

State tomography using weak continuous quantum measurements

Shoumik Chowdhury[†]

Personal Section

Although I began studying physics at school in the 7th grade, my interest in the subject came largely from reading popular science books. One such book was called *Einstein for Everyone*; a slim volume, it provided a non-technical rundown of all of the major areas in modern physics, from atomic physics all the way ‘up’ to cosmology. I was particularly intrigued with the book’s description of quantum mechanics – the thought of ‘Schrödinger’s cat’ being both alive and dead simultaneously was fascinating to me, as was the notion that measuring something could change the outcome.

Over the next several years I read many more books, and attended several summer programs through which I learned more about quantum physics. My freshman year of high-school, I also began taking computer science; this is how I first heard of quantum computers. Given my continued interest in this area over the years, I was excited to - in my junior year - obtain a research position with Dr. Rajamani Vijayaraghavan at the Quantum Measurement and Control Laboratory, of Tata Institute of Fundamental Research (TIFR). The project we decided upon was a study on types of quantum measurements, the very problem I had been puzzled by since I first read about quantum physics.

Before beginning, however, I had to familiarize myself with linear algebra, the mathematical language of quantum mechanics. It was certainly enriching to be able to study quantum mechanical ideas from a mathematical, rather than purely conceptual, standpoint. As my project lay at the intersection of theoretical and experimental work, I also became acquainted with computational physics, and the notion of programming computer models to test predictions. Indeed, much of my work involved simulating quantum measurements using C code. Thus, I began to work my way up, studying examples about weak measurements, such as in quantum trajectories. Writing a simulation to model this proved to be a fruitful exercise in developing the intuition needed for the main problem I wanted to tackle: simulating a new measurement scheme for quantum state tomography proposed in a paper by D. Das and Arvind. (Don’t worry, I will explain what state tomography is in the next section!)

It was not easy recreating a measurement scheme from its description in scientific literature, and I learned to seek out advice from my mentor, senior project student Malay Singh. Our method of simulation relied on quantum Bayesian statistics (a modified form of Bayes’ probability theorem applicable to quantum measurements. Numerous difficulties arose, the most fundamental of which was the need for a high-quality random number generator – this is obviously critical to simulating quantum phenomena. Furthermore, translating matrix operations from linear algebra into computer code proved to be more challenging than I had expected; it took several weeks of searching for suitable C libraries, writing, and rewriting before the program gave the correct output in modeling the tomographic process. The last stage of my research involved verifying Das and Arvind’s results, while also placing it into the context of other scientific literature – this meant comparing it to descriptions in other papers relating to weak measurement tomography.

In completing this project, I realize that the lines between different fields of study are becoming increasingly blurred in research. My research problem, though it had implications for theory, was very much involved with the experimental aspects of physics, as well as computational physics. I highly recommend simulation based projects to any high-school student interested in advanced physics, or even computer science! Overall, my biggest takeaway from this was being able apply myself to solve complex problems, and persevering through the (sometimes arduous) process! My project was immensely rewarding, and I look forward to further work in this area.

[†]Email: shoumikdc@gmail.com

Research Section

I. Introduction

In quantum mechanics, the state of a system can be fully represented by its wave function. Knowledge of the wave function of a quantum system serves as a direct means to predict the outcome of a measurement on an observable of that system. The ability to accurately determine the wave function of a possibly unknown quantum state is, therefore, of fundamental importance.

However, as we know, it is no easy task to ‘measure’ the state of a quantum system. The notion of measurement itself carries a different connotation in quantum mechanics than in classical physics. Measurements performed on quantum systems are invasive, irrevocably disturbing the system. We say that measurements cause backaction on the state being measured. For ‘projective’ measurements in particular, the wave function appears to ‘collapse’ into only one of the eigenstates of the initial superposition. It is observed that subsequent measurements on a ‘collapsed’ state give the same measurement result. That is to say, no further information can be gained. This is none other than the famous ‘measurement problem’ of quantum mechanics [1]. As the outcome of one single projective measurement cannot be accurately predicted, measurements must be performed on many identical copies of the quantum system, known as an *ensemble*. This process is known as quantum state tomography, where a series of measurements is performed on an ensemble in order to estimate the quantum state [2].

