}} = \int_{-\infty}^{\infty} dy \, \frac{e^{-(x_0'-y)^2/4(\Delta x_S')^2}}{[2\pi(\Delta x_S')^2]^{1/4}} |y\rangle \, . \tag{3.8}$$

This can be seen by straightforward calculation, using the fact that the product of two Gaussians is a new Gaussian. The center and width have shifted:

$$\begin{aligned} x_0 \to x_0' &= x_0 \frac{\Delta x_M^2}{\Delta x_{\rm tot}^2} + x \frac{\Delta x_S^2}{\Delta x_{\rm tot}^2} \\ \Delta x_S^2 \to (\Delta x_S')^2 &= \frac{\Delta x_S^2 \Delta x_M^2}{\Delta x_{\rm tot}^2}. \end{aligned} \tag{3.9}$$

These results seem natural. The post-measurement state is centered somewhere between the original state's center x_0 and the measurement outcome x, with relative weights depending on the uncertainty in the two factors.

For example, suppose we do a reasonably "strong" measurement, where the measurement uncertainty is much less than the initial uncertainty in the state: $\Delta x_M \ll \Delta x_S$, and get outcome x. Then we have $\Delta x_{\rm tot}^2 \approx \Delta x_S^2$, and so

$$\begin{aligned} x_0' &\approx x \\ (\Delta x_S')^2 &\approx \Delta x_M^2. \end{aligned} \tag{3.10}$$

This is basically what you would expect from a projective measurement of position, except that the resulting state has a finite-width distribution Δx_M^2, which heuristically goes to zero in the projective limit.

3.2. Heisenberg's microscope and measurement back-action

Let's connect the discussion above to a commonly-repeated interpretation of Heisenberg's uncertainty relation

$$\Delta x^2 \Delta p^2 \geq \frac{1}{4}. \tag{3.11}$$

One often hears that Eq. (3.11) means that if you first measure the position of something, then measure its momentum, that the product of the uncertainties of the

measurements is lower-bounded by 1/4. This is incorrect. The right interpretation of Eq. (3.11) is to first write it as $\langle \Delta x^2 \Delta p^2 \rangle \geq 1/4$, where the expectation value is taken *in some particular quantum state at one fixed time*. This is a statement about joint probability distributions in a *fixed state* and has nothing to do with a pair of measurements on the same system.

We are, however, now in position to directly analyze Heisenberg's microscope experiment, where one literally measures the position of a system S, gets some outcome, and then measures the momentum. We prepare the system S in a localized Gaussian state $|\psi_{x_0}\rangle$ as above, measure the position with the POVM of width Δx_M^2, and then measure it again using the analogous momentum-space POVM. Let's work in the limit of a "strong" measurement $\Delta x_M^2 \ll \Delta x_S^2$, as in Eq. (3.10). The first set of measurement outcomes has variance

$$\Delta x_{\text{meas}}^2 = \Delta x_{\text{tot}}^2 \approx \Delta x_S^2, \tag{3.12}$$

where Δx_S^2 is the uncertainty of the initial state itself. The approximation is justified by the assumption that we made a strong measurement. Now, the state after the measurement has a position space uncertainty

$$\Delta x_S^2 \to (\Delta x_S')^2 \approx \Delta x_M^2 \ll \Delta x_S^2, \tag{3.13}$$

as in Eq. (3.10). The position space wavefunction has narrowed. But then the *momentum* space wavefunction must have increased,

$$\Delta p_S^2 \to (\Delta p_S')^2 = \frac{1}{4(\Delta x_S')^2} \gg \Delta p_S^2. \tag{3.14}$$

This follows by direct calculation [e.g., by inserting a complete set of momentum eigenstates into Eq. (3.8)], but is also clear from the usual Heisenberg uncertainty relation.

Suppose that we now measure p, where we describe the measurement using a POVM of width Δp_M^2. Then the distribution of outcomes is identical to the position-space results (since everything here is Gaussian), and the analogous width of the outcome distribution is

$$(\Delta p')_{\text{meas}}^2 = (\Delta p_S')^2 + \Delta p_M^2 = \frac{1}{4(\Delta x_S')^2} + \Delta p_M^2 \geq \frac{1}{4(\Delta x_S')^2}, \tag{3.15}$$

using Eq. (3.14). Thus we have the product of uncertainties in the two measurements:

$$(\Delta x)_{\text{meas}}^2 (\Delta p')_{\text{meas}}^2 \geq \frac{\Delta x_S^2}{4(\Delta x_S')^2} \gg \frac{1}{4}, \tag{3.16}$$

where the $\gg$ comes from our use of a strong initial measurement, Eq. (3.10), which says that $(\Delta x'_S)^2 \approx \Delta x_M^2 \ll \Delta x_S^2$.

What Eq. (3.16) tells us is that the product of uncertainties on this pair of measurements is actually *substantially worse* than the simple product formula in the Heisenberg uncertainty relation, Eq. (3.11). Notice that here everything was done with Gaussian states that minimize Heisenberg uncertainty at each time step; we are not doing anything sketchy by introducing large uncertainties by hand. One might be skeptical that we somehow picked a sub-optimal pair of measurements; you will explore this in a homework problem.

3.3. Number and quadrature bases; homodyne measurements

In this section we move on to measurements of harmonic oscillators, described by the Hamiltonian

$$H = \omega a^\dagger a, \tag{3.17}$$

where as usual $[a, a^\dagger] = 1$. The most important example of these in practice are the modes of the electromagnetic field, where a mode means excitations of the field with a specific frequency ω and spatial momentum $\mathbf{k}$. See Appendix A for a review of the quantization of the electromagnetic field and its decomposition into such modes. In the rest of this section we will often discuss the interpretation of the oscillator in terms of a single mode of the electromagnetic field, in which case it describes photons of fixed frequency and momentum. There are a number of important bases for harmonic oscillators, which have very different properties in measurements. The most commonly encountered bases are:

Number basis: These are defined as eigenstates of the number operator $\hat{n}\,|n\rangle = n\,|n\rangle$, where $\hat{n} = a^\dagger a$ and $n = 0, 1, \ldots$ is an integer. These are energy eigenstates [i.e., eigenstates of the Hamiltonian, Eq. (3.17)]. For EM field modes, these are states with a definite number of photons of the same fixed frequency and momentum.

Quadrature basis: Quadratures are the generalization of the position and momentum bases for a textbook harmonic oscillator. For any angle θ, we define the Hermitian operator

$$X_\theta := \frac{1}{\sqrt{2}}(a e^{i\theta} + a^\dagger e^{-i\theta}), \tag{3.18}$$

which is known as the quadrature variable at that angle. In particular, the quadratures $\theta = 0$ and $\theta = \pi$ are often used and given particular names,

$$X = \frac{1}{\sqrt{2}}(a + a^\dagger), \quad Y = -\frac{1}{\sqrt{2}}(a - a^\dagger). \tag{3.19}$$

For a mechanical oscillator, these are of course the position and momentum variables, rescaled by their zero-point fluctuations to be dimensionless operators. In the EM case, these are referred to as the amplitude and phase quadratures, for reasons discussed in App. A. They do not commute, but satisfy the canonical commutation relation

$$[X, Y] = i. \tag{3.20}$$

Any θ quadrature can be used to define a continuous basis, via $\hat{X}_\theta |X_\theta\rangle = X_\theta |X_\theta\rangle$, where the eigenvalues X_θ run over the real numbers. These are just like the position or momentum basis of a mechanical oscillator, so in particular the states are Dirac-normalized and can be used to write the wavefunction of a general pure state. For example, in the amplitude basis,

$$\langle X|X'\rangle = \delta(X - X'), \quad |\psi\rangle = \int_{-\infty}^{\infty} dX\, \psi(X)\, |X\rangle \tag{3.21}$$

with analogous expressions holding for any other value of θ. In particular we have standard Fourier relations like $\langle X|Y\rangle = e^{iXY}/\sqrt{2\pi}$.

Coherent state basis: The number and quadratures bases defined above are the most common bases in which an oscillator is actually measured, as we will discuss below. However, the coherent states, defined as eigenstates of the annihilation operator,

$$\hat{a}\, |\alpha\rangle = \alpha\, |\alpha\rangle\,, \qquad \alpha \in \mathbf{C} \tag{3.22}$$

are also commonly encountered. These states are not orthogonal to each other but rather satisfy

$$\langle\alpha|\alpha'\rangle = \exp\left\{-\frac{1}{2}\left(|\alpha|^2 + |\alpha'|^2 - 2\alpha^*\alpha'\right)\right\}, \quad \text{i.e. } |\langle\alpha|\alpha'\rangle|^2 = \exp\left\{-\frac{1}{2}\left(|\alpha - \alpha'|^2\right)\right\}, \tag{3.23}$$

which turns out to mean that they form an *over-complete* basis.[3] The key property of coherent states we will need immediately is that they behave somewhat like classical states. To be precise, note that a coherent state does not have a definite number

[3] In particular, these give an extremely useful representation of an arbitrary density matrix

$$\rho = \int_{\mathbf{C}} d\alpha\, P(\alpha)\, |\alpha\rangle\langle\alpha|\,, \quad \int_{\mathbf{C}} d\alpha\, P(\alpha) = 1, \tag{3.24}$$

known as the Glauber-Sudarshan representation. This suggests that we can interpret an arbitrary state of one mode like a classical ensemble of coherent states with probability distribution $P(\alpha)$, except that $P(\alpha)$ is not always positive, and thus not always an actual distribution! We will discuss this much more later in class when we talk about "non-classical" states.

of photons (it is not a number eigenstate), but the average number of photons in a coherent state is

$$\overline{n} = \langle\alpha|n|\alpha\rangle = |\alpha|^2 \tag{3.25}$$

while the variance is

$$\Delta n^2 = \langle\alpha|(n-\overline{n})^2|\alpha\rangle = |\alpha|^2 = \overline{n}. \tag{3.26}$$

This means that the probability distribution of the photon number is Poisson ($\overline{n} = \Delta n^2$), and importantly, $\Delta n/\overline{n} \sim 1/\sqrt{\overline{n}} \to 0$ as $|\alpha|^2 = \overline{n} \to \infty$. Thus a large-amplitude coherent state has vanishingly small number fluctuations. Another "classical"-like behavior is that the expectation values of the quadratures obey classical equations in these states. For example, one can calculate easily using Eqs. (3.17) and (3.22) that

$$\langle\alpha|X(t)|\alpha\rangle = \frac{1}{\sqrt{2}}|\alpha|\cos(\omega t + \phi), \tag{3.27}$$

where $\alpha = |\alpha|e^{i\phi}$. For a mechanical oscillator for example, this says that the position x on average moves like a classical harmonic oscillator.

Some properties of these bases are collected in Appendix A. Let us now consider actual measurements of the field in these bases and how they are implemented.

In some sense the most elementary basis for measurements is the number basis, i.e., the energy basis. In the EM case, this corresponds to measurement of the energy content of an EM mode, which in its fundamental limit means counting up the number of photons. Another common terminology is that these are "intensity" or "square law" measurements, because they are sensitive to the energy, which is quadratic in the EM field $\sim \mathbf{E}^2 \sim \omega n$. This is distinguished from quadrature measurements which are linear in the field.

For now, we will make a simple model of an intensity detector, essentially following the discussion in Sec. 1. There, we considered a toy model of a detector based on the photoelectric effect, where we use excitations of electrons from a ground state $|0\rangle_M$ into a conduction band $|e\rangle_M$ with an energy gap tuned to the frequency of incoming light, $\Delta_{e0} = \omega$. Here we can just extend this by letting the detector have some large number $N \gg 1$ of electrons so that it can detect more than one photon. Let $|n\rangle_M$ denote the state of the detector with n excited electrons, and model the interaction between the EM mode and these electrons as perfectly absorbing:

$$U\,|n\rangle_S\,|0\rangle_M = |0\rangle_S\,|n\rangle_M\,. \tag{3.28}$$

Here S is the electromagnetic mode. The state $|0\rangle_S$ is the vacuum state $\hat{a}\,|0\rangle_S = 0$,

i.e., the state with exactly zero photons. The Kraus operators are simple:

$$K_n = \langle n|U|0\rangle_M = |0\rangle\,|n\rangle_S\,. \tag{3.29}$$

This n labels the number of clicks seen in the electron count. The meaning of this Kraus operator is clear: it takes in the state of n photons and outputs the vacuum, meaning, the photons are all absorbed and the post-measurement state is the vacuum $|\psi\rangle_S \to K_n\,|\psi\rangle_S\,/\sqrt{P(n)} = |0\rangle_S$. The POVM is

$$E_n = K_n^\dagger K_n = |n\rangle\,\langle n|_S\,, \tag{3.30}$$

so this enacts a projective, destructive measurement of the number of photons.

Measurements in the quadrature bases are slightly more subtle. The simplest way to implement them (in the EM field) is by doing what is called a homodyne measurement. This consists of first interfering light of known amplitude and phase, for example the output of a laser or maser, with the mode of interest, and then performing an intensity/number measurement on the resulting light. See Fig. 5.

To model this, we need to build up a key ingredient: the action of a beamsplitter.[4] For our purposes, we can just describe this by its action on coherent states. Let a, b be the annihilation operators of the two input modes to the beamsplitter. In general, the beamsplitter can be modeled as a unitary which implements the following transformation on a pair of incoming coherent states:

$$U_{\rm BS}\,|\alpha\rangle\,|\beta\rangle = |\alpha'\rangle\,|\beta'\rangle \tag{3.31}$$

with

$$\begin{aligned}\alpha' &= e^{-i\phi_1/2}(\alpha\cos\phi_2 e^{-i\phi_3/2} - \beta\sin\phi_2 e^{i\phi_3/2})\\ \beta' &= e^{-i\phi_1/2}(\alpha\sin\phi_2 e^{-i\phi_3/2} + \beta\cos\phi_2 e^{i\phi_3/2}),\end{aligned} \tag{3.32}$$

where $\phi_{1,2,3}$ are three arbitrary angles (corresponding to the three generators of $SU(2)$, which is unimportant for us [**?**]). For example, a 50-50 beamsplitter has $\phi_1 = \phi_3 = 0$ and $\phi_2 = \pi/4$, and outputs

$$\alpha' = (\alpha - \beta)/\sqrt{2}, \quad \beta' = (\alpha + \beta)/\sqrt{2}. \tag{3.33}$$

Another useful example that we will use here is a mostly-transmissive mirror, where $\phi_1 = \phi_3 = 0$ but $\phi_2 = \epsilon \ll 0$, which gives

$$\alpha' \approx (1-\epsilon^2)\alpha - \epsilon\beta, \quad \beta' \approx \epsilon\alpha + (1-\epsilon^2)\beta \tag{3.34}$$

[4]The discussion here is based loosely on a combination of Mandel and Wolf, Chapter 21, and a paper of Tyc and Sanders arXiv:quant-ph/0404090.

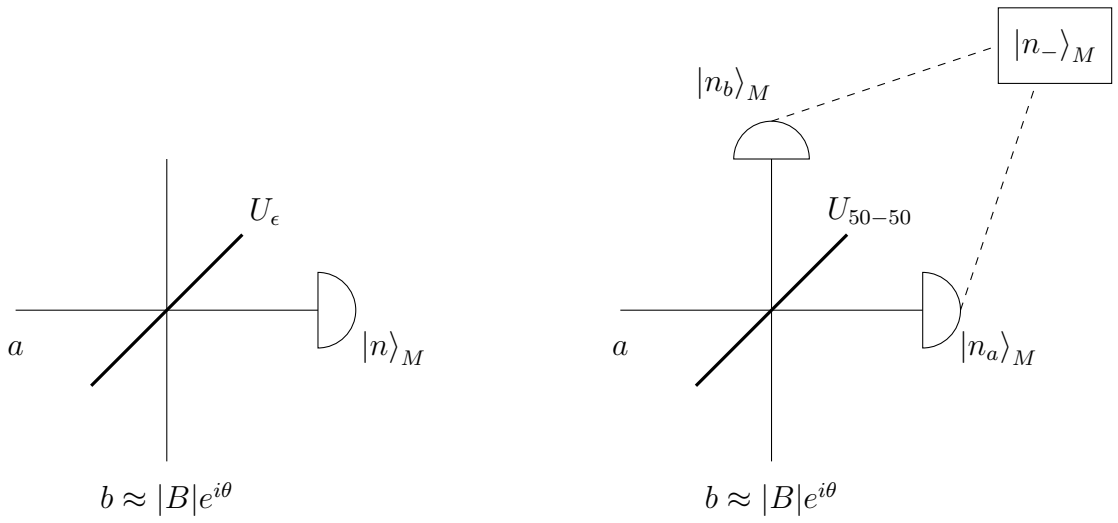

Figure 5: **Homodyne measurements of an electromagnetic field mode.** Here we show two common architectures for homodyne measurements. In both cases, we start with both the mode of interest in some unknown state $|\psi\rangle_S$, and another mode of the same frequency prepared in a coherent state $|B\rangle$ with large amplitude $|B|^2 = \langle n_{\rm in}\rangle \gg 1$ and known phase $B = |B|e^{i\theta}$. These modes are then sent through a beamsplitter, which produces a new beam consisting of the superposition of $|\psi\rangle_S$ and $|B\rangle$, i.e., encoding an interference pattern. This intefered beam is then measured with intensity measurements. Left: unbalanced homodyne, where the signal is extracted from a single intensity measurement after being sent through a mostly-transmissive beamsplitter U_ϵ with reflectivity $\epsilon \ll 1$. Right: balanced homodyne, where the signal is extracted from a differential measurement of two intensity detectors after a 50-50 beamsplitter.

Since this is the action of $U_{\rm BS}$ on coherent states, and the coherent states are eigenvectors of the annihilation operators $a\,|\alpha\rangle = \alpha\,|\alpha\rangle$, $b\,|\beta\rangle = \beta\,|\beta\rangle$, it follows that the Heisenberg-picture version of the transformation acts in similar fashion:

$$\begin{aligned} a \to a' = U_{\rm BS}^\dagger a U_{\rm BS} &= e^{-i\phi_1/2}(a\cos\phi_2 e^{-i\phi_3/2} - b\sin\phi_2 e^{i\phi_3/2}) \\ b \to b' = U_{\rm BS}^\dagger b U_{\rm BS} &= e^{-i\phi_1/2}(a\sin\phi_2 e^{-i\phi_3/2} + b\cos\phi_2 e^{i\phi_3/2}). \end{aligned} \tag{3.35}$$

One can give an explicit form of $U_{\rm BS}$ in terms of an $SU(2)$ matrix, see **??**, but this will not be needed here.

Now let's analyze the unbalanced homodyne measurement (see Fig. 5, left). In homework you will analyze the balanced version, which is more practical. The essential idea is that we interfere the state of interest $|\psi\rangle_S$, which is the mode a, with

a large-amplitude coherent state $|B\rangle$ with known phase $B = |B|e^{i\theta}$. We also assume that the reflectivity ϵ is small, but not so small, so that the average number of B photons which is reflected and mixed with $|\psi\rangle_S$ is $n_{\rm LO} = \epsilon^2|B|^2 \gg 1$. Here the LO notation stands for "local oscillator". Since this field is a large-amplitude coherent state, we can approximate its mode operator as just a classical c-number. In particular, the a' mode after the beamsplitter has number operator

$$\begin{aligned} n_a \to n'_a &= a'^\dagger a' \\ &= n_a(1-\epsilon^2)^2 a^\dagger a - \epsilon(ab^\dagger + a^\dagger b) + \epsilon^2 b^\dagger b \\ &\approx n_a(1-2\epsilon^2) + n_{\rm LO} - \epsilon|B|(ae^{i\theta} + a^\dagger e^{-i\theta}), \end{aligned} \tag{3.36}$$

up to terms of order ϵ^3. Notice that this last term in parentheses is the quadrature operator X_θ. So by interfering the input light with a LO of known phase θ, we are imprinting the quadrature information onto the number operator. Let's assume that n_a itself (the number of photons in the signal) is small compared to $n_{\rm LO}$. This means that we can drop this term and approximate

$$n'_a \approx n_{\rm LO} - \sqrt{n_{\rm LO}} X_\theta \implies X_\theta = \frac{n'_a - n_{\rm LO}}{\sqrt{n_{\rm LO}}}. \tag{3.37}$$

That is, we can invert the relation: the value of X_θ is given by the number of photons n'_a in the beam after the beamsplitter, with $X_\theta = 0$ corresponding to $n'_a = n_{\rm LO}$. Notice that Eq. (3.37) is an operator relation between the number operator and quadrature operator.

Let's make the measurement itself precise, by working out Kraus operators and the POVM corresponding to this. The system (S = mode a) first goes through the beamsplitter and is interfered, and then interacts with the measurement apparatus (M = photodetector). Notice that $U_{\rm BS} = U_{\rm BS} \otimes 1_M$ does not act on photodetector M at all, only on S. So, the whole Kraus operator is

$$K_n = \langle n|U|0\rangle_M = \langle n|U_{\gamma-pd}U_{\rm BS}|0\rangle_M = \langle n|U_{\gamma-pd}|0\rangle_M U_{\rm BS} = |0\rangle\langle n|_S U_{\rm BS}. \tag{3.38}$$

Here $U_{\gamma-pd}$ is the unitary describing absorption of the photons by the photodetector, Eq. (3.28). The label n on the Kraus operator has the same meaning as before: the measurement outcome is that *n electrons are excited in the photodetector*. Now, however, that outcome gives us very different information about S, because of the action of $U_{\rm BS}$. Based on the argument in the previous paragraph, one might expect this Kraus operator to look like some kind of quadrature-basis object. Indeed, one

can show[5] that

$$U_{\mathrm{BS}}^{\dagger} \left|n\right\rangle_S \approx \left|X_\theta^n\right\rangle_S , \tag{3.40}$$

where $\left|X_\theta^n\right\rangle$ is a *quadrature* eigenstate $\hat{X}_\theta \left|X_\theta^n\right\rangle = X_\theta^n \left|X_\theta^n\right\rangle$, with eigenvalue

$$X_\theta^n = \frac{n - n_{\mathrm{LO}}}{\sqrt{n_{\mathrm{LO}}}}. \tag{3.41}$$

The Kraus operator is thus

$$K_n = \left|0\right\rangle \left\langle X_\theta^n\right|_S \tag{3.42}$$

with corresponding POVM

$$E_n = \left|X_\theta^n\right\rangle \left\langle X_\theta^n\right|_S . \tag{3.43}$$

This means that the total measurement acts as a projective measurement *in the* θ *quadrature* basis. The post-measurement state of the photon is still the vacuum, since the photons are absorbed by the photodetector.

We can get some intuition for this set of formulas by working through a simple example. Consider an arbitrary pure state of the mode, which we can express in the amplitude quadrature basis as

$$\left|\psi\right\rangle_S = \int_{-\infty}^{\infty} dX \, \psi(X) \left|X\right\rangle_S \tag{3.44}$$

for some wavefunction $\psi(X)$. Tuning the local oscillator to $\theta = 0$ so that we measure in the amplitude quadrature basis, the probability that we see n electron excitations in the photodetector is just

$$P(n) = \left\langle\psi\right|E_n\left|\psi\right\rangle_S = \left|\psi(X^n)\right|^2. \tag{3.45}$$

For example, consider preparing a near-eigenstate of the amplitude operator, $\left|\psi(X)\right|^2 \approx \delta(X - X_0)$. Then we get a distribution of clicks on the detector

$$P(n) \approx \delta(X^n - X_0) \tag{3.46}$$

which is a sharply peaked distribution centered at $n = n_{\mathrm{LO}} + \sqrt{n_{\mathrm{LO}}} X_0$. So, when you look at the number of clicks in the photodiode, you can immediately read off the value of X_0 by just seeing how far the click distribution is from $n = n_{\mathrm{LO}}$. With a general state, Eq. (3.45) says that the click distribution will essentially reproduce a discretized image of $\left|\psi(X)\right|^2$, shifted and scaled so that $X = 0$ is centered at $n = n_{\mathrm{LO}}$.

[5]Here is the proof. Let $U = U_{\mathrm{BS}}$. We have that $U^{\dagger} n U = n'$. Thus, an eigenstate $\left|n\right\rangle$ of n transforms into an eigenstate of n' with the same eigenvalue:

$$n' U^{\dagger} \left|n\right\rangle = U^{\dagger} n U U^{\dagger} \left|n\right\rangle = n U^{\dagger} \left|n\right\rangle . \tag{3.39}$$

Then according to Eq. (3.37), $U^{\dagger} \left|n\right\rangle$ is approximately an eigenstate of X_θ, with eigenvalue $X_\theta = X_\theta^n$.

Problems

1. **Minimal uncertainty in successive measurements.** Consider measurements of a particle in one spatial dimension. Assume that you can make position measurements described by a Gaussian POVM with uncertainty Δx_M (i.e., variance Δx_M^2), and momentum measurements described by a Gaussian POVM with uncertainty Δp_M. [This problem is based on a paper of Distler and Paban arXiv:1211.4169]

   (a) Consider an initial Gaussian state $\psi_S(x) = Ne^{-x^2/4\Delta x_S^2}$. For general Δx_S and POVM widths, calculate the uncertainty in the outcomes in a position measurement $(\Delta x)^2_{\rm meas}$.

   (b) Calculate the post-measurement state after this position measurement, in particular its width.

   (c) Now perform a momentum measurement; calculate the uncertainty in the outcomes $(\Delta p')^2_{\rm meas}$, again for general values of the widths.

   (d) Prove that the product of the two measurement uncertainties satisfies the bound
   $$(\Delta x)^2_{\rm meas}(\Delta p)^2_{\rm meas} \geq \frac{1}{4}\left(1+\sqrt{1+4\Delta x_M^2\Delta p_M^2}\right)^2. \qquad (3.47)$$

   (e) Find the initial state $\psi_S(x)$ that saturates this inequality (i.e., its width).

2. **Balanced homodyne detection.** Consider the balanced homodyne detector shown in Fig. 5, right. Here the signal beam a is mixed with a local oscillator $B = |B|e^{i\theta}$ with large occupation $|B|^2 \gg 1$ as in the unbalanced case. The difference is that we mix these beams with a 50-50 beamsplitter, and then measure the differential photocurrent.

   (a) Let a,b denote the annihilation operators for the input beams and a',b' the output beams. Work out the expressions for the number operators in the beams after the beamsplitter, n'_a and n'_b, in the notation of the main text.

   (b) Now approximate $b \to |B|e^{i\theta}$ as a classical external field. Calculate the effective number operators n'_a and n'_b, and their difference $n'_- = n'_a - n'_b$. Give n_- in terms of X_θ. Make the same approximations as in the main text.

   (c) Suppose we measure n_-. Calculate the Kraus operators and POVM elements for the outcomes $n_- = -\infty, \ldots, \infty$, in terms of the X_θ basis.

(d) Put your experimentalist hat on and give a clear technical reason that the balanced homodyne is superior to an unbalanced one (think about laser noise).

4. Tomography of quantum states

In the previous chapters, we studied what it means to do a measurement of a quantum system. This was taken as something of obvious interest. A measurement is defined to mean that we observe a system and obtain some outcomes with some probability distribution. We now begin our study of various tasks that can be accomplished with this information.

A fundamental task is the following: perform a series of measurements on a system prepared in an unknown quantum state ρ, and use the outcomes of these measurements to determine ρ. This problem is called "state tomography": we want to take an "image" (a tomograph) of the state. This is, in some sense, the most primitive or fundamental problem in measurement theory. See Fig. 6.

To formalize the problem, consider the following scenario: we have some black-box system with a button on it, and every time we hit the button, the box spits out a system S prepared in a state ρ. The state is the same every time we hit the button, but can in general be a mixed state. For example, the black box could be Stern-Gerlach apparatus that prepares $|0\rangle + |1\rangle$ of a spin, or it could be a lightbulb that prepares ρ_T the thermal state of some electromagnetic field modes. Our problem is to design a measurement scheme (i.e., an appatus M, coupling $U = U_{SM}$, and basis $|a\rangle_M$, or equivalently a POVM $\{E_a\}$ on $\mathcal{H}_S$) and an algorithm which converts the measurement data (the distribution of outcomes $P(a)$) into an estimate $\tilde{\rho}$ of ρ.

We begin with a few examples on simple systems, and then state a fundamentally important fact: tomography is always possible with *local measurements*, even on an arbitrarily large and complex multipartite system! This surprising statement ("axiom of local tomography") is given in general, and then demonstrated with a constructive example on n qubits known as Pauli tomography (an algorithm which is commonly used in practice). We then discuss some inefficiencies with this protocol and use them to motivate the study of *informationally-complete measurements* (IC-POVMs), and show how these can be used to perform more efficient state reconstructions. Finally we show how the tomography problem works in continuous variable settings, particularly homodyne tomography, which historically was the origin for the idea of tomography.

hello

Figure 6: **The state tomography problem.** Left: ...

4.1. Examples on a single degree of freedom

To highlight the basic procedures, we begin by describing concrete tomography algorithms on a single qubit and a single harmonic oscillator. The latter, in particular, will also give us an excuse to discuss "phase-space" descriptions of quantum oscillators, such as the Wigner representation.

4.1.1. Pauli tomography on one qubit

Consider an arbitrary state ρ of a single qubit. This is a 2×2 complex matrix

$$\rho = \begin{pmatrix} \rho_{00} & \rho_{01} \\ \rho_{10} & \rho_{11} \end{pmatrix} \tag{4.1}$$

subject to the constraints that it is Hermitian $\rho^\dagger = \rho$ and normalized $\operatorname{tr} \rho = 1$. This means that it has $8 - 4 - 1 = 3$ real parameters. There is a convenient basis for such matrices, formed by the Pauli matrices $\sigma_{x,y,z}$ and the identity 1, which concretely means that we can express a general qubit state as

$$\rho = \frac{1}{2} \sum_{i=0}^{3} S_i \sigma_i = \frac{1}{2} \begin{pmatrix} 1 + S_3 & S_1 - iS_2 \\ S_1 + iS_2 & 1 - S_3 \end{pmatrix} \tag{4.2}$$

where

$$\sigma_i = (1, \sigma_x, \sigma_y, \sigma_z), \quad S_i = \operatorname{tr} \sigma_i \rho. \tag{4.3}$$

Note that $S_0 = 1$ by definition, so the other three $S_{1,2,3}$ are a set of three real numbers which can parametrize the space of states ρ. These matrices in fact define an orthonormal basis, where we define the inner product between two matrices A, B as

$$\langle A, B \rangle := \operatorname{tr} A^\dagger B \tag{4.4}$$

in which sense we have

$$\langle \sigma_i, \sigma_j \rangle = 2\delta_{ij} \tag{4.5}$$

which follows directly from their anticommutation relations. We will make repeated use of this inner product on matrices in this chapter. The 2 here is the reason for the $1/2$ in Eq. (4.2).

Our goal now is to find a series of measurements on the qubit from which we can reconstruct ρ.[6] The reason for introducing the parameterization (4.2) is that

[6]The description in this section is based on [**?**], rewritten into a language that we can generalize to Sec. **??**, which in turn is based largely on conversations with Jacob Beckey and the thesis of Angus Lowe [**?**].

this is particularly convenient for this task: we need a series of measurements $\{E_a\}$ from which the measured output $\{P(a)\}$ can be used to infer the numbers $S_{1,2,3}$. Let $|\pm x\rangle = (|0\rangle \pm |1\rangle)/\sqrt{2}$ be the eigenstates of σ_x, and then notice that

$$\begin{aligned} S_1 &= \operatorname{tr} \sigma_x \rho \\ &= \langle +x|\rho|+x\rangle - \langle -x|\rho|-x\rangle \\ &= \operatorname{tr}\left[|+x\rangle\langle +x|\,\rho\right] - \operatorname{tr}\left[|-x\rangle\langle -x|\,\rho\right] \\ &= P(+x) - P(-x), \end{aligned} \tag{4.6}$$

where we introduced the simple, projective, two-outcome POVM with elements $E_{x,+} = |+x\rangle\langle +x|$, $E_{x,-} = |-x\rangle\langle -x|$. By identical manipulations one finds that

$$S_2 = P(+y) - P(-y), \quad S_3 = P(+z) - P(-z) \tag{4.7}$$

where the two, two-outcome POVMs here are $\{E_{y,\pm} = |\pm y\rangle\langle \pm y|\}$ and $\{E_{\pm z} = |z,\pm\rangle\langle \pm z|\}$ in terms of the $\sigma_{y,z}$ eigenstates. What these equations show us is that we can find a series of three POVMs, one each for x, y, z, whose measured outcomes [the set of six real numbers $P(+x), \ldots, P(-z)$] are sufficient to reconstruct the density matrix coefficients $S_{1,2,3}$. For later use, notice that we can write these POVM elements as

$$E_{i,\pm} = \frac{1 \pm \sigma_i}{2}, \tag{4.8}$$

and label the outcomes with ± 1 (the eigenvalue of the appropriate σ_i). This can be seen by explicit calculation with the σ_i eigenstates.