As opposed to projective measurements, another class of generalized measurement is known as the ‘weak’ measurement. Here, the coupling between the system and the measuring device is made weak, thus causing proportionately less backaction. The superposition state does not fully collapse into one of the eigenstates, but instead the system may be reused for subsequent measurements [3–5]. Consequently, however, this means that only small amounts of information can be extracted from each measurement [6]. It is important to understand however, that both weak and projective measurements fall under the more general class of Positive Operator Valued Measurements (POVMs), which are often interpreted as projective measurements on a higher-dimension Hilbert space [1, 6].

Several schemes have been proposed that make use of weak measurements for various applications [7–11]. Specifically, several examples of quantum state tomography using weak measurements can be found in Refs. [12–14]. Most notably, however, weak continuous measurements are finding applications in quantum error correction [16] and for feedback in superconducting quantum circuits [15, 17] where projective measurements cause too much backaction on the system.

Building upon these, this study aims to explore the apparent dichotomy between weak and projective measurements. We consider a scheme proposed by Das and Arvind, which uses weak measurements for quantum state tomography. Das and Arvind believe that their protocol can outperform ‘standard’ quantum state tomography based on projective measurements [18]. This result is surprising, and thus we seek to reproduce their protocol and place it into the context of current knowledge about weak measurements. We test the scheme on known quantum states using quantum Bayesian statistics and use the expectation values to reconstruct the state. This process is then compared for both weak and projective measurements.

This paper is organized as follows. In Section II, we review certain key concepts from quantum measurement theory. In Section III, we present the Das-Arvind scheme and our simulation used to test it for both weak and projective measurements. We present the results of these simulated measurements in Section IV. Section V compares the results to other protocols, and Section VI discusses future work.

II. Background

A. The von Neumann Model of Measurement

Although the von Neumann model initially set out to describe projective measurements, it can easily be applied to weak measurements as well [19]. Let us consider a generalized quantum system with initial state vector $|\psi\rangle$, which can be expressed as

$$|\psi\rangle = \sum_{i=1}^n c_i |a_i\rangle \quad (1)$$

with complex coefficients c_i and eigenstates $|a_1\rangle, |a_2\rangle, \dots |a_n\rangle$. These are the *basis states* described by the von Neumann model. As we are working with qubits, we will consider the specific case of a quantum two-level system, characterized by a two-dimensional Hilbert space \mathcal{H}_2 . We can label the basis states of this system $|0\rangle$ and $|1\rangle$, which refer to the eigenstates of the z-component of spin, σ_z . Any general state in this system can then be written as

$$|\psi\rangle = \alpha |0\rangle + \beta |1\rangle \quad (2)$$

By convention, the eigenvalues of this system are $+1$ and -1 , respectively. We know from Born's rule that the probability of the measurement outcome $|0\rangle$ is $|\alpha|^2$ and for outcome $|1\rangle$ is $|\beta|^2$. In order to satisfy the normalization condition $\langle\psi|\psi\rangle = 1$, clearly $|\alpha|^2 + |\beta|^2 = 1$.

Let us now consider a measurement device with a continuous meter describing the variable q . We follow the description given in [18], wherein the initial quantum state $|\phi\rangle$ for this device follows a Gaussian distribution, with a spread (i.e. standard deviation) Δq . The state, then, is given by

$$|\phi\rangle = \sqrt{\frac{\epsilon}{2\pi}} \int_{-\infty}^{\infty} dq e^{-\epsilon q^2/4} |q\rangle \quad (3)$$

where $\epsilon = 1/(\Delta q)^2$. The ‘weakness’ of the measurement is characterized by $1/\sqrt{\epsilon}$. We see that $|\phi\rangle^2$ give a standard Gaussian probability density distribution corresponding to the likeliness of obtaining each meter reading q . Real measurement devices have characteristics that are modeled well by a Gaussian quantum state, and thus this choice of distribution is apt [20]. After the measurement apparatus interacts with the system, the two wave functions $|\psi\rangle$ and $|\phi\rangle$ become entangled, and the resulting output state $|\psi_{out}\rangle$ can be given by