The tomography algorithm then proceeds as follows. We prepare a copy of ρ, then measure it with the POVM $E_{x,\pm}$ (i.e., by doing a Stern-Gerlach experiment, since $E_{x,\pm} = |\pm x\rangle\langle \pm x|$ is just a projective measurement). Doing this M times, we get a list of data points like

$$x_m = (1, -1, -1, 1, \ldots), \quad m = 1, \ldots, M \tag{4.9}$$

from which we can estimate the probabilities

$$\tilde{P}(+x) = \frac{\#\ +1'\text{s}}{M}, \quad \tilde{P}(-x) = \frac{\#\ -1'\text{s}}{M}, \tag{4.10}$$

and likewise form an estimator for S_1:

$$\tilde{S}_1 = \tilde{P}(+x) - \tilde{P}(-x) = \frac{1}{M}\sum_{m=1}^{M} x_m. \tag{4.11}$$

The tildes here represent the fact that these are sampled from a finite number of outcomes; their average over many measurements must converge to the actual quantum distribution:

$$\mathbb{E}\left[\tilde{P}(\pm x)\right] = P(\pm x) = \operatorname{tr} E_{x,\pm}\rho, \quad \mathbb{E}\left[\tilde{S}_x\right] = S_x = \operatorname{tr} \sigma_x \rho. \tag{4.12}$$

However, with a finite number of samples M, the inferred probabilities have a standard error. Since the outcomes are binary, it is easy to estimate this error. The expected variance of S_1 satisfies

$$(\Delta S_1)^2 \leq 1, \tag{4.13}$$

since x_m forms a binary measurement with outcomes ± 1. This means that the standard error

$$\Delta S_1 = \sqrt{\frac{(\Delta S_1)^2}{M}} \leq \frac{1}{\sqrt{M}}, \tag{4.14}$$

the usual result for a binary-outcome measurement. So, the result of these M measurements is an estimate $\tilde{S}_1$ for the density matrix coefficient S_1, which we know has an error no worse than $1/\sqrt{M}$.

Note that this error is not due to imperfections in the measurement, noise, or anything else "practical". It is a fundamental quantum mechanical effect that comes from the fact that the measurement outcomes are random. In the limit $M \to \infty$ these errors vanish and we can perfectly reconstruct S_1, but with a finite number of measurements M, the randomness of quantum mechanical measurement outcomes only allows us to reconstruct the original density matrix up to some error.

To finish stating the algorithm, we next prepare another copy ρ of the state, but now measure in the $E_{y,\pm}$ basis. From this, we similarly infer an estimate $\tilde{S}_2$, with the same upper bound on its error, by doing *another* M measurements in this basis. We then repeat the process for $E_{z,\pm}$, obtaining $\tilde{S}_3$. In the end, we obtain an estimator

$$\tilde{\rho} = \frac{1}{2}\sum_{i=0}^{3} \tilde{S}_i \sigma_i = \frac{1}{2}\begin{pmatrix} 1+\tilde{S}_3 & \tilde{S}_1 - i\tilde{S}_2 \\ \tilde{S}_1 + i\tilde{S}_2 & 1-\tilde{S}_3 \end{pmatrix}. \tag{4.15}$$

Note that we had to do $3M$ total measurements here, in 3 different bases (i.e., with 3 different POVMs). We will discuss these resource requirements more in the multipartite case.

How good is the estimator? By construction, we know that this estimator is unbiased, in the sense that averaging over many measurement outcomes we will converge to the right answer

$$\mathbb{E}\left[\tilde{\rho}\right] = \rho. \tag{4.16}$$

How fast will it converge? We know each of the S_i with precision $1/\sqrt{M}$, where M is the number of measurements *of each POVM*. To characterize the total error in the density matrix, a number of options are possible. A commonly used definition is called the sample complexity, where the terminology refers to the number of samples ($N = 3M$, in this case) that have to be prepared. To define the sample complexity, we use our inner product on matrices, Eq. (4.4). This inner product can be used to define a few useful norms. One is

$$||A||_F := \sqrt{\text{tr}(A^\dagger A)} \tag{4.17}$$

where the F refers to Frobenius, who was some math guy. Another useful one is the so-called 1-norm

$$||A||_1 := \text{tr}\,\sqrt{A^\dagger A}. \tag{4.18}$$

As you can imagine, the Frobenius norm is a lot easier to deal with in practice since you don't have to find the square root of a matrix, but the 1-norm is more statistically relevant because it generalizes the usual Kolmogorov notion of distance between two probability measures. That is, we define the "trace distance" between two states as

$$d_1(\rho, \rho') := ||\rho - \rho'||_1 = \text{tr}\,\sqrt{(\rho - \rho')^\dagger(\rho - \rho')}. \tag{4.19}$$

In practice, it is most useful to calculate distances with the Frobenius norm, and use this to bound the trace distance, by the inequality

$$||A||_1 \leq \sqrt{d||A||_2^2}, \tag{4.20}$$

for a $d \times d$ matrix A. Finally, we can define the sample complexity: it is the number of total measurements N such that our reconstructed density matrix is sufficiently close to the true state:

$$d_1(\rho, \tilde{\rho}) \lesssim \epsilon. \tag{4.21}$$

The approximation here is so that we can sweep the truly precise definition under the rug, which says that we demand this inequality strictly hold with probability $1 - \delta$; here we will just make order-of-magnitude arguments.

Now let's consider the sample complexity in the specific tomographic protocol we just defined, for one qubit. Let $\epsilon_{ij} = \rho_{ij} - \tilde{\rho}_{ij} \sim 1/\sqrt{M}$ be the error on each density matrix element after M measurements of each Pauli operator. These are random numbers of order $1/\sqrt{M}$ but can be positive or negative. We have

$$||\rho - \tilde{\rho}||_F^2 = \text{tr}\left[\begin{pmatrix} \epsilon_{00} & \epsilon_{01} \\ \epsilon_{10} & \epsilon_{11} \end{pmatrix}\begin{pmatrix} \epsilon_{00} & \epsilon_{01} \\ \epsilon_{10} & \epsilon_{11} \end{pmatrix}\right] = \epsilon_{00}^2 + 2\epsilon_{01}\epsilon_{10} + \epsilon_{11}^2 \lesssim \frac{2}{M}. \tag{4.22}$$

The estimate here is that the two squared terms always contribute positively like $1/M$ while the middle term can be positive or negative. We could write this somewhat prosaically as d/M, since here $d = 2$. We will see that this accounting generalizes to d dimensions when we do an n-qubit version of this, in which case $d = 2^n$. Using the bound on the trace distance Eq. (4.20), we get

$$d_1(\rho, \tilde{\rho}) \leq \sqrt{2||\rho - \tilde{\rho}||_2^2} \approx 2/\sqrt{M}. \tag{4.23}$$

If we want our estimate to be within ϵ in this distance, this requires

$$M \gtrsim \frac{4}{\epsilon^2} \implies N_{\text{samples}} = 3M \gtrsim \frac{3 \times 4}{\epsilon^2}. \tag{4.24}$$

We will see later that this generalizes to $N_{\text{samples}} \gtrsim d^4/\epsilon^2$ in higher dimensions. The factor of 3 entered because we have to measure with 3 different POVMs (x, y, z bases); this means that we have to set up the experiment with $N_{\text{settings}} = 3$ different configurations. For one qubit, this says that to get an estimate of the density matrix accurate to 1% requires around $N_{\text{samples}} \approx 10^5$ measurements. The requirements quickly become very challenging once we have $n \gg 1$ qubits, in which case $N_{\text{samples}} \sim 2^{4n}$. We discuss this in detail in Sec. 4.2.1.

4.1.2. Wigner representations; homodyne tomography of an oscillator

Tomography for a continuous-variable system is somewhat different, since $\dim \mathcal{H} = \infty$. In this section, we consider tomography of a single harmonic oscillator, using the homodyne measurements introduced in Sec. 3.3. It turns out that measuring quadrature variables at different angles is sufficient to reconstruct an oscillator state. The method we will describe here is a workhorse in many disparate systems, including optical fields, microwave fields, and the mechanical motion of trapped atoms and ions.

Much like the qubit example of the previous section, it turns out to be very useful to first pick a representation of an arbitrary oscillator state which is easily tied to easily implemented measurements. This is analogous to the Pauli decomposition of Eq. (4.2). Of course, we could expand an arbitrary state ρ in, say, the position or number bases:

$$\rho = \sum_{n,n'} \rho_{nn'} |n\rangle \langle n'| = \int_{-\infty}^{\infty} dx dx' \, \rho(x, x') |x\rangle \langle x'| . \tag{4.25}$$

However, much like the qubit, there turns out to be a much more convenient representation. Consider an arbitrary oscillator state ρ and define the function

$$W(x, p) = \frac{1}{2\pi} \int_{-\infty}^{\infty} dy \, e^{ipy} \left\langle x - \frac{y}{2} \middle| \rho \middle| x + \frac{y}{2} \right\rangle . \tag{4.26}$$

This function is known as the Wigner distribution or Wigner function. The kets in the integral are position-space eigenstates. This integral implements a Radon transformation, which is a transformation which takes a 2d density function [here, the position-space density matrix elements $\rho(x, x')$], performs the Fourier transformation over a 1d line in this plane, and outputs a number. This kind of transformation is the basis for real tomography (e.g., in medical imaging) and is the historical origin of the term tomography. We will soon show that Eq. (4.26) is invertible, which means that an arbitrary state ρ can be equivalently described by a Wigner function W.

The Wigner function $W(x, p)$ is very similar to a probability distribution on phase space. To emphasize the connection, we will mostly label the coordinates (x, p), rather than (X, Y) as in Sec. 3.3, but this is just notation. Let's discuss some basic properties of this function. Most importantly, it can be used to compute probabilities in the state ρ. For example, we have

$$\begin{aligned}\int_{-\infty}^{\infty} dp\, W(x,p) &= \frac{1}{2\pi}\int_{-\infty}^{\infty} dp dy\, e^{ipy}\left\langle x-\frac{y}{2}\right|\rho\left|x+\frac{y}{2}\right\rangle \\ &= \int_{-\infty}^{\infty} dy\,\delta(y)\left\langle x-\frac{y}{2}\right|\rho\left|x+\frac{y}{2}\right\rangle \\ &= \langle x|\rho|x\rangle \\ &= P(x),\end{aligned} \tag{4.27}$$

which shows that $P(x)$ can be calculated as a partial integral over $W(x, p)$, sometimes called a marginal (in the sense of marginal distribution). A similar but slightly more involved calculation shows that we can do the same to get $P(p)$, the distribution of momentum measurement outcomes:

$$\begin{aligned}\int_{-\infty}^{\infty} dx\, W(x,p) &= \frac{1}{2\pi}\int_{-\infty}^{\infty} dx dy\, e^{ipy}\left\langle x-\frac{y}{2}\right|\rho\left|x+\frac{y}{2}\right\rangle \\ &= \frac{1}{2\pi}\int_{-\infty}^{\infty} dx dy dp' dp''\, e^{ipy}\left\langle x-\frac{y}{2}\right|p'\left\rangle\langle p'|\rho|p''\rangle\left\langle p''\right|x+\frac{y}{2}\right\rangle \\ &= \frac{1}{(2\pi)^2}\int_{-\infty}^{\infty} dx dy dp' dp''\, e^{i(p-p'/2-p''/2)y}e^{ix(p'-p'')}\langle p'|\rho|p''\rangle \\ &= \frac{1}{2\pi}\int_{-\infty}^{\infty} dp' dp''\,\delta(p-p'/2-p''/2)\delta(p'-p'')\langle p'|\rho|p''\rangle \\ &= \langle p|\rho|p\rangle \\ &= P(p).\end{aligned} \tag{4.28}$$

To get the second line we inserted two factors of the identity $1 = \int dp\ |p\rangle\langle p|$, and to get the third we used $\langle x|p\rangle = e^{ipx}/\sqrt{2\pi}$. A very similar calculation shows that, in

general, the probability distribution of any quadrature X_θ [see Eq. (3.18)] is given by

$$P(X,\theta) := \langle X_\theta|\rho|X_\theta\rangle = \int_{-\infty}^{\infty} dP\, W(X\cos\theta - P\sin\theta, X\sin\theta + P\cos\theta). \qquad (4.29)$$

It is this connection to homodyne distributions that will enable a straightforward tomography algorithm: essentially, we just invert this to get an expression for W as a function of (measurable) homodyne data.

Historically, Wigner actually derived Eq. (4.26) by demanding that W obeys Eq. (4.29), as well as a normalization condition:

$$\int_{-\infty}^{\infty} dx dp\, W(x,p) = 1. \qquad (4.30)$$

Since the function W is normalized like a probability distribution, and can be used to compute the various marginal distributions $P(X,\theta)$, it is tempting to think of it as an actual probability distribution on phase space (x,p). However, it turns out that $W(x,p)$ can take *negative values*, which is obviously inconsistent with being a probability. Thus sometimes one sees terms like "quasiprobability distribution" used for W. More interesting than the terminology is the physical interpretation of $W < 0$ (called "Wigner negativity"). We will see here and again later in our discussion of "non-classical" light in Ch. **??** that Wigner negativity is a hallmark sign of quantum behavior, meaning behavior that cannot arise in a classical description. Here we will state this somewhat loosely; in Ch. **??** we will give much more precise, operational characterizations of this non-classicality.

Before giving the tomography algorithm, let's study a few examples of Wigner functions to get some intuition.

- The vacuum $|\psi\rangle = |0\rangle$, where $a|0\rangle = 0$. An easy calculation shows that W is a Gaussian $W(x,p) \sim e^{-x^2/2\Delta x_0^2} e^{-p^2/2\Delta p_0^2}$, where $\Delta x_0^2 = 1/2m\omega$ and $\Delta p_0^2 = m\omega/2$ are the ground-state oscillator uncertainties. Here $W(x,p) > 0$ everywhere, so in this sense the vacuum is "classical"; again, we will have more to say about this when we discuss non-classical states.

- The canonical thermal state $\rho = \rho_T = e^{-H/T}/Z$, with $H = \omega a^\dagger a$ the Hamiltonian and $Z = \mathrm{tr}\, e^{-H/T}$ the partition function. This turns out to be another Gaussian $W(x,p) \sim e^{-x^2/2\Delta x_T^2} e^{-p^2/2\Delta p_T^2}$, except now the uncertainties are the usual thermal values, like $\Delta x_T^2 = \coth(\omega/2T)/2m\omega$. As $T \to 0$ this reproduces the vacuum state while as $T \to \infty$ the Gaussian becomes a flat distribution, as one would expect. Again $W(x,p) > 0$ everywhere.

- A coherent state $|\alpha\rangle$ [see Eq. (3.22)]. Once again the state is a Gaussian, but now non-trivially centered: $W(x,p) \sim e^{-(x-x_\alpha)^2/2\Delta x_0^2} e^{-(p-p_\alpha)^2/2\Delta p_0^2}$. The uncertainties are the same as those in the vacuum, but the center values $x_\alpha = \text{Re } \alpha$ and $p_\alpha = \text{Im } \alpha$. Yet again, $W(x,p) > 0$ everywhere.

- Finally, an example with Wigner negativity: a "cat state" $|\psi\rangle \sim |\alpha\rangle + |-\alpha\rangle$ made up of the superposition of two coherent states with opposing center values. If $|\alpha|$ is not too large, then these two states have a non-trivial overlap since $|\langle\alpha| - \alpha\rangle|^2 = e^{-|\alpha|^2/2}$ [see Eq. (3.23)]. In position space, this overlap occurs around the origin $x = 0$, and it turns out that this leads to an interference pattern with Wigner negativity $W < 0$ in the overlap region. This is consistent with our heuristic statement about Wigner negativity and non-classical behavior. For a mechanical oscillator for example, this cat state represents the coherent superposition of two states where the oscillator is localized at disjoint locations, which is as non-classical an object as one can imagine in a single degree of freedom.

See Figs. **??** and **??** for some graphical representations.

Finally, let's describe the homodyne tomography algorithm. Here I will essentially follow the original prescription due to Risken and Vogel [], which was experimentally demonstrated in an optical mode almost immediately []. This method is somewhat naive and has poor convergence properties; the goal here is mostly to give a proof-of-principle construction that tomography is (approximately) possible even in infinite dimensions. For a thorough discussion of the changes that need to be made to make the algorithm really work in practice, see [].

The basic strategy is simple. First, we find the inverse of Eq. (4.26), which means a formula that gives the density matrix ρ in terms of its Wigner function $W(x,p)$. Second, we find a formula that gives the Wigner function $W(x,p)$ as a function of homodyne outcome distributions $P(X,\theta)$. Since this last item can be measured experimentally, we then obtain a formula for the state as a function of measurement outcomes $\rho = \rho[P(X,\theta)]$. With actual data, one then forms an estimator $\tilde{\rho} = \tilde{\rho}[\tilde{P}(X,\theta)]$, as desired.

The first step, inversion of Eq. (4.26), is easy. Just invert the Fourier transform:

$$\begin{aligned}\int_{-\infty}^{\infty} dp e^{-ip2\Delta x}\, W(x,p) &= \frac{1}{2\pi}\int_{-\infty}^{\infty} dp dy\, e^{ip(y-2\Delta x)} \left\langle x-\frac{y}{2}\,\middle|\,\rho\,\middle|\,x+\frac{y}{2}\right\rangle \\ &= \langle x-\Delta x|\rho|x+\Delta x\rangle \\ &= \rho(x-\Delta x, x+\Delta x).\end{aligned} \tag{4.31}$$

Since Δx is just an arbitrary parameter, this gives all the position-space matrix elements of the Wigner function, i.e., the density matrix itself, as a function of $W(x,p)$.

The second step, determining a formula that gives $W = W[P(X,\theta)]$, is slightly more involved. We define some Fourier transforms as follows:

$$\begin{aligned}P(X,\theta) &= \int_{-\infty}^{\infty} d\xi\, e^{i\xi X}\hat{P}(\xi,\theta) \\ W(x,p) &= \int_{-\infty}^{\infty} du dv\, e^{iux+ivp}\hat{W}(x,p).\end{aligned} \tag{4.32}$$

This second formula can be conveniently rewritten in a non-standard form of polar coordinates:

$$W(x,p) = \int_{-\infty}^{\infty} d\xi \int_0^{\pi} d\theta\, |\xi| e^{i\xi(x\cos\theta + p\sin\theta)}\hat{W}(\xi\cos\theta, \xi\sin\theta). \tag{4.33}$$

These are non-standard in the sense that we're parametrizing the plane in terms of rays, not radii, so for example the usual point $r=1, \theta = 3\pi/2$ is here $\xi = -1, \theta = \pi/2$. This parametrization is useful precisely because the Wigner function is a Radon transformation. Using these definitions and Eq. (4.29), it's straightforward to show $\hat{P}$ and $\hat{W}$ are related through the simple formula

$$\hat{P}(\xi,\theta) = \hat{W}(\xi\cos\theta, \xi\sin\theta). \tag{4.34}$$

This expresses the Wigner function W as a function of homodyne data $\hat{P}$, which is what we wanted. Explicitly, inserting (4.34) back into (4.33) and using the inverse Fourier transform of the first line in (4.32), we obtain

$$W(x,p) = \int_{-\infty}^{\infty} d\xi \int_0^{\pi} d\theta \int_{-\infty}^{\infty} dX\, |\xi| e^{i\xi(x\cos\theta + p\sin\theta - X)} P(X,\theta). \tag{4.35}$$

This gives us $W = W[P(X,\theta)]$, and then we can insert it into Eq. (4.31) to get $\rho = \rho[P(X,\theta)]$, which gives us a formula for the density matrix in terms of the homodyne data.

Consider actually doing homodyne tomography based on the inversion formula Eq. (4.35). This formula says that we need to prepare the state ρ, send it into a

homodyne detector with some fixed angle θ_1, and repeat this many times to form an estimate of the distribution $P(X, \theta_1)$. We then need to repeat this process at another angle θ_2, another θ_3, and so forth, and in principle we actually need to get the homodyne distribution *at every value of* θ. In the language of the Pauli example of the previous section, this requires $N_{\text{settings}} = \infty$ and therefore also $N_{\text{samples}} = \infty$. What this means of course is that one has to discretize the set of angles and perform some type of splining. A detailed analysis of the errors induced by this is beyond the scope of this course, but [] contains a thorough treatment.

4.2. Axiom of local tomography

In Sec. 1.3, we gave a precise definition of local measurements on a multipartite system $\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2 \otimes \cdots$, as a measurement whose POVM elements factor $E_{\mathbf{a}} = E_{a_1} \otimes E_{a_2} \otimes \cdots$. Less abstractly, this means a measurement that can be described by coupling a local measurement apparatus M_i to each system factor $\mathcal{H}_i$ through a single-site unitary U_i, and making a projective measurement on each M_i separately.

Armed with the appropriate definition of a local measurement, we can state an incredible fact about quantum mechanics:

> **Theorem ("Axiom of local tomography"):** For any state ρ in a Hilbert space $\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2 \otimes \cdots$, there always exists a set of local measurements $\{E_{\mathbf{a}} = E_{a_1} \otimes E_{a_2} \otimes \cdots\}, \{E_{\mathbf{b}} = E_{b_1} \otimes E_{b_2} \otimes \cdots\}, \ldots$ (possibly involving more than one POVM) which give sufficient information to reconstruct ρ to arbitrary accuracy.

This is a theorem in quantum mechanics. The name "axiom of local tomography" comes from various attempts to *derive* quantum mechanics by first assuming that this statement is true in some more general context (for some examples, see []). We stated this as possibly involving multiple POVMs because that is how tomography often actually works (e.g., in our examples above), but in fact it is always possible to do tomography with a single POVM, as we will discuss in Sec. 4.3.[7]

We are stating the theorem somewhat loosely here, and it can be tightened in various ways, but the key point is this: not only can you always reconstruct a given state by measurements, you can *always do it with local measurements*! This is very non-obvious, since quantum states themselves can contain non-local in formation (e.g., entanglement). Nevertheless it is true, and important in practice. This fact implies

[7] **(Actually, is it always possible to do LOCAL tomography with one POVM...? –dc)**

in particular that we can detect non-local information like entanglement by doing local measurements; for example, Bell tests with photons can be done by sending two entangled photons very far apart from each other and measuring them locally with photodiodes separated by kilometer-scale distances []. We will discuss this more in Ch. **??**.

Although we could prove this in general, let's instead prove it constructively in an example of clear interest: n qubits. This will highlight a number of practical concerns about efficiency in actually implementing a tomography algorithm.

4.2.1. Example: Pauli tomography of n qubits

Consider now an arbitrary state ρ of n qubits. The total Hilbert space can be factored $\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2 \otimes \cdots \otimes \mathcal{H}_n$, and has dimension $d = 2^n$. We can parametrize a general state using a simple generalization of Eq. (4.2):

$$\rho = \frac{1}{d} \sum_{i=1}^{d^2} S_i \Sigma_i, \quad S_i := \operatorname{tr} \Sigma_i \rho, \tag{4.36}$$

where the Σ_i are called Pauli strings, formed by choosing one Pauli operator (or the identity) on each qubit, for example

$$\Sigma = \sigma_x \otimes 1 \otimes \sigma_z \otimes \cdots, \text{ or } \Sigma = \sigma_z \otimes \sigma_z \otimes \sigma_y \otimes \cdots, \text{ etc.} \tag{4.37}$$

Since there are four choices of operators on this n-site list, there are $4^n = d^2$ of these. They form an orthogonal basis, $\langle \Sigma_i, \Sigma_j \rangle = d\delta_{ij}$, on the space of n-qubit density matrices.

To perform tomography on this system, we can make a reasonably straightforward generalization of the single-qubit protocol described in Sec. 4.1.1. We form a set of two-outcome POVMs, one POVM for each Pauli string Σ_i, as follows:

$$E_{i,\pm} = \frac{1 \pm \Sigma_i}{2}. \tag{4.38}$$

The outcomes are labeled

$$z_{i,\pm} = \pm 1, \tag{4.39}$$

and correspond to **(Under construction! Need to fix local measurement def in ch 1 –dc)**

4.3. Informationally complete measurements

4.3.1. Example: Random unitaries and SIC-POVMs

Problems

1. **Simulating Pauli tomography.** In this exercise you will work through the basic single-qubit Pauli tomography protocol by simulating it on a computer. This is a great exercise to make sure you understand the basics of what is random, etc.

   (a) Consider some arbitrary density matrix, say $\rho = ((1/4, i/4), (-i/4, 3/4))$ in the σ_z basis. Calculate the probabilities $P(+x), \ldots, P(-z)$ in this state. You can do this by hand or on a computer.

   (b) Using these probabilities, write a Mathematica (or similar) code that simulates an experiment in which you perform M measurements in each of these bases, by generating some data $\{x_1, x_2, \ldots, x_M\}$ and so forth. Start with something like $M = 100$.

   (c) Check your simulation by plotting a histogram of the $+z$ outcomes against the expected number of $+z$ outcomes. Does the error make sense at the order-of-magnitude level?

   (d) Using this simulated data, form the estimator $\tilde{\rho}$ for the density matrix. Compute (using the actual data) the standard errors on the coefficients $\tilde{S}_i$.

   (e) Now use the estimates in the text above to estimate the number of measurements M you need to do in order to get a total error (measured in trace distance) of order $\epsilon \approx 1\%$. Go back in your code and set M to this value. Run the simulation again, compute the trace distance between the true state ρ and your estimator $\tilde{\rho}$, and show that the error is of order 1%.

5. Time evolution and its measurement

In standard quantum mechanics courses time evolution is presented in terms of unitary evolution, for example through the Schrödinger equation. But this is only a valid description of reversible systems, such as those which are totally isolated from the environment. In measurement theory we are always, at minimum, interested in the evolution of a system S which at some point we will couple to a measuring apparatus. In general the system will also couple to some unmonitored part of the environment. Describing the time evolution of S thus requires going beyond the unitary evolution paradigm.

5.1. Quantum channels, operations, and instruments

Suppose we have a system S with Hilbert space $\mathcal{H}_S$, initialized to some state ρ. We want to describe time evolution as some map

$$\rho \to \rho' = L[\rho]. \tag{5.1}$$

What are the most bare-bones properties that this map L needs to satisfy in order for quantum mechanics to make sense, in its normal interpretation? There are a few very simple rules we can state, which essentially are just the demands that ρ' is a valid density matrix:

1. **Conservation of total probability:** Any density matrix has to satisfy $\operatorname{tr}\rho = 1$ so that $\sum_a P(a) = 1$ for any measurement $\{E_a\}$. Thus we need $\operatorname{tr} L[\rho] = 1$ for density matrices, or more generally preservation of the trace $\operatorname{tr} L[\mathcal{O}] = \operatorname{tr}\mathcal{O}$ for any operator.

2. **Hermiticity:** The density matrix has to be Hermitian $\rho = \rho^\dagger$ and positive $\rho \geq 0$. The first condition means that the state is diagonalizable and the second means it has positive eigenvalues, which together ensure that ρ can be interpreted as an ensemble of pure states $\rho = \sum_i p_i |\psi_i\rangle \langle\psi_i|$ with $\{p_i\}$ a probability distribution. Thus we require $L[\mathcal{O}]^\dagger = L[\mathcal{O}]$ for any Hermitian operator $\mathcal{O}^\dagger = \mathcal{O}$, as well as preservation of positivity:

3. **Positivity:** $L[\mathcal{O}] \geq 0$ for any positive operator $\mathcal{O} \geq 0$.

4. **Linearity:** Suppose we know the action of L on some particular states, say $\rho'_\psi = L[|\psi\rangle \langle\psi|]$ and $\rho'_\phi = L[|\phi\rangle \langle\phi|]$. By the usual rules of quantum mechanics, we can

form an initial state that's a statistical mixture, say $\rho = p_\psi |\psi\rangle\langle\psi| + p_\phi |\phi\rangle\langle\phi|$, which represents the scenario where we think we prepared $|\psi\rangle$ with probability p_ψ and similarly for $|\phi\rangle$. Then with probability p_ψ, the final state should be ρ'_ψ, and similarly for ρ'_ϕ. That is, we should have $L[p_a\rho_a + p_b\rho_b] = p_a L[\rho_a] + p_b L[\rho_b]$ for any two states $\rho_{a,b}$. Generalizing, we will require that L is linear in the usual sense that $L[a\mathcal{O}_a + b\mathcal{O}_b] = aL[\mathcal{O}_a] + bL[\mathcal{O}_b]$ for any operators and complex coefficients.

In addition to these, we need a rule that says that if $\mathcal{H}_S$ is part of some larger total Hilbert space $\mathcal{H} = \mathcal{H}_S \otimes \mathcal{H}_{\text{other}}$, that we can "trivially" time evolve the rest of the total state, and still get a valid total density matrix. Here $\mathcal{H}_{\text{other}}$ can represent the measurement apparatus, environment, etc. Let the total initial state be $\rho \otimes \rho_{\text{other}}$, then the appropriate rule is

3′. **Complete positivity:** Consider the channel $L \otimes 1$ acting on $\mathcal{H}_S \otimes \mathcal{H}_{\text{other}}$ as the original channel L on $\mathcal{H}_S$ and the trivial map 1 on $\mathcal{H}_{\text{other}}$. Then we require that $(L \otimes 1)[\rho_{\text{total}}] \geq 0$ for any global state ρ_{total} on the whole $\mathcal{H} = \mathcal{H}_S \otimes \mathcal{H}_{\text{other}}$, i.e., the total final state is still a positive state and thus a valid density matrix, for all choices of $\mathcal{H}_{\text{other}}$ and ρ_{other}.

This last condition obviously implies that the channel L itself is a positive map by just taking $\mathcal{H}_{\text{other}}$ to be empty, so it replaces the original item 3. Complete positivity is however strictly stronger: there are linear, trace-preserving, positive maps which are not completely positive. A famous example is the map sending a state to its transpose $L_T[\rho] = \rho^T$, which is worth writing out because it also explains the notation $L \otimes I$. On a qubit, the transpose map L_T acts as

$$L_T\left[|i\rangle\langle j|\right] = |j\rangle\langle i|, \quad i,j = 0,1. \tag{5.2}$$

Consider now appending another qubit as $\mathcal{H}_{\text{other}}$, and sending in the maximally entangled state $|\psi\rangle = (|00\rangle + |11\rangle)/\sqrt{2}$. The total channel $L_T \otimes 1$ acts as

$$(L_T \otimes 1)[|\psi\rangle\langle\psi|] = |00\rangle\langle 00| + |10\rangle\langle 01| + |01\rangle\langle 10| + |11\rangle\langle 11| = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \tag{5.3}$$

which is just a SWAP gate. One can check easily that this has eigenvalues $-1, +1, +1, +1$. Thus $L_T \otimes 1$ is not a positive map, and so L_T is not completely positive.

The set of linear maps $\rho \to \rho' = L[\rho]$ satisfying the above rules are called CPTP maps (completely positive, trace preserving), and sometimes are called quantum channels. We are going to take this as a basic model of time evolution of a quantum system $\mathcal{H}_S$. These rules are often presented as a set of "obvious truths", but it should be emphasized that these are *assumptions* we are making about the way time evolution works and what the density matrix actually means. One can imagine conspiratorial dynamics which actually violates the above, although this almost always forces a lot of mental gymnastics about the ensuing interpretation. For example, violations of the assumption of linearity have long been studied: [1, 2].

Besides conspiracies, there are some simple settings in which we need to further generalize the above picture. One is if the dimension of $\mathcal{H}_S$ changes in time; this is easy to write down and still covered by a CPTP map/channel. A more interesting case is if we are making a measuerment and learn the outcome, in which case linearity is lost as we will re-emphasize below; this forms a set of maps called "quantum operations" (or sometimes "quantum instruments").

One can additionally demand that time evolution can be undone: this would be the demand that for any channel L, there has to be some inverse L^{-1} such that $L^{-1}[L[\rho]] = \rho$ for any initial state ρ. If L^{-1} exists mathematically, we say that the channel is *invertible*: given knowledge of the final state ρ', we can retrodict the initial state $\rho = L^{-1}[\rho']$. For this inverse to be a physical operation that we can implement with time evolution, we also have to demand that L^{-1} itself is a channel. This is the usual assumption made (implicitly or explicitly) in standard quantum mechanics books. Clearly, under unitary evolution,

$$L[\rho] = U\rho U^\dagger \implies L^{-1}[\rho'] = U^\dagger \rho' U, \tag{5.4}$$

the inverse map exists, and is a channel. We can just "run time backwards". What about the converse? It turns out that an inverse channel $\tilde{L}$ exists *only if* the channel L is unitary

$$L^{-1} \text{ exists and is a channel} \implies L[\rho] = U\rho U^\dagger. \tag{5.5}$$

For a proof, see Preskill Ch. 3. Thus, *time evolution is reversible if and only if it is described by a unitary map.* So what more general channels L describe is non-reversible evolution, such as the evolution of a subsystem coupled to other stuff.