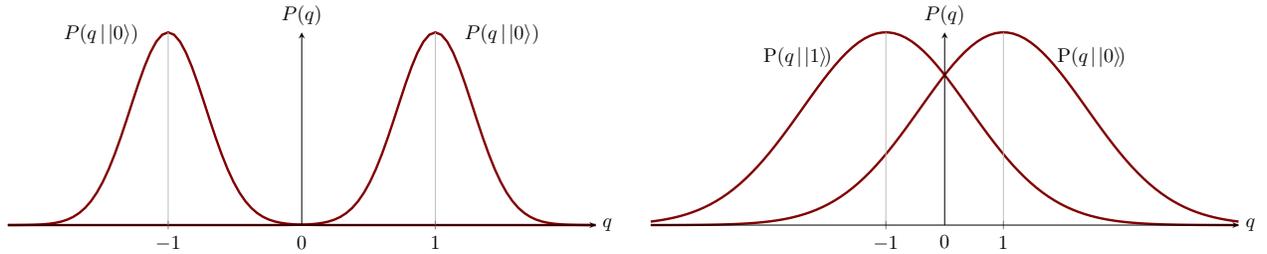
$$|\psi_{out}\rangle = \sqrt{\frac{\epsilon}{2\pi}} \int_{-\infty}^{\infty} dq \alpha e^{-\frac{\epsilon(q-1)^2}{4}} |0\rangle \otimes |q\rangle + \sqrt{\frac{\epsilon}{2\pi}} \int_{-\infty}^{\infty} dq \beta e^{-\frac{\epsilon(q+1)^2}{4}} |1\rangle \otimes |q\rangle \quad (4)$$

The state $|\psi_{out}\rangle$ corresponds to two Gaussian distributions centered around the respective eigenvalues of $|0\rangle$ and $|1\rangle$. We must state explicitly, however, that the distributions given by Eq. (4) only represent the probability *amplitudes*. We can find the corresponding probability density functions by taking $||\psi_{out}\rangle|^2$, where the expectation value is $+1$ for the outcome $|0\rangle$ and -1 for the outcome $|1\rangle$. The conventional way to depict this is to plot the two probability density distributions as sketched in Fig. 1a.

We use the notation $P(\cdot|\cdot)$ to express conditional probability. Thus, $P(q|0)$ on Fig. 1a should be read as ‘the probability of obtaining a meter reading q given that the system is in the eigenstate $|0\rangle$.’ From Bayes’ theorem, and in particular its quantum counterpart, we can *update* the state based on the measurement result q [21]. We know that

$$P(|0\rangle|q) = \frac{P(|0\rangle) \cdot P(q|0)}{P(q)} \quad (5)$$

where $P(|0\rangle)$ is the *a priori* probability of the system collapsing into the $|0\rangle$ eigenstate and $P(|0\rangle|q)$ is the *a posteriori* probability after obtaining a pointer value q . The probability $P(q)$ is the combined probability distribution of obtaining a measurement value in either state, given simply by $P(q) = P(|0\rangle) \cdot P(q|0) + P(|1\rangle) \cdot P(q|1)$



(a) Illustration of quantum measurement outcomes described in Eq. (4)

(b) Illustration of a weak measurement, with a large initial spread Δq .

Fig. 1: Some nice pictures

Qualitatively, we see that the two distributions in Fig. 1a do not overlap appreciably, and there is little ambiguity in the measurement result: a positive pointer value implies the outcome $|0\rangle$, whereas a negative value implies $|1\rangle$. If, however, we were to increase the initial spread of the pointer (Δq), the resultant Gaussians would look more like those in Fig. 1b. Here, there is an ambiguity in the measurement outcome, as the two Gaussians overlap significantly. Interestingly, a pointer value of 0 gives us no information about the quantum state but consequently causes no backaction, leaving the system unchanged! We will test the Das-Arvind protocol for various values of Δq (i.e by varying ϵ) to test whether weak measurements possess any advantages over projective measurements (which can be represented in our model as ‘strong’ measurements where Δq is very small).

B. The Bloch Sphere and the Density Operator

Recall that a Bloch sphere is a unit sphere where every point inside and on the surface of the sphere represents a unique quantum state [1]. Points on the surface represent pure states, whereas points within the sphere are indicative of mixed states; the latter exist as a mixture of states, and are often represented as density operators. The Bloch sphere is a useful tool to visualize qubits, and can be obtained by parametrizing Eq. (2), while still satisfying the normalization condition, as

$$|\psi\rangle = \cos\frac{\theta}{2}|0\rangle + e^{i\varphi}\sin\frac{\theta}{2}|1\rangle \quad (6)$$

where the parameters θ and φ specify the spherical coordinates on the surface of a unit sphere in \mathbb{R}^3 . We see that the poles of this sphere correspond to the qubit eigenstates $|0\rangle$ and $|1\rangle$. ‘Pure states’ are those whose coordinates lie on the surface of the Bloch sphere, and can be denoted by a position vector $\vec{a} = (x, y, z) = (\sin\theta\cos\varphi, \sin\theta\sin\varphi, \cos\theta)$. An example of a state vector drawn on the Bloch sphere is given in Fig. 2.

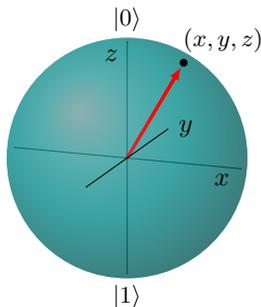


Fig. 2: Bloch sphere representation of a qubit, with $\vec{a} = (x, y, z)$

We can also draw a link between the Bloch sphere representation of a quantum state and the density operator, also known as a ‘density matrix’. Recall now that the density operator ρ of a quantum state $|\psi\rangle$ is described most generally as

$$\rho = \sum_i k_i |\psi_i\rangle \langle\psi_i| \quad (7)$$

This is a compact representation for a quantum state and was devised by von Neumann in 1927 [22]. As a shorthand to Eq. (7), ρ can also be written as a matrix with elements k_i . The dimensions of this matrix correspond to the number of Hilbert dimensions of the quantum state: a wave function in \mathcal{H}_n requires an $n \times n$ matrix. For the specific case of a qubit, we can use the Bloch sphere vector \vec{a} to arrive at the density matrix, given that

$$\rho = \frac{1}{2}(I + x\sigma_x + y\sigma_y + z\sigma_z) = \frac{1}{2} \begin{pmatrix} 1+z & x-iy \\ x+iy & 1-z \end{pmatrix} \quad (8)$$

Here, I is the identity matrix, and σ_x, σ_y , and σ_z are the 2×2 Pauli matrices.

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \quad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

Using this new notation, let us write Eq. (2) as a density matrix. The qubit state is given by $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, and thus correspondingly $\langle\psi| = \alpha^*\langle 0| + \beta^*\langle 1|$. Here, and throughout this study, the notation $*$ signifies the complex conjugate. We see that

$$\begin{aligned}\rho &= |\alpha|^2 |0\rangle\langle 0| + \beta^*\alpha |0\rangle\langle 1| + \alpha^*\beta |1\rangle\langle 0| + |\beta|^2 |1\rangle\langle 1| \\ &= \begin{pmatrix} |\alpha|^2 & \beta^*\alpha \\ \alpha^*\beta & |\beta|^2 \end{pmatrix}\end{aligned}$$

Clearly, the diagonal elements of ρ give the probabilities of obtaining the outcomes $|0\rangle$ and $|1\rangle$, respectively. Thus, the normalization condition for a 2×2 density matrix is that its trace sums to 1.