This interpretation can be made much more clear by a set of simple formal results. The first is called the Kraus representation, which states that any channel can be

expressed as

$$\rho \to \rho' = L[\rho] = \sum_a K_a \rho K_a^\dagger, \tag{5.6}$$

for some set of operators satisfying the completeness relation

$$\sum_a K_a^\dagger K_a = 1_S \tag{5.7}$$

This completeness relation should be familiar: it is the same thing we demanded for the Kraus operators used to described the measurement of a system S by coupling it to an apparatus M and measuring the apparatus in the basis $\{|a\rangle_M\}$ [see Eq. (1.15)]. This is no accident! Using the same formal result as we did there, Stinespring's dilation theorem, we can always find some other system $\mathcal{H}_B$, coupling unitary U_{SB} acting on $\mathcal{H}_S \otimes \mathcal{H}_B$, reference state $|0\rangle_B$ and basis $\{|a\rangle_B\}$, such that these Kraus operators can be written

$$K_a = \langle a|U_{SB}|0\rangle_B \,. \tag{5.8}$$

As in the case of measurements, these are operators on $\mathcal{H}_S$ only. The notation B here stands for "bath", rather than using M for "measurement apparatus". The reason is because of the sum over outcomes $\sum_a$ in Eq. (5.6). Recall that in the case of a measurement, once we know the outcome a, the apparatus is projected into $|a\rangle_M$ while the system is updated according to the usual rule,

$$\rho \to \frac{K_a \rho K_a^\dagger}{P(a)}, \qquad \text{(after measurement given outcome } a\text{)}. \tag{5.9}$$

What Eq. (5.6) says is that we update ρ by summing over all the possible outcomes of such a measurement. Notice that Eq. (5.9) is *non-linear*, because $P(a) = \operatorname{tr} E_a \rho$ depends on ρ. Thus measurement itself is a quantum "operation", rather than a channel. In contrast, the $P(a)$ factors are not needed in Eq. (5.6), since ρ' is automatically normalized due to the completeness relation $\sum_a K_a^\dagger K_a = 1$. So Eq. (5.6) is linear, as a channel should be.

This picture is very intuitive. What it says is that a general time evolution described by a channel can always be interpreted as time evolution of a non-isolated system S, which is coupled to an unmonitored bath B. See Fig. 5.1. The "unmonitored" part is crucial: if we are measuring the bath state, then we go back to the usual story of measurements, in which case the S state is instead updated conditioned on our known outcomes as in Eq. (5.9). We will give some more detailed interpretation of this in terms of local and non-local information when we discuss decoherence in the next section.

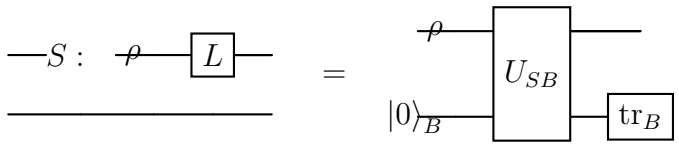

Figure 7: Decomposition of an arbitrary channel into a system-bath interaction.

Example: amplitude damping channel. Way back in our study of Schrödinger's cat in Ch. 1, we introduced a simple model of the decay of an unstable radioisotope. There we modeled the atom as having an unstable excited state $|e\rangle$ which can decay into a stable ground state $|g\rangle$, with some probability $1-p$, where $p = p(t) \sim e^{-t/t_{1/2}}$. This system has a very obvious bath: the other fields of nature which get excited when the decay occurs. For example, an actual ^{238}U atom usually undergoes α decay, in which it emits a helium nucleus. More commonly in quantum optics classes one uses an example of spontaneous photoemission from the decay of an electronic atomic state. In any case, the basic point is that we have dynamics that can be modeled as

$$\begin{aligned} U_{SB} \left|g\right\rangle_S \left|0\right\rangle_B &= \left|g\right\rangle_S \left|0\right\rangle_B \\ U_{SB} \left|e\right\rangle_S \left|0\right\rangle_B &= \sqrt{p} \left|g\right\rangle_S \left|1\right\rangle_B + \sqrt{1-p} \left|e\right\rangle_S \left|0\right\rangle_B . \end{aligned} \tag{5.10}$$

Here, $\left|1\right\rangle_B$ represents whatever excited state of the bath occurs after the decay (e.g., the state of one photon after spontaneous emission). The only important part is that $\left\langle 0|1\right\rangle_B = 0$ so the bath has total knowledge of whether the decay occurred or not.

However, suppose that we can't monitor the bath. What is the channel describing the time evolution we predict for S alone? The Kraus operators [Eq. (5.8)] are

$$\begin{aligned} K_0 &= \left\langle 0|U_{SB}|0\right\rangle_B = \left|g\right\rangle \left\langle g\right|_S + \sqrt{1-p} \left|e\right\rangle \left\langle e\right|_S \\ K_1 &= \left\langle 1|U_{SB}|0\right\rangle_B = \sqrt{p} \left|g\right\rangle \left\langle e\right|_S . \end{aligned} \tag{5.11}$$

These are complete $\sum_{a=0,1} K_a^\dagger K_a = 1_S$. The Kraus representation [Eq. (5.6)] for the channel is therefore

$$\begin{aligned} \rho \to \rho' = L[\rho] = &\left[\left\langle g|\rho|g\right\rangle + p \left\langle e|\rho|e\right\rangle\right] \left|g\right\rangle \left\langle g\right| + (1-p) \left\langle e|\rho|e\right\rangle \left|e\right\rangle \left\langle e\right| \\ &+ \sqrt{1-p} \left[\left\langle g|\rho|e\right\rangle \left|g\right\rangle \left\langle e\right| + \left\langle e|\rho|g\right\rangle \left|e\right\rangle \left\langle g\right|\right] . \end{aligned} \tag{5.12}$$

Each term has a simple interpretation. The first says that, with probability $P = P_{\text{in}}(g) + p \times P_{\text{in}}(e)$, the atom ends up in the ground state, either because it started there or because it decayed from the excited state. The second says that with probability

$(1-p) \times P_{\text{in}}(e)$, the atom was excited but does not decay. The final terms, on the second line, lead to decay of the off-diagonal density matrix elements, a phenomenon called dephasing (or sometimes decoherence) that we will study in more detail in the next section. It is nice to check that $\operatorname{tr} \rho' = 1$ as it should be.

Notice that there is no term proportional to $|e\rangle \langle g|$ in Eq. (5.11), which would represent the atom being excited from its ground state. This is because the bath is initially in the vacuum, so there is no photon for the atom to absorb. However, if we instead considered the bath to be initially in some other state like the thermal state (e.g., an atom surrounded by a thermal bath of photons), then there will be a $|e\rangle \langle g|$ term. We will see this later in a detailed model of an atom in cavity with a hot bath.

It should be physically obvious that this channel is non-reversible. For example, we can look at the purity of the density matrix. Consider a pure initial state $\rho = |e\rangle \langle e|$. Then the purity evolves as

$$\operatorname{tr}\left(\rho^2\right) = 1 \to \operatorname{tr}\left(L[\rho]^2\right) = \operatorname{tr}\left([(1-p)\,|e\rangle \langle e| + p\,|g\rangle \langle g|]^2\right) = 1 - 2p(1-p) < 1. \quad (5.13)$$

So for any decay probability $0 < p < 1$, this channel evolves an initial pure state into a mixed state. In an actual decay, this is not quite the full story: since $p(t) \to 1$ as $t \to \infty$, the final state is *always* the pure state $|g\rangle$, and indeed the purity returns to 1. So, viewed as a time dependent process, starting from the pure state $|e\rangle \langle e|$, we first transition into a mixed state with, and then ultimately end up in a pure state again!

5.2. Decoherence and dephasing

Decoherence and dephasing are phenomena in which the quantum coherence of a system is lost. At the level of an individual system, they are indistinguishable phenomena, although the term decoherence typically carries some further meaning regarding irreversible couplings to an external system. We will discuss this in detail below. First, let's talk about the actual observable signature of these processes.

Consider a qubit prepared in the general pure state

$$|\psi\rangle = a\,|{+z}\rangle + b\,|{-z}\rangle\,, \quad (5.14)$$

where as usual these states are the σ_z basis elements. In this state, the probability of measuring the spin up or down on the z-axis is of course $P(+z) = |a|^2$, $P(-z) = |b|^2$. In a non-commuting basis, the answer is more interesting. The probability to measure,

say, the spin in $+x$ is

$$P(+x) = |\langle +x|\psi\rangle|^2 = \frac{1}{2}\left[1 + (a^*b + ab^*)\right]. \tag{5.15}$$

The terms $\sim a^*b$ represent interference of the $|\pm z\rangle$ states: if $a = 0$ or $b = 0$ (i.e., if we had prepared either $|+z\rangle$ or $|-z\rangle$), these terms go away and we get simply $P(+x) = 1/2$. This interference is analogous to the double-slit experiment, where we prepare a free particle in a superposition of two spherical waves $|\psi\rangle = (|\psi_L\rangle + |\psi_R\rangle)/\sqrt{2}$. The pattern of hits observed on the screen is given by the position distribution

$$P(x) = |\langle x|\psi\rangle|^2 = \frac{1}{2}\left[|\psi_L(x)|^2 + |\psi_R(x)|^2 + \psi_L^*(x)\psi_R(x) + \psi_L(x)\psi_R^*(x)\right]. \tag{5.16}$$

The first two terms are just simple peaks at either $x = x_L$ or $x = x_R$, while the last two terms show interference, $\psi_L\psi_R \sim \exp(-ik(x - [x_L - x_R]))$ where k is the wavevector of the particle. These interference terms, famously, lead to a sinusoidal modulation of the hits on the screen.

The pure qubit state in Eq. (5.14) has density matrix

$$\rho_{\text{pure}} = \begin{pmatrix} |a|^2 & a^*b \\ ab^* & |b|^2 \end{pmatrix}. \tag{5.17}$$

Consider instead the mixed state

$$\rho_{\text{mixed}} = \begin{pmatrix} |a|^2 & 0 \\ 0 & |b|^2 \end{pmatrix}, \tag{5.18}$$

where we set the off-diagonal matrix elements to zero. The probabilities of measuring the spin in $\pm z$ are the same in either of these states, $P(+z) = |a|^2$, $P(-z) = |b|^2$. However, in the mixed state, the probability of measuring $+x$ is

$$\begin{aligned} P(+x) &= \text{tr}\left[|+x\rangle\langle +x|\,\rho_{\text{mixed}}\right] = \text{tr}\left[\frac{1}{2}\begin{pmatrix}1 & 1\\ 1 & 1\end{pmatrix}\begin{pmatrix}|a|^2 & 0\\ 0 & |b|^2\end{pmatrix}\right] \\ &= \frac{1}{2}\,\text{tr}\left[\begin{pmatrix}|a|^2 & |b|^2\\ |a|^2 & |b|^2\end{pmatrix}\right] = \frac{1}{2}\left(|a|^2 + |b|^2\right) \\ &= \frac{1}{2}. \end{aligned} \tag{5.19}$$

One can easily find a similar result for $P(-x)$ and $P(\pm y)$. There are no interference terms, as expected. Indeed, in a general qubit state

$$\rho = \begin{pmatrix} \rho_{00} & \rho_{01} \\ \rho_{01}^* & \rho_{11} \end{pmatrix} \tag{5.20}$$

one finds

$$P(+x) = \frac{1}{2}\left(1 + \rho_{01} + \rho_{01}^*\right), \tag{5.21}$$

so the interference effects (in x measurements) are given precisely by the off-diagonal terms (in the z basis representation). For this reason people sometimes refer to the off-diagonal density matrix elements as "coherences": they quantify the level at which the state is coherently superposed.

Decoherence and dephasing are processes by which an initial state with some degree of coherence [for example, a pure state such as (5.17)] evolves dynamically into a state with less coherence [for example, the mixed state (5.18)], thus destroying observable interference effects. In this simple single-qubit example, we can model this process by the channel:

$$\rho \to \rho' = L[\rho] = \begin{pmatrix} \rho_{00} & (1-p)\rho_{01} \\ (1-p)\rho_{10}^* & \rho_{11} \end{pmatrix}. \tag{5.22}$$

In particular, for large p, this channel produces a totally diagonal density matrix, with no interference terms, just a statistical mixture of $|0\rangle = |+z\rangle$ and $|1\rangle = |-z\rangle$. However, the probabilities of the $\pm z$ states remain fixed. So this channel represents "pure" dephasing/decoherence (and is thus often called "the" dephasing channel).

What kind of system-bath interaction produces this channel? Just as in the case of pure measurements, a given channel does not admit a unique Kraus representation, but rather a unitary family of them. However, we can easily just guess a form for a particular Kraus representation of this channel. The physics here is that the bath should be measuring whether the system is in $|0\rangle$ or $|1\rangle$, but in such a way that this measurement does not actually cause transitions between these states. For example, consider a bath made of a three-level system $|0,1,2\rangle_B$, and

$$\begin{aligned} U_{SB}\,|0\rangle_S\,|0\rangle_B &= \sqrt{1-p}\,|0\rangle_S\,|0\rangle_B + \sqrt{1-p}\,|0\rangle_S\,|1\rangle_B \\ U_{SB}\,|1\rangle_S\,|0\rangle_B &= \sqrt{1-p}\,|1\rangle_S\,|0\rangle_B + \sqrt{p}\,|0\rangle_S\,|2\rangle_B\,. \end{aligned} \tag{5.23}$$

The idea is that the bath starts in $|0\rangle_B$, and then with some probability p, it interacts non-trivially with the system. We can think of it like a scattering event, where with probability p the bath particle scatters and with probability $1-p$ the bath particle just flies past. In the scattering event, if the system is in $|0\rangle_S$, then the bath transitions $|0\rangle_B \to |1\rangle_B$, while if the system is in $|0\rangle_B$, the bath transitions $|0\rangle_B \to |2\rangle_B$, so the final bath state knows the qubit state (if the interaction occurs), but in either case

the system state is left alone. There are three Kraus operators,

$$\begin{aligned} K_0 &= \langle 0|U_{SB}|0\rangle_B = \sqrt{1-p}\,1_S \\ K_1 &= \langle 1|U_{SB}|0\rangle_B = \sqrt{p}\,|0\rangle\langle 0|_S \\ K_2 &= \langle 2|U_{SB}|0\rangle_B = \sqrt{p}\,|1\rangle\langle 1|_S\,. \end{aligned} \tag{5.24}$$

This gives the Kraus representation of the channel

$$\begin{aligned} \rho \to \rho' &= L[\rho] \\ &= (1-p)\rho + p\,\langle 0|\rho|0\rangle\,|0\rangle\langle 0| + p\,\langle 1|\rho|1\rangle\,|1\rangle\langle 1| \\ &= \begin{pmatrix} \rho_{00} & (1-p)\rho_{01} \\ (1-p)\rho_{01}^* & \rho_{11} \end{pmatrix}, \end{aligned} \tag{5.25}$$

which is exactly the dephasing channel of Eq. (5.22).

This system-bath model gives a very intuitive picture of how decoherence works, and more generally illustrates the physics of general lossy channels. The system and bath evolve together unitarily,

$$\rho_{\text{total}} = \rho_S \otimes |0\rangle\langle 0|_B \to rho'_{\text{total}} = U_{SB}\,(\rho_S \otimes |0\rangle\langle 0|_B)\,U_{SB}^\dagger, \tag{5.26}$$

which in general is some entangled joint state of the system and bath. Since the time evolution is unitary, the total entropy is conserved,[8] and the amount of information in the universe is conserved. No information is lost, rather, the system starts to share non-local information with the bath. Now, the issue is, if we're an experimentalist who can't measure the bath, what state should we use to make predictions about the outcomes on S alone? The answer is the reduced density matrix,[9]

$$\rho'_S = \text{tr}_B\,\rho'_{\text{total}} = \sum_a \langle a|U_{SB}\,(\rho_S \otimes |0\rangle\langle 0|_B)\,U_{SB}^\dagger|0\rangle_B = \sum_a K_a \rho_S K_a^\dagger. \tag{5.27}$$

But this is precisely the form of a general channel. So this picture of decoherence of a system S is simply the idea that S becomes entangled with some bath, thus sharing non-local quantum information, and so observations made solely on S will lose interference effects.