C. Simulating weak measurements using quantum trajectories

The action of weak measurements on a quantum state can be demonstrated using a quantum trajectory. Studying this example will help introduce several key methods that are useful for simulating quantum state tomography. If a series of weak measurements are performed on a quantum state, each measurement will cause some amount of backaction on the system. We can describe the discrete time evolution for a state ρ as a function $\rho(t)$, corresponding to successive measurements being made on the state. This evolution is known as a quantum trajectory and was first introduced as a theoretical tool to study open quantum systems [23–25]. A single weak measurement results in a probability distribution similar to that depicted in Fig. 1b. For a continuous measurement signal, the result is binned into discrete time intervals τ , corresponding to successive weak measurements. Using either a stochastic master equation [26] or Bayesian inference [15, 27], the state is updated after each ‘measurement.’

The process by which weak measurements can be simulated is twofold. Let us take a qubit in the state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$. After we entangle the qubit with the measurement apparatus, the result of the measurement can be given by a wave function $|\psi_{out}\rangle$ described in Eq. (4). To compute the respective probability densities of obtaining a particular pointer value, we take $||\psi_{out}\rangle|^2$ and observe that the outcome $|0\rangle$ is obtained with a probability $|\alpha|^2$ and the outcome $|1\rangle$ is obtained with a probability $|\beta|^2$. Thus, in order to simulate the result a weak quantum measurement, we must first choose a distribution based on the probabilities $|\alpha|^2$ and $|\beta|^2$. Using the chosen distribution, we can generate a measurement outcome. A simple model for choosing the Gaussian is a *biased coin toss*; this method may be more familiar as the key step used to generate binomial random variables. In this case, a uniformly distributed random number r between 0 and 1 is generated. Given a quantum state ρ , the probability that r is less than ρ_{00} is $|\alpha|^2$. This can be implemented using the commands as shown in Fig. 3.

Input: Quantum state given by $\rho = \begin{pmatrix} \rho_{00} & \rho_{01} \\ \rho_{10} & \rho_{11} \end{pmatrix}$

Output: $\mu \leftarrow$ Mean of the chosen Gaussian

- 1: **procedure** COINTOSS(ρ_{00})
- 2: $\mathcal{U}[a, b] \leftarrow$ uniform random number between $[a, b]$
- 3: $r \leftarrow \mathcal{U}[0, 1]$
- 4: **if** $r \leq \rho_{00}$ **then** $\mu \leftarrow +1$ ▷ For outcome $|0\rangle$
- 5: **else if** $r > \rho_{00}$ **then** $\mu \leftarrow -1$ ▷ For outcome $|1\rangle$

Fig. 3: Algorithm for choosing a Gaussian using a biased coin toss

In actual simulations, however, the accuracy of the *biased coin toss* will depend on the quality of random numbers generated by $\mathcal{U}[a, b]$. Through the result of COINTOSS, we have effectively ‘chosen’ a distribution: the $|0\rangle$ Gaussian has a mean $+1$ and the $|1\rangle$ Gaussian has a mean -1 .

We can use these Gaussians described by $\mathcal{N}(\mu, \sigma)$ to generate measurement outcomes M following a normal distribution. Subsequently, we can use M to *update* the state by calculating the conditional probabilities $P(M||0\rangle)$ and $P(M||1\rangle)$. As described earlier, these in turn give the updated state via Bayes’ theorem. Given that at time t , $\rho_{00}(t)$ and $\rho_{11}(t)$ are simply $P(|0\rangle)$ and $P(|1\rangle)$ respectively, we can rewrite Eq. (5) using the following equations [28]:

$$\rho_{00}(t + \tau) = \frac{\rho_{00}(t)P(M||0\rangle)}{P(M)} \quad (9)$$

and correspondingly,

$$\rho_{11}(t + \tau) = \frac{\rho_{11}(t)P(M || 1)}{P(M)} \quad (10)$$

On the other hand, ρ_{01} is updated by

$$\rho_{01}(t + \tau) = \rho_{01}(t) \sqrt{\frac{\rho_{00}(t + \tau)\rho_{11}(t + \tau)}{\rho_{00}(t)\rho_{11}(t)}} \quad (11)$$

Once we have found ρ_{01} , we can find ρ_{10} by taking the complex conjugate. As t is a discrete variable here, we can set $\tau = 1$ and then treat t as the index of an array for the purposes of simulation. The procedure described by Eqs. (9), (10) and (11) is summarized below in Fig. 4.

Input:

Initial quantum state given by $\rho(0) = \begin{pmatrix} \rho_{00}(0) & \rho_{01}(0) \\ \rho_{10}(0) & \rho_{11}(0) \end{pmatrix}$

$N \leftarrow$ number of measurements

$\sigma \leftarrow$ standard deviation, gives measurement strength

$\mathcal{N}(\mu, \sigma) \leftarrow$ generates Gaussian random number

Output: $\rho(t)$

```

1: procedure TRAJECTORY
2:    $t \leftarrow 0$ 
3:   for  $t < N$  do
4:      $\mu \leftarrow \text{COINTOSS}(\rho_{00}(t))$  ▷ Choosing a Gaussian
5:      $M \leftarrow \mathcal{N}(\mu, \sigma)$  ▷ Measurement outcome
6:      $P(M || 0) \leftarrow \frac{1}{\sqrt{2\sigma^2\pi}} e^{-(M-1)^2/2\sigma^2}$ 
7:      $P(M || 1) \leftarrow \frac{1}{\sqrt{2\sigma^2\pi}} e^{-(M+1)^2/2\sigma^2}$ 
8:      $P(M) \leftarrow \rho_{00}(t)P(M || 0) + \rho_{11}(t)P(M || 1)$ 
9:      $\rho_{00}(t + 1) \leftarrow \frac{\rho_{00}(t)P(M || 0)}{P(M)}$ 
10:     $\rho_{11}(t + 1) \leftarrow \frac{\rho_{11}(t)P(M || 1)}{P(M)}$ 
11:     $\rho_{01}(t + 1) \leftarrow \rho_{01}(t) \sqrt{\frac{\rho_{00}(t+1)\rho_{11}(t+1)}{\rho_{00}(t)\rho_{11}(t)}}$ 
12:     $\rho_{10}(t + 1) \leftarrow (\rho_{01}(t + 1))^*$ 
13:     $t \leftarrow t + 1$ 

```

Fig. 4: Pseudocode for generating a quantum trajectory $\rho(t)$

In TRAJECTORY, as with in COINTOSS, the quality of the random number generators will directly influence the accuracy of the method. (A convenient method for generating Gaussian random numbers is to transform a $\mathcal{U}[-1, 1]$ distribution into a $\mathcal{N}(0, 1)$ distribution. The Gaussian can then be adjusted to a specified μ and σ by taking $\sigma \times \mathcal{N}(0, 1) + \mu$. This is known as the Marsaglia method - for more details see Ref. [29]).

As the ρ_{00} and ρ_{11} must always sum to 1, the trajectory $\rho_{11}(t) = 1 - \rho_{00}(t)$. Figures 5a and 5b show two such trajectories for a qubit initialized to the state $\rho_1 = \begin{pmatrix} 0.5 & 0.5 \\ 0.5 & 0.5 \end{pmatrix}$.

In both examples above $N = 100,000$. However, in the case of Fig. 5a, we have (arbitrarily) $\sigma = 22.36$, implying a relatively weak measurement. Thus, here, the system ‘collapses’ into one of its eigenstates after approximately 600 time steps (i.e. 600 measurements). As $\rho_{00}(t) \rightarrow 1$, it can be said with near certainty that the system eventually collapses into the $|0\rangle$ state. Conversely, in Fig. 5b, the system collapses into the $|1\rangle$ eigenstate, as indicated by the value of $\rho_{11}(t) \rightarrow 1$. In this case, however, the measurement is relatively stronger, with $\sigma = 5$. Consequently, the ‘collapse’ occurs after only 25 time steps: that is to say, after 25 measurements the system is almost certainly in the $|1\rangle$ state. Of course, in order to arrive at a general conclusion about the average ‘collapse’ time, one must average over many trajectories. As the measurement strength increases, this average collapse time will reduce. Intuitively this makes sense, as for