[8] Let U be a unitary and ρ a general state of some Hilbert space. Then $\text{tr}\left(U\rho U^\dagger\right)^n = \text{tr}\left(U\rho U^\dagger U\rho U^\dagger \cdots U\rho U^\dagger\right) = \text{tr}\,\rho^n$, using cyclicity of the trace. So every Renyi entropy is preserved and thus so is the von Neumann entropy.

[9] For any pair of systems $\mathcal{H}_A \otimes \mathcal{H}_B$, let $\rho_A = \text{tr}_B\,\rho_{\text{total}}$. Then for any observable $\mathcal{O}_A$, we have $\langle \mathcal{O}_A\rangle = \text{tr}_A(\mathcal{O}_A\rho_A) = \text{tr}_{\text{total}}(\mathcal{O}_A \otimes 1_B\rho_{\text{total}})$. This is the formal statement that the reduced density matrix contains all the information about measurements made only on A.

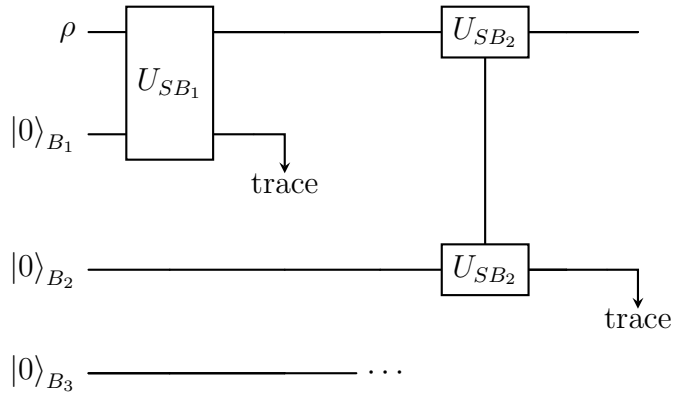

Figure 8: **Markov approximation** with an infinite series of different bath objects.

Sometimes decoherence is confused with wavefunction collapse. Let's clear this up quickly. Consider an initial state $|\psi\rangle = a\,|0\rangle + b\,|1\rangle$ of a qubit. In the dephasing channel example, say in the strongest case $p = 1$, the channel acts as

$$|\psi\rangle\langle\psi| \to \rho' = \begin{pmatrix} |a|^2 & 0 \\ 0 & |b|^2 \end{pmatrix}, \tag{5.28}$$

which is a classical statistical ensemble representing a qubit which we think is prepared in $|0\rangle$ with probability $|a|^2$ and prepared in $|1\rangle$ with probability $|b|^2$. So what the environment has done here is to take the initial pure state and reduce its coherence to the point that we have a simple classical ensemble. Wavefunction collapse, on the other hand, would be the case in which we (not the environment!) projectively measure the qubit, and we know the outcome is either $|0\rangle$ or $|1\rangle$. Once we have the outcome, and assuming the measurement was non-destructive in the sense of Ch. 1, then *we know with complete certainty* which pure state the qubit is in, say

$$|\psi\rangle\langle\psi| \to \rho' = |0\rangle\langle 0| = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \tag{5.29}$$

if we got the outcome $|0\rangle$. The predictions we would make using the state (5.28) and (5.29) are very different. For example, if we projectively measure the qubit in the computational basis, (5.28) predicts that we will get either $|0\rangle$ or $|1\rangle$ with some probabilities, whereas (5.29) says that we get $|0\rangle$ every time. So decoherence is very different than wavefunction collapse.

This discussion of decoherence and dephasing has been in an extremely simplified setting, where a single bath object interacts with the system once, causing some amount of decoherence. In reality, the much more common scenario is instead that we

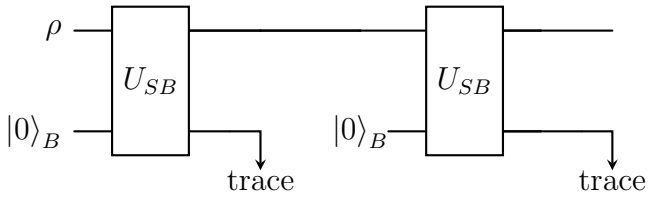

Figure 9: **Markov approximation** with a bath resetting at each timestep. At the level of the channel acting on S, this is equivalent to the circuit of Fig. 5.2.

have a large bath comprised of many degrees of freedom, and the repeated interactions cause decoherence. The large bath is critical because it make the whole process fundamentally irreversible: once the system information is shared with the bath, it is usually unlikely to leak back into the system, although this depends somewhat on how the bath evolution itself works. We will make this idea very precise in the next section, but first let's motivate a basic concept, the Markov approximation.

Consider now a qubit S which interacts with some large number N of bath particles $B_1, \ldots, B_N$, each through the dephasing interaction of Eq. (5.23). We can imagine these happening sequentially, one per timestep $\Delta t = t/N$ during some total time t. See Fig. 5.2. This sequence of interactions causes the system state to evolve as

$$\rho \to L[\rho_S] \to L[L[\rho_S]] \to \cdots \to L^N[\rho_S], \tag{5.30}$$

or rather

$$\rho \to \rho' = \begin{pmatrix} \rho_{00} & (1-p)^N \rho_{01} \\ e^{-\Gamma t}\rho_{10}^* & \rho_{11} \end{pmatrix}. \tag{5.31}$$

At each timestep the probability that the associated B_i scatters is p, so we can write $p = \Gamma \Delta t = \Gamma t/N$ where Γ is the collision rate (in units of events/time, e.g., Hz). So then in the limit of a large bath $N \to \infty$ we have $(1-p)^N = (1 - \Gamma t/N)^N \to e^{-\Gamma t}$ and the whole channel acts as

$$\rho \to \rho' = \begin{pmatrix} \rho_{00} & e^{-\Gamma t}\rho_{01} \\ e^{-\Gamma t}\rho_{10}^* & \rho_{11} \end{pmatrix}. \tag{5.32}$$

After times $t \gtrsim 1/\Gamma$, the off-diagonal matrix elements are quickly decaying to zero, and the state is decohered.

In this N-body bath example, we assumed that at each timestep a new bath particle interacts with the system. At the level of the evolution of the system state, we could equivalently have described this by having just a single bath particle B. The idea in this second picture would be that at each timestep, B is prepared in

	Global	Local
Closed system	$\rho(t) = U\rho(0)U^\dagger$	$\dot{\rho}(t) = i[H, \rho(t)]$
Open system	$\rho(t) = \sum_a K_a \rho(0) K_a^\dagger$	$\dot{\rho}(t) = i[H, \rho(t)] + \sum_a L_a \rho(t) L_a^\dagger$ $- \frac{1}{2} L_a^\dagger L_a \rho(t) - \frac{1}{2} \rho(t) L_a^\dagger L_a$

Table 3: Time evolution in its "global" and "local" forms in quantum mechanics. In the right column, we have assumed that the dynamics are *stationary*: the Hamiltonian H and Lindblad operators L_a are time-independent. One can make this table more general by including time dependence in these terms, as discussed in the text.

$|0\rangle_B$, interacts with S, is measured, and then the outcome of this measurement is "lost". We then proceed to the next timstep, where B's state is "reset" to $|0\rangle_B$, and so forth. See Fig. 5.2. This is purely a mathematical equivalence here; using this second picture is a simple example of a Markov approximation, which says that the bath state can be viewed as resetting to some initial state with some characteristic timescale. We will study this idea in great detail in the next two sections.

5.3. Local time evolution

In the previous sections, we described the general rules for evolution of a system from t_1 to t_2, separated by a finite interval. This generalizes the usual finite-time evolution operator $U(t)$ in a closed system. In this section, we turn to local, differential time evolution: the generalization of the Schrödinger equation $i\partial\psi = H\psi$ or Heisenberg equation $i\partial_t\rho = [H, \rho]$. The general situation is summarized in Table 3.

By local time evolution, we mean that $\dot{\rho}(t)$ at time t depends only on the state at the same time $\rho(t)$, as opposed to something like an integral over all the past values. Integrating the local evolution laws shown here over time, one obtains global evolution rules like those in the left column. However, in general, one can imagine global evolution of the form shown in the left column which does not arise from a local time evolution law.

The equation in the bottom-right corner is known as the Lindblad equation. Next, in Sec. 5.3.1, we derive the Lindblad equation and study how it can be used to model processes like the dephasing channels studied in the previous sections. We then turn in Sec. 5.3.2 to some more general, non-stationary and/or non-Markovian noise models.

5.3.1. The Lindblad equation

Consider a channel $\rho(t) = L[\rho(0)] = E_{\Delta t}[E_{\Delta t}[\cdots E_{\Delta t}[\rho(0)]\cdots]]$ that acts as a sequence of the same "differential" channel $E_{\Delta t}$ applied repeatedly. From our discussion in Sec. 5.1, we know that we can write $E_{\Delta t}$ as unitary evolution of the system with a bath, with the bath state then lost and reset in each time step. Equivalently, we can view this as an infinite set of copies of the same bath evolving unitarily forever. See Fig. **??**. In either case, this channel represents a Markov process: at each time step Δt, a new random process acts on the system state. Here we have also taken this process to be stationary, i.e., the channel $E_{\Delta t}$ is the same at each time step, but we could generalize this further, as discussed below.

In this stationary Markov process, we can derive a simple and important differential equation governing the system state's evolution known as the Lindblad equation. In one time step the system evolves as

$$\rho(t + \Delta t) = E_{\Delta t}[\rho(t)] = \sum_a K_a \rho(t) K_a^\dagger \tag{5.33}$$

by our general results on channels above. The left hand side can be expanded as $\rho(t + \Delta t) = \rho(t) + \mathcal{O}(\Delta t)$, and similarly we can expand the Krauss operators K_a. We will be interested in the $\mathcal{O}(\Delta t)$ terms in order to get a differential equation, which means that K_a will start at $\mathcal{O}(\Delta t^{1/2})$. We can always unitarily rotate the Krauss operators to put the $\mathcal{O}(\Delta t^0)$ term uniquely into it, in which case we have

$$\begin{aligned} K_0 &= 1 + (-iH + \tilde{L})\Delta t \\ K_{a\geq 1} &= L_a \sqrt{\Delta t}. \end{aligned} \tag{5.34}$$

Here, the H and $\tilde{L}$ are Hermitian, so all we did was write the $\mathcal{O}(\Delta t)$ term in terms of its real and imaginary parts; we will see shortly that H has the interpretation of being a Hamiltonian. In the other Krauss operators, the L_a are not necessarily Hermitian; these are called Lindblad or "jump" operators for reasons which will become clear. We can eliminate $\tilde{L}$ by appealing to the completeness relations. To order Δt, these read

$$1 = \sum_{a\geq 0} K_a^\dagger K_a = 1 + 2\tilde{L}\Delta t + \sum_{a\geq 1} L_a^\dagger L_a \Delta t \tag{5.35}$$

which implies

$$\tilde{L} = \frac{1}{2} \sum_{a\geq 1} L_a^\dagger L_a. \tag{5.36}$$

This is like the optical theorem in scattering, which gives a relation between the forward-scattered wave and the total cross section. Both are consequences of unitarity,

where here we mean unitary evolution on the total system plus bath state. In any case, plugging the expansion Eq. (5.34) back into Eq. (5.33), using Eq. (5.36), and taking the continuum limit $\Delta t \to 0$, we obtain

$$\partial_t \rho(t) = -i[H, \rho(t)] + \sum_{a \geq 1} L_a \rho(t) L_a^\dagger - \frac{1}{2} L_a^\dagger L_a \rho(t) - \frac{1}{2} \rho(t) L_a^\dagger L_a. \tag{5.37}$$

This is, as advertised, the Lindblad equation.

Let's see an example: the dephasing channel, in the limit that it acts continuously in time. The global version of this channel was given in Eq. (5.22). Suppose the bath acts in this manner at each timestep Δt, where the probability $p = \Gamma \Delta t$ now represents the probability of the bath particle "scattering" in a given timestep, so Γ is a rate (collisions per time). We can expand the Kraus operators of Eq. (5.24) in the $\Gamma \Delta t \ll 1$ limit:

$$\begin{aligned} K_0 &= 1 - \frac{\Gamma \Delta t}{2} \\ K_1 &= \sqrt{\Gamma \Delta t} \left|0\right\rangle \left\langle 0\right| \\ K_2 &= \sqrt{\Gamma \Delta t} \left|1\right\rangle \left\langle 1\right|. \end{aligned} \tag{5.38}$$

The differential evolution of the system state is

$$\begin{aligned} \rho(t + \Delta t) = E_{\Delta t}[\rho(t)] &= \sum_{a \geq 0} K_a \rho(t) K_a^\dagger \\ &= \rho(t) - \Gamma \Delta t \rho(t) + \mathcal{O}(\Delta t^2) + \Gamma \Delta t \left|0\right\rangle \left\langle 0|\rho|0\right\rangle \left\langle 0\right| + \Gamma \Delta t \left|1\right\rangle \left\langle 1|\rho|1\right\rangle \left\langle 1\right|. \end{aligned} \tag{5.39}$$

From Eq. (5.38) we can identify the jump operators

$$L_1 = \sqrt{\Gamma} \left|0\right\rangle \left\langle 0\right|, \quad L_2 = \sqrt{\Gamma} \left|1\right\rangle \left\langle 1\right| \tag{5.40}$$

and from the first term in Eq. (5.39), we can read off $\tilde{L} = \Gamma \Delta t/2 = \Gamma \Delta t/2(\left|0\right\rangle \left\langle 0\right| + \left|1\right\rangle \left\langle 1\right|)$. Clearly these satisfy the relation $2\tilde{L} = L_1^\dagger L_1 + L_2^\dagger L_2$ [Eq. (5.36)]. Using these observations, one can re-write Eq. (5.39) as, to order Δt,

$$\begin{aligned} \rho(t + \Delta t) - \rho(t) &= -\tilde{L}\rho(t) - \rho\tilde{L} + L_1 \rho L_1^\dagger + L_2 \rho L_2^\dagger \\ &= \sum_{a=1,2} L_a \rho L_a^\dagger - \frac{1}{2} L_a^\dagger L_a \rho - \frac{1}{2} \rho L_a^\dagger L_a, \end{aligned} \tag{5.41}$$

which in the limit $\Delta t \to 0$ is just Eq. (5.3.1), the Lindblad equation with a trivial Hamiltonian $H = 0$.

We can make the connection to our dephasing model even more clear by solving Eq. (5.41). This is an equation for the whole density matrix. It is easy to work out the equations for the individual components. The 00 diagonal element obeys

$$\begin{aligned}\dot{\rho}_{00} &= \langle 0|\dot{\rho}|0\rangle \\ &= \langle 0| \left(\sum_{a=1,2} L_a \rho L_a^\dagger - \frac{1}{2} L_a^\dagger L_a \rho - \frac{1}{2} \rho L_a^\dagger L_a \right) |0\rangle \\ &= \Gamma \left(\rho_{00} - \frac{1}{2}\rho_{00} - \frac{1}{2}\rho_{00} \right) \\ &= 0,\end{aligned} \tag{5.42}$$

i.e., it is a constant $\rho_{00}(t) = \rho_{00}(0)$. The 11 component is similarly a constant. The off-diagonal elements, however, obey

$$\begin{aligned}\dot{\rho}_{01} &= \langle 0|\dot{\rho}|1\rangle \\ &= \langle 0| \left(\sum_{a=1,2} L_a \rho L_a^\dagger - \frac{1}{2} L_a^\dagger L_a \rho - \frac{1}{2} \rho L_a^\dagger L_a \right) |1\rangle \\ &= \Gamma \left(0 - \frac{1}{2}\rho_{01} - \frac{1}{2}\rho_{01} \right) \\ &= -\Gamma \rho_{01},\end{aligned} \tag{5.43}$$

that is, they decay $\rho_{01}(t) = e^{-\Gamma t}\rho_{01}(0)$. Thus the whole density matrix evolves as

$$\rho(t) = \begin{pmatrix} \rho_{00}(0) & e^{-\Gamma t}\rho_{01}(0) \\ e^{-\Gamma t}\rho_{01}^*(0) & \rho_{11}(0) \end{pmatrix} \tag{5.44}$$

which is exactly what we found in Eq. (5.32), where we had the dephasing channel act repeatedly. This is no surprise since that is precisely what this Lindblad equation models.