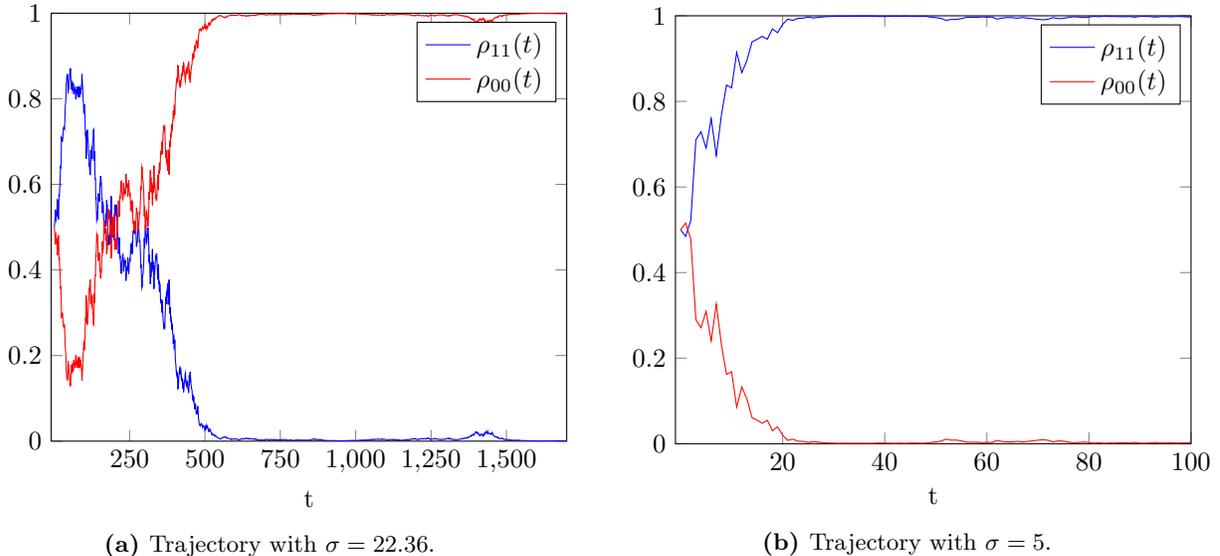


Fig. 5: Sample quantum trajectories of a qubit set to initial state ρ_1 .

small values of σ , the state would immediately collapse into one eigenstate upon measurement. Likewise, by observing Figs. 5a and 5b, we notice that successive weak measurements eventually give the same result as a strong measurement: the wave function collapsing into one of the eigenstates. This fact is well-supported theoretically: Oreshkov et. al. showed that any generalized measurement can be written as a series of weak measurements [30].

III. Method: Single-qubit Quantum State Tomography

Quantum state tomography is the process by which an unknown state is ‘reconstructed’ from a series of measurements on the state [13]. Of course, by performing state tomography on a known state, and comparing this to the estimated state, we can determine the efficacy of the tomography process. As a classic example given in Refs. [1,31], imagine that Alice prepares an ensemble of qubits and gives it to Bob to observe. Bob, skeptical of Alice’s description, can perform quantum state tomography on the ensemble to determine the true state of the qubits. In an ideal scenario, Bob would be able to carry out measurements on an infinite number of copies of the qubit. However, in practice, experiments are limited by a fixed ensemble size.

From literature, we know that projective measurements outperform weak measurements for state reconstruction in actual experiments, where ensemble sizes are large (with over thousands of copies of the qubit state). However, Das and Arvind raise the question of whether weak measurements possess an advantage for small ensemble sizes. This is prompted by the possibility of recycling states for subsequent measurements. Das and Arvind propose the following scheme: weak measurements are performed on the z and x Pauli components of spin, followed by a projective measurement along the y component. This scheme is tested, and it is shown that under certain conditions, weak measurements can outperform projective measurements. Das and Arvind claim this scheme offers a “novel way of extracting information from a quantum system”, as the two weak measurements extract information without greatly disturbing the state.

The protocol considers a fixed ensemble of identical qubits in the state ρ . On this ensemble, a σ_z measurement is carried out with a coupling strength ϵ_1 . For each qubit the meter reading is observed. Due to the associated backaction of the measurements, the state of the qubits is changed, denoted by the density matrix ρ_1 . Next, a σ_x measurement is performed with coupling strength ϵ_2 , following the same process as for σ_z . This gives an updated state ρ_2 . Lastly, a projective σ_y measurement is carried out on each of the qubits in the ensemble. This subroutine is repeated a fixed number of times, and the results are averaged in order to avoid statistical error. A summary of this subroutine is given in Fig. 6. In the first two measurements, the Gaussians overlap, indicating that the measurement strength is weak. These are similar to the illustration in Fig. 1b. In the σ_y measurements, however, the Gaussians are well separated. By computing the expectation

values for σ_x , σ_y , and σ_z , we arrive at the Bloch sphere vector (x, y, z) of the state.

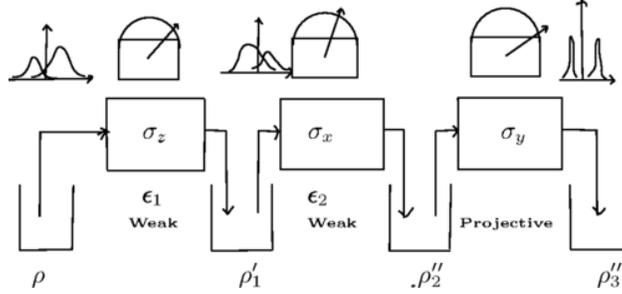


Fig. 6: Depiction of the Das-Arvind protocol. Reprinted from Ref. [18]

When the Gaussians overlap significantly, and the system has similar probability of being in either eigenstate, it becomes difficult to ascribe a meter reading to its associated eigenstate. Das and Arvind address this by introducing a *discard parameter*. This is defined as a region in between the two Gaussians of width $2a$. Any meter reading that falls in this region is not considered; however, the state is preserved and recycled for subsequent measurements. Meter readings that do lie outside of this ‘discard region’ are binned into two categories. For example, in the case of the σ_x measurement, readings to the right of this region indicate σ_x has a value $+1$, whereas readings to the left indicate a value -1 .

This study seeks to reproduce their results. As we will see in Section IV, it is surprising that a weak measurement scheme should outperform a protocol based on projective measurements. Our first step involves recreating the procedure used by Das and Arvind. In their scheme, Das and Arvind calculate the probability density function using the reduced density matrix of the measurement apparatus, which itself is obtained by taking the partial trace of the state. We, however, make use of Bayesian inference to update the density matrices. In theory, this is functionally equivalent and does *not* degrade the quality of the state tomography. Furthermore, we work with the assumption that the strength of both weak measurements is equal – that is, $\epsilon_1 = \epsilon_2 = \epsilon$. Imagine now that the ensemble of n qubits is prepared in a state given by

$$\rho = \rho_{00} |0\rangle \langle 0| + \rho_{01} |0\rangle \langle 1| + \rho_{10} |1\rangle \langle 0| + \rho_{11} |1\rangle \langle 1| \quad (12)$$

Referring to Eq. 8, we can compute the coordinates of the Bloch sphere vector (x, y, z) as

$$x = \rho_{01} + \rho_{10} \quad (13)$$

$$y = i(\rho_{01} - \rho_{10}) \quad (14)$$

$$z = \rho_{00} - \rho_{11} \quad (15)$$

On each member of the ensemble, three measurements are performed. For σ_z the measurement outcome M_z is recorded and used to update the state using Bayesian inference (the process is the same used as in COINTOSS and TRAJECTORY). However, the value of M_z may fall inside the discard region for some of the n qubits in the ensemble. Thus, we define the number of ‘valid measurement outcomes’ C_{σ_z} . If a measurement is ‘valid’, then we increment a sum S_{σ_z} by $+1$ for values of M_z to the right of the discard region, and by -1 for values of M_z to the left of the discard region. The expectation value $\langle \sigma_z \rangle$ is then given by

$$\langle \sigma_z \rangle = \frac{S_{\sigma_z}}{C_{\sigma_z}} \quad (16)$$

Next, before