One could say many more things about the Lindblad equation, but the above captures the essence of what we need in measurement theory. In particular, we could study other examples of the Lindblad equation, for example those which couple the various density matrix elements amongst themselves. Instead, we are going to now study more complex and detailed dynamical models of pure dephasing/decoherence, which will generalize this picture further, and show us how to derive models like these from detailed microscopic system-bath interactions.

5.3.2. Markovian and non-Markovian noise

So far, our discussion of local time evolution has focused on stationary, Markovian models, where a single channel $E_{\Delta t}$ acts repeatedly on the system. Let's now step outside this picture and instead look at a much more general and realistic scenario, in which the bath has arbitrary time dependence and the dynamics can be non-Markovian. We will analyze this in a specific qubit model so that we can connect it to our dephasing models above.

Consider a qubit (say, a spin or two-level atom) which couples to an external magnetic field $B = B(t)$ via

$$H(t) = \frac{1}{2}\left(\omega + \mu B(t)\right)\sigma_z. \tag{5.45}$$

Here μ is a coupling constant (e.g., the magnetic moment), and ω is the bare frequency of the qubit; the $1/2$ is for convenience in what follows. In a real lab, the external field $B(t)$ is random and varies in time; we are thinking of it as, for example, random stray fields while any known DC fields are encoded in ω. This model is assuming that $B = B_z$ is the dominant source of fluctuations; a more realistic coupling would be to consider

$$V_{\text{int}} = \mu\boldsymbol{\sigma}\cdot\mathbf{B}(t) \tag{5.46}$$

but this is much more complicated to analyze. We will do so in Sec. 5.4.

In this section, we will treat $B(t)$ as a classical random field; in the next section, we will quantize it. In the classical picture, let's take B to obey stationary Gaussian statistics with zero mean:

$$\begin{aligned}\langle B(t)\rangle_{\text{noise}} &= 0\\ \langle B(t)B(t')\rangle_{\text{noise}} &= K(t-t').\end{aligned} \tag{5.47}$$

The second equation is what we mean by stationary: it says that the noise is correlated over time, but the amount of correlation depends only on the relative time. The function $K(t-t')$ is called the covariance of the Gaussian distribution. The average here is a classical average,

$$\langle\cdots\rangle_{\text{noise}} = \int DB\, P[B(t)](\cdots) \tag{5.48}$$

where the integral is over all possible time-dependent configurations $B(t)$ and $P[B(t)]$ is the probability distribution of these configurations. We will not need this formalism in any detail; this equation just highlights the idea that we are doing ordinary classical probability theory in the noise field $B(t)$.

We will need the noise statistics in a different form, one which we will encounter repeatedly in the rest of this course: the power spectrum. Define

$$S_{BB}(\nu) = \int_{-\infty}^{\infty} \frac{dt}{2\pi} e^{-i\nu t} \left\langle B(t)B(0)\right\rangle_{\text{noise}}. \tag{5.49}$$

This function, the *noise power spectral density* (PSD), is a function of frequency ν, with dimensions of B^2/ν, e.g., Tesla^2/Hz. What it represents is the noise power. The utility of this function will become clear shortly. We can re-organize this information again to understand what the stationary noise assumption means in frequency domain. Consider the B correlation function:

$$\begin{aligned}
\left\langle B(\nu)B(\nu')\right\rangle_{\text{noise}} &= \int_{-\infty}^{\infty} \frac{dtdt'}{(2\pi)^2} e^{-i\nu t} e^{-i\nu' t'} \left\langle B(t)B(t')\right\rangle_{\text{noise}} \\
&= \int_{-\infty}^{\infty} \frac{dtdt'}{(2\pi)^2} e^{-i\nu t} e^{-i\nu' t'} \int_{-\infty}^{\infty} d\omega e^{i\omega(t-t')} S_{BB}(\omega) \\
&= S_{BB}(\nu)\delta(\nu+\nu').
\end{aligned} \tag{5.50}$$

In other words, the noise is uncorrelated other than $\nu \leftrightarrow -\nu$.

Armed with our noise statistics, let's compute the evolution of a qubit evolving under Eq. (5.45). This particular model is special because we can actually derive the exact time evolution for a fixed realization of the noise, meaning some fixed function $B(t)$. The equation of motion is of course

$$\dot{\rho} = i[H(t), \rho] \tag{5.51}$$

which has solution

$$\rho(t) = U(t)\rho(0)U^\dagger(t) \tag{5.52}$$

with

$$\begin{aligned}
U(t) &= T\exp\left\{-i\int_0^t dt'\, H(t')\right\} \\
&= \exp\left\{-\frac{i}{2}\int_0^t dt'\, \left(\omega + \mu B(t')\right)\sigma_z\right\} \\
&= \exp\left\{-\frac{i}{2}\left(\omega t + \Phi(t)\right)\sigma_z\right\}.
\end{aligned} \tag{5.53}$$

The first equation is the general solution of Eq. (5.51) with a time-dependent Hamiltonian; the T means time ordered exponential. To get to the second line we used the fact that $[H(t), H(t')] = 0$, which makes the time ordering trivial and turns the answer into a simple exponential. The final line has an accumulated phase factor

$$\Phi(t) = \mu\int_0^t dt'\, B(t'). \tag{5.54}$$

Given a specific $B(t)$, this is a real number. With $B(t)$ random, $\Phi(t)$ itself becomes a random variable, with statistics determined by Eq. (5.47).

Since $U(t)$ in Eq. (5.52) only involves the σ_z operator in the exponential, we can easily work out the evolution of the density matrix components in the σ_z basis. The diagonal elements evolve trivially:

$$\rho_{00}(t) = e^{-i(\omega t+\Phi(t))/2}\rho_{00}(0)e^{+i(\omega t+\Phi(t))/2} = \rho_{00}(0), \tag{5.55}$$

while the off-diagonal elements have non-trivial evolution:

$$\rho_{01}(t) = e^{-i(\omega t+\Phi(t))/2}\rho_{01}(0)e^{-i(\omega t+\Phi(t))/2} = e^{-i\omega t}e^{-i\Phi(t)}\rho_{01}(0), \tag{5.56}$$

where the key is the opposite sign in the second exponential from $\sigma_z |1\rangle = -|1\rangle$. Again, this is with $B(t)$ some specific realization of the noise. To predict what we actually see in the lab, we have to do many experiments, each with its own random noise realization, and thus we want the average over noise of this density matrix. This is

$$\langle \rho_{01}(t) \rangle_{\text{noise}} = e^{-i\omega t} \langle e^{-i\Phi(t)} \rangle_{\text{noise}} \rho_{01}(0). \tag{5.57}$$

We should be clear about the interpretation of these symbols. This is still a density matrix (i.e., an operator); the average does *not* mean expectation value over qubit states. The $\rho_{01}(0)$ on the right hand side is also not random: we assume that we can prepare the same initial condition each experiment. What this equation is predicting is the average over time-evolved copies of this initial state.

Let's calculate the phase factor in more detail. We have

$$\begin{aligned} \langle e^{-i\Phi(t)} \rangle_{\text{noise}} &= \int DB\, P[B(t)] \left(1 - i\Phi(t) - \frac{1}{2}\Phi^2(t) + \cdots\right) \\ &= 1 - \frac{1}{2}\langle \Phi^2(t) \rangle_{\text{noise}} + \cdots \\ &= e^{-D(t)}. \end{aligned} \tag{5.58}$$

The first line is the definition of the average and a Taylor expansion of the exponential. The linear Φ term vanishes by our assumption that the noise is zero-mean. The odd powers vanish because the noise is Gaussian. The remaining terms, the even powers of Φ, can be re-summed into the exponential again by assumption of Gaussianity, which in particular says that the higher moments $\langle \Phi^4 \rangle \sim \langle \Phi^2 \rangle \langle \Phi^2 \rangle$, etc. The final

answer is expressed in terms of the function $D(t)$, which is

$$\begin{aligned} D(t) &= \frac{1}{2}\left\langle \Phi^2(t)\right\rangle_{\text{noise}} \\ &= \frac{\mu^2}{2}\int_0^t dt'dt''\left\langle B(t)B(t')\right\rangle_{\text{noise}} \\ &= \frac{\mu^2}{2}\int_0^t dt'dt''\int_{-\infty}^{\infty} e^{i\nu(t'-t'')}S_{BB}(\nu) \\ &= 2\mu^2\int_{-\infty}^{\infty} d\nu\,\frac{\sin^2(\nu t/2)}{\nu^2}S_{BB}(\nu), \end{aligned} \tag{5.59}$$

in terms of the noise PSD, Eq. (5.49).

To summarize, we have found that, averaging over the noise, the qubit density matrix evolves as

$$\rho(t) = \begin{pmatrix} \rho_{00}(0) & e^{-D(t)}\rho_{01}(0) \\ e^{-D(t)}\rho_{01}^*(0) & \rho_{11}(0) \end{pmatrix}. \tag{5.60}$$

This is similar to our continuous dephasing model [Eq. (5.44)], except now we have a somewhat arbitrary-looking function $D(t)$ rather than simple exponential decay $D(t) \sim \Gamma t$. In Eq. (5.60), we have suppressed the notation $\langle\cdots\rangle_{\text{noise}}$ on the left-hand side, to make comparison to Eq. (5.44) clear. One might wonder what the equivalent "noise averaging" was in the old channel model. The answer is that in the dephasing channel and Lindblad calculations, we were averaging over the bath state using tr_B to compute the Kraus operators.

In general, Eq. (5.60) is non-Markovian. If the bath (i.e., the external field $B(t)$) has long auto-correlations in time, we cannot simply approximate it as acting randomly at each time step; its action in one time step can become correlated with its action at a later time step. As a very simple example, consider modeling the noise as a monochromatic field with a random phase,

$$B(t) = B_0\cos(\omega_0 t + \phi), \tag{5.61}$$

where $0 \le \phi \le 2\pi$ is evenly distributed $P[\phi] = 1/2\pi$. The autocorrelator is easy to calculate:

$$\left\langle B(t)B(t')\right\rangle_{\text{noise}} = \int_0^{2\pi}\frac{d\phi}{2\pi}\,\cos(\omega_0 t+\phi)\cos(\omega_0 t'+\phi) = \frac{1}{2}\cos[\omega_0(t-t')] \tag{5.62}$$

which in particular confirms that this model produces stationary noise. The noise PSD is similarly easy to obtain:

$$S_{BB}(\nu) = \int_{-\infty}^{\infty}\frac{dt}{2\pi}\,e^{-i\nu t}\left\langle B(t)B(0)\right\rangle_{\text{noise}} = \frac{B_0^2}{2}\left[\delta(\nu-\omega_0)+\delta(\nu+\omega_0)\right]. \tag{5.63}$$

As one would expect, the noise comes in only at the frequency of the external field. This gives a simple behavior for the density matrix of the qubit:

$$D(t) = 2\mu^2 \int_{-\infty}^{\infty} d\nu \, \frac{\sin^2(\nu t/2)}{\nu^2} S_{BB}(\nu) = \frac{2\mu^2 B_0^2}{\omega_0^2} \sin^2(\omega_0 t/2). \tag{5.64}$$

Here, in contrast to our dephasing examples, the qubit does not simply decohere into a diagonal density matrix. Rather, its off-diagonal elements *oscillate*, and in particular $\rho_{01}(t) = e^{-D(t)}\rho_{01}(0)$ becomes periodic with period set by the noise field. This is highly non-Markovian behavior: the bath actually acts to *re-cohere* the qubit state.

What kind of noise model gives rise to a Markov-type behavior? Suppose that the noise PSD is roughly flat around $\nu = 0$, say with some bandwidth Γ_c.[10] If we are interested in the density matrix's value at times $t \gg 1/\Gamma_c$, then the "window" function in the noise integral Eq. (5.59) is essentially a Dirac delta, in the sense that

$$W_t(\nu) = \frac{\sin^2(\nu t/2)}{\nu^2} \approx \frac{t}{\pi}\delta(\nu). \tag{5.66}$$

In this case, we obtain the simple result

$$D(t) \approx -\Gamma t, \quad \Gamma = \frac{2\mu^2}{\pi} S_{BB}(0). \tag{5.67}$$

The time evolution reduces to the simple Markovian result Eq. (5.44). We also get a microscopic derivation of the decay constant Γ, in terms of the noise power spectral density evaluated at $\nu = 0$. The basic physics is clear once again: if the bath state resets itself fast compared to our observation scale, then the noise acts independently in each time step, and we have a stationary Markov process.

5.3.3. Quantum noise: spin-boson model

Finally, let's consider what happens when the bath is treated quantum mechanically. In the previous section, we modeled the external random field $B(t)$ as classical. Suppose instead that we quantize it into modes, as in Appendix A. To make contact with

[10] For example, consider a simple exponentially decaying bath correlation model

$$\langle B(t)B(t')\rangle_{\text{noise}} = B_0^2 e^{-\Gamma_c(t-t')} \implies S_{BB}(\nu) = \frac{B_0^2\Gamma_c}{\nu^2 + \Gamma_c^2}, \tag{5.65}$$

which is roughly given by $B_0^2/\Gamma_c = \text{const}$ for $\nu \lesssim 1/\Gamma_c$ and then falls off sharply.

those expansions, recall that $B = B_z$ in Eq. (5.45) really means the z-component of the field. This field component can be expanded into modes

$$B(t) = \sum_{k=-\infty}^{\infty} g_k e^{i\omega_k t} b_k + g_k^* e^{-i\omega_k t} b_k^\dagger, \quad g_k = i\sqrt{\frac{\omega_k}{2L^3}}. \tag{5.68}$$

Here, we are still quantizing the field in a large box of volume L^3, so the sum is discrete, over modes k propagating in a single direction (say, the x-axis). We are also implicitly using a basis of polarization vectors $\epsilon_\pm(\mathbf{k})$ so that only one of the modes has $B_z \neq 0$, say the + polarization, so that we don't have to include a sum over polarizations in Eq. (5.68). As discussed above, this set of approximations is somewhat contrived if we are thinking about this as a spin in a vacuum, but it is reasonably accurate for something like a qubit embedded in or at the end of a transmission line.

We wrote $B = B(t)$ here as a time-dependent operator, and will insert this into the Hamiltonian (5.45), in order to make direct contact with the external classical $B(t)$ calculations. This is a little awkward since that is a Schrödinger picture Hamiltonian for the qubit. A more formal way of justifying this would be to write B in the Schrödinger picture, include the free-field Hamiltonian $H_{\text{field}} = \sum_k \omega_k b_k^\dagger b_k$, and then time evolve everything together. What we are doing here is working in the "frame rotating with the field", which just means we transform everything by $U_{\text{rf}} = e^{-iH_{\text{field}}t}$ to get rid of the free-field Hamiltonian, instead putting the field time evolution directly into the $B(t)$ operator. This is essentially the same maneuver used to define the interaction picture in perturbation theory; we will discuss this more in Sec. 5.4.

The main complication here is that $B(t)$ is now a many-body operator involving all the modes b_k. To define its noise statistics we need to specify a state for the field. Let's take it to be the thermal state at temperature T. For a single oscillator $H = \omega a^\dagger a$, the thermal state is

$$\rho_T = \frac{1}{Z}\sum_{n=0}^{\infty} e^{-\omega n/T} |n\rangle \langle n|, \quad Z = \sum_{n=0}^{\infty} e^{-\omega n/T} = \frac{1}{1-e^{-\omega/T}}. \tag{5.69}$$

The creation and annihilation operators in this state are easy to calculate:

$$\begin{aligned} \langle a^\dagger a\rangle &= n_T(\omega) \\ \langle aa^\dagger\rangle &= \langle [a, a^\dagger] + a^\dagger a\rangle = 1 + n_T(\omega) \\ \langle aa\rangle &= \langle a^\dagger a^\dagger\rangle = 0, \end{aligned} \tag{5.70}$$

where the thermal occupation number

$$n_T(\omega) = \frac{1}{e^{\omega/T} - 1}. \tag{5.71}$$

From these results, we can directly calculate the correlations of the field operator of Eq. (5.68):

$$\begin{aligned} \langle B(t)B(t')\rangle &= \sum_{k,k'} g_k g_{k'}^* e^{i\omega_k t} e^{-i\omega_{k'} t'} \left\langle b_k b_{k'}^\dagger \right\rangle + g_k^* g_{k'} e^{-i\omega_k t} e^{i\omega_{k'} t'} \left\langle b_k^\dagger b_{k'} \right\rangle \\ &= \sum_k |g_k|^2 \left[e^{i\omega_k (t-t')} n_T(\omega_k) + e^{-i\omega_k (t-t')} \left(1 + n_T(\omega_k)\right) \right]. \end{aligned} \tag{5.72}$$

To get to the second line, we used the fact that the partition function of a bunch of oscillators factors $Z = \prod_k Z_k$, which follows from the fact that $H = \sum_k H_k$, so that the modes have no correlations amongst themselves. Notice that this noise is stationary.

From this autocorrelation function, we can easily calculate the noise PSD using the same definition Eq. (5.49) given above,

$$\begin{aligned} S_{BB}(\nu) &= \int_{-\infty}^{\infty} \frac{dt}{2\pi} e^{-i\nu t} \langle B(t)B(0)\rangle \\ &= \sum_k |g_k|^2 \left[n_T(\omega_k)\delta(\nu - \omega_k) + \left(1 + n_T(\omega_k)\right) \delta(\nu + \omega_k) \right]. \end{aligned} \tag{5.73}$$

The presence of the $+1$ term, which traces back to the commutator $[b_k, b_{k'}^\dagger] = \delta_{kk'}$, is a hallmark of quantum behavior. It represents noise due to "vacuum fluctuations", and in particular is still present in the limit $T \to 0$. In the classical noise case, there is no corresponding term. This can be formalized by observing that in the classical case, $B(t)$ is a real variable which commutes with $B(0)$, which implies that $S_{BB}(\nu) = S_{BB}(-\nu)$ is a symmetric function (just switch $t \to -t$ and translate the arguments by $2t$ in the definition). Once quantized, we have instead $[B(t), B(0)] \neq 0$ and we find that the noise PSD is no longer symmetric in frequency, but rather now contains an antisymmetric part.

What is the impact of this noise on the qubit? We can solve the qubit's time evolution exactly, using nearly the same formulas from the discussion starting at Eq. (5.52). This is not immediately obvious, because now $[H(t), H(t')] \neq 0$, and so we can't just drop the time-ordering in the exponential. However, we can still reduce the full $U(t)$ to an ordinary exponential in this simple model. Consider writing $U(t)$ as a Magnus expansion, rather than a time-ordered exponential:

$$U(t) = \exp\left\{ -i \int_0^t dt'\, H(t') + (-i)^2 \int_0^t dt' \int_0^{t'} dt''\, [H(t'), H(t'')] + \cdots \right\} \tag{5.74}$$

where the dots represent terms of $\mathcal{O}(H^3)$, which are a series of nested commutators [3]. Here, we have

$$\begin{aligned}[H(t), H(t')] &= \frac{\mu^2}{4}[B(t), B(t')] \\ &= \sum_k |g_k|^2 \left(e^{i\omega_k(t-t')} - e^{-i\omega_k(t-t')}\right) \\ &= 2i \sum_k |g_k|^2 \sin \omega_k (t-t')\end{aligned} \tag{5.75}$$

The first line is a simple consequence of Eq. (5.45), and then the rest shows that this commutator is just an imaginary number [as opposed to proportional to σ_z, like the first term in Eq. (5.74)]. Thus the higher order Magnus terms are zero, and this second order term will have no effect on any of the density matrix elements. We find the same structure as the classical noise case:

$$\rho(t) = \begin{pmatrix} \rho_{00}(0) & e^{-D(t)}\rho_{01}(0) \\ e^{-D(t)}\rho_{01}^*(0) & \rho_{11}(0) \end{pmatrix}, \tag{5.76}$$

where $D(t)$ is given by the same general formula

$$D(t) = 2\mu^2 \int_{-\infty}^{\infty} d\nu \, \frac{\sin^2(\nu t/2)}{\nu^2} S_{BB}(\nu). \tag{5.77}$$

The difference now comes in the form of this noise power.

To connect to more general models, we can make the approximation that the bath is a continuum, so that $\sum_k$ becomes an integral. We define the bath spectral density $J(\nu)$ by

$$\sum_k |g_k|^2 = \int_{-\infty}^{\infty} d\nu \, J(\nu), \tag{5.78}$$

in which case we can write

$$\begin{aligned} S_{BB}(\nu) &= \int_{-\infty}^{\infty} d\nu' \, J(\nu') \left[n_T(\nu')\delta(\nu - \nu') + (1 + n_T(\nu'))\, \delta(\nu + \nu')\right] \\ &= \begin{cases} J(\nu) n_T(\nu), & \nu < 0 \\ J(\nu)[1 + n_T(\nu)], & \nu > 0. \end{cases} \end{aligned} \tag{5.79}$$

Here, we see again the asymmetry in the PSD as a function of frequency. The terms $J(\nu)$ essentially encode the density of states of the bath, while the terms $n_T(\nu)$ give the thermal occupancy of the bath modes. The $+1$ term represents extra noise from bath vacuum fluctuations. Plugging this into $D(t)$, we get

$$D(t) = 2\mu^2 \int_{-\infty}^{\infty} d\nu \, \frac{\sin^2(\nu t/2)}{\nu^2} J(\nu) n_T(\nu) + D_0(t), \tag{5.80}$$

where

$$D_0(t) = 2\mu^2 \int_0^\infty d\nu \, \frac{\sin^2(\nu t/2)}{\nu^2} J(\nu) \tag{5.81}$$

is attenuation from vacuum fluctuations (in particular, generically, $D_0 \neq 0$ in the $T \to 0$ limit).

Much like the classical case, these formulas show that the qubit evolves in a generically non-Markovian fashion. The Markov limit arises in similar regimes as analyzed there: when the PSD is sufficiently flat in some regime $\nu \lesssim \Gamma_c$ and we are interested in $t \gtrsim 1/\Gamma_c$. In this case, for $t \gtrsim 1/\Gamma_c$, we have

$$D(t) \approx \Gamma t, \tag{5.82}$$

with the decay constant

$$\Gamma = \frac{2\mu^2}{\pi} J(0) \left[n_T(0) + \frac{1}{2} \right]. \tag{5.83}$$

The extra $1/2$ comes from the D_0 term, where we normalize $\int_0^\infty dx \, \delta(x) = 1/2$. A prototypical case is that of Ohmic damping, where we take $J(\nu) = A\nu$ at low frequencies $\nu \lesssim \Gamma_c$, where A is some constant, and then falls off like some power of $1/\nu$ above Γ_c. Since $J(\nu) n_T(\nu) \to A\nu * T/\nu = AT$ at $T \gtrsim \nu$, while $J(\nu) \to 0$ as $\nu \to 0$, this gives the simple result that $\Gamma = 2\mu^2 AT/\pi$, and the zero-point fluctuations drop out.

5.4. Thermalization: optical Bloch equations

Problems

1. **Decoherence and irreversibility.** In this problem you will explore how decoherence arises from irreversible interactions with an explicit bath, and the way that this requires a large bath. Consider a system S consisting of a single qubit, coupled to a bath of N other qubits, through the Hamiltonian

$$H = \sum_{i=1}^{N} g_i Z_S Z_i \tag{5.84}$$

where $Z = \sigma_z$ is a Pauli operator, and the g_i are real couplings. Notice that there is no free evolution of the system qubit, for simplicity.

   (a) Start with $N = 1$. Prepare the bath qubit in the state $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$. Assume the system starts in an arbitrary pure state $|\psi\rangle = a\,|0\rangle + b\,|1\rangle$. Compute the reduced density matrix $\rho_S(t)$.

(b) Compute the von Neumann entropy of S as a function of time. Pick some value for g and plot this function. Give an interpretation for what happens at $t = 2\pi/g$ in terms of non-local information shared between S and the bath qubit.

(c) Now let $N \gg 1$ be arbitrary (but finite). Prepare the bath in $|+\rangle_1 |+\rangle_2 \cdots |+\rangle_N$. Compute the reduced density matrix $\rho_S(t)$.

(d) Find a time t_{rec} where $\rho_S(t_{\text{rec}}) = \rho_S(0)$. This is called a Poincaré recurrence.

(e) Suppose the g_i are randomly drawn from a normal distribution with width σ_g centered around $g = 0$. Write a code that draws these values for $N = 1, 10, 100$ and plot the von Neumann entropies of S for each case as a function of time (i.e., there should be three plots). What happens to t_{rec}?

(f) Using your result for the reduced state $\rho_S(t)$, write an estimate for the off-diagonal elements in the short-time limit $t \approx 0$, as an exponential function of the couplings g_i [hint: $\cos x \approx (1 - x^2) \approx e^{-x^2}$ for small x]. Let N be very large, and find Γ_{dec} such that $\rho_{01}(t) \approx e^{-\Gamma_{\text{dec}} t} \rho_{01}(0)$.

2. **Decoherence from gas collisions.** Suppose you prepare an atom, say rubidium, in a superposition of two locations $|\psi\rangle = (|x_L\rangle + |x_R\rangle)\sqrt{2}$. You can model these states as Gaussians in position space with narrow widths. We can think about doing this by setting up some optical trap with two minima. For details ask anybody in Holger Müller's group.

(a) There will be ambient gas in the chamber. Let's model it as a thermal bath of helium atoms at room temperature $T = 300$ K and some pressure P which we can try to lower with a vacuum pump. Estimate the de Broglie wavelength of these gas atoms.

(b) Suppose the spatial superposition is such that $\Delta x = |x_L - x_R|$ is order microns. Compare this to the de Broglie wavelength of the bath. Suppose one of these gas particles collides with the trapped atom. Will the bath particle learn the location of the atom? What will the reduced state of the atoms become if this collision occurs? Keep in mind that $m_{Rb} \gg m_{He}$, and the Rb is in a trap. Does its location change? Assuming that there is some probability p for the scattering event in the first place, model the collision as a channel acting on the atom.

(c) Assume that the scattering cross section is approximately given by the geometric cross section $\sigma \sim A_{\text{He}} \sim (0.1 \text{ nm})^2$. Given pressure P, estimate the decoherence time of the atom. Evaluate this numerically for something reasonable, say $P \sim 10^{-8}$ Pa.

References

[1] S. Weinberg, "Testing quantum mechanics," *Annals of Physics* **194** no. 2, (1989) 336–386.

[2] S. Weinberg, "Precision tests of quantum mechanics," *Physical Review Letters* **62** no. 5, (1989) 485.

[3] S. Blanes, F. Casas, J.-A. Oteo, and J. Ros, "The magnus expansion and some of its applications," *Physics reports* **470** no. 5-6, (2009) 151–238.

A. Quantum mechanics of the electromagnetic field

In this appendix we collect some basic facts about the quantum treatment of the electromagnetic field, which will be used repeatedly through the rest of the notes.

A.1. Quantization of the field

The energy content of the electromagnetic field is given by its Hamiltonian

$$H = \frac{1}{2}\int d^3\mathbf{x}\, \mathbf{E}^2(\mathbf{x}) + \mathbf{B}^2(\mathbf{x}). \tag{A.1}$$

This expression is true both classically and in the quantum theory. When we quantize the electromagnetic field, in the absence of interactions, it becomes a bunch of harmonic oscillators. Here we review the basic procedure.

Recall that we can write the EM field in terms of a gauge potential $A_\mu(\mathbf{x}, t)$, a Lorentz 4-vector, as

$$\begin{aligned} E_i(\mathbf{x}, t) &= \partial_0 A_i(\mathbf{x}, t) - \partial_i A_0(\mathbf{x}, t) \\ B_i(\mathbf{x}, t) &= \sum_{j,k} \epsilon_{ijk}\partial_j A_k(\mathbf{x}, t). \end{aligned} \tag{A.2}$$

Here, $\mu = 0, x, y, z$ is the full spacetime index, 0 is the time component, and $i = x, y, z$ label the spatial components. There is freedom in this decomposition ("gauge symmetry"), which in particular means we can set $A_0 = 0$, and further choose $\sum_\mu \partial_\mu A^\mu = 0$ ("Lorentz gauge"). The classical equations of motion for A_μ in this gauge are simply

$$\left[\partial_0^2 - \nabla^2\right] A_\mu(\mathbf{x}, t) = 0 \tag{A.3}$$

which have plane wave solutions ("modes") of the form

$$A_\mu(\mathbf{x}, t) \sim \epsilon_\mu(\mathbf{k}) u_\mathbf{k}(\mathbf{x}) e^{-i\omega_k t}, \tag{A.4}$$

where $\mathbf{k}$ is the momentum vector of the plane wave, $\omega_k = |\mathbf{k}|$ is the energy (thus frequency in our $\hbar = c = 1$ units), $\epsilon_\mu(\mathbf{k})$ is a polarization vector which is orthogonal to the direction of propagation $\sum_\mu k^\mu \epsilon_\mu(\mathbf{k}) = 0$ with $k_0 := \omega_k$, and finally $u_\mathbf{k}(\mathbf{x}) = e^{i\mathbf{k}\cdot\mathbf{x}}/\sqrt{2\omega_k L^3}$ are the spatial profiles of the waves. For simplicity we are assuming that we are looking at fields in a cubic box of linear size L with periodic boundary conditions; free space is recovered by $L \to 0$. In particular, this means that the set of allowed momenta $\mathbf{k}$ is discrete, $\mathbf{k} = \pi\mathbf{n}/L$ with $\mathbf{n} = (n_x, n_y, n_z)$ a vector of integers.

The plane wave solutions form a complete set, and so we can write the general solution of the equations of motion as a Fourier sum

$$A_\mu(\mathbf{x},t) = \sum_{\mathbf{k},s} \epsilon_{s,\mu}(\mathbf{k}) u_{\mathbf{k}}(\mathbf{x}) e^{-i\omega_k t} a^\dagger_{\mathbf{k},s} + \epsilon^*_{s,\mu}(\mathbf{k}) u^*_{\mathbf{k}}(\mathbf{x}) e^{i\omega_k t} a_{\mathbf{k},s}. \tag{A.5}$$

The conjugated terms ensure that A_μ is real. The sum $s = 1, 2$ is over the two independent polarization vectors for each direction $\mathbf{k}$. The coefficients $a_{\mathbf{k},s}$ and their conjugates give the Fourier coefficients for each mode in the expansion. Classically, these are just complex numbers.

Now, to move to the quantum theory, we first note that Eq. (A.3) is not just the classical equation of motion for A_μ. It is also the Heisenberg-picture equation once we turn A_μ into an operator (or more accurately, a set of operators $A_\mu(\mathbf{x})$, one for each spatial location $\mathbf{x}$). That is, it can be derived from $\dot{A}_\mu = i[H, A_\mu]$. Thus in Eq. (A.5), we can promote the $a_{\mathbf{k},s}$ and their conjugates to operators. We assign the commutation relations

$$[a_{\mathbf{k},s}, a^\dagger_{\mathbf{k}',s'}] = \delta_{\mathbf{k}\mathbf{k}'}\delta_{ss'}. \tag{A.6}$$

The reason for doing this is better justified by starting with the canonical commutation relations on A_μ and its conjugate momentum $\pi_\mu \sim \dot{A}_\mu$, but this is a little subtle in a gauge theory and beyond the scope of what we need here. Using Eqs. (A.5), (A.6), (A.2), one finds that the Hamiltonian of the field [Eq. (A.1)] reduces to

$$\begin{aligned} H &= \sum_{\mathbf{k},s} \frac{\omega_k}{2} \left[a_{\mathbf{k},s} a^\dagger_{\mathbf{k},s} + a^\dagger_{\mathbf{k},s} a_{\mathbf{k},s} \right] \\ &= \Lambda + \sum_{\mathbf{k},s} \omega_k a^\dagger_{\mathbf{k},s} a_{\mathbf{k},s}. \end{aligned} \tag{A.7}$$

This is, of course, just the sum of a bunch of harmonic oscillator Hamiltonians, one for each wavevector $\mathbf{k}$ and polarization s. The constant

$$\Lambda = \sum_{\mathbf{k},s} \frac{\omega_k}{2} \tag{A.8}$$

arises from the commutation relation, and is very infinite (it diverges like $1/L^4$ as $L \to \infty$, which you can see by turning the sum into an integral in the continuum limit). However, it is just a number and has no effect unless we are considering gravity — which we are not — so we will ignore it. Physically, it corresponds to zero-point energies of all the field modes.

In many problems of interest, we have sources and/or detectors that produce and/or detect the electromagnetic field in a very narrow band of frequencies, say

$\Delta\omega \lesssim 1/L$. In this case, we will often focus on just a single mode of the field, say wavevector $\mathbf{k}_0$ and polarization s_0. The Hamiltonian of this single mode is simply

$$H_{\text{one mode}} = \omega a^\dagger a, \tag{A.9}$$

where $a = a_{\mathbf{k}_0,s_0}$ and $\omega = \omega_{k_0}$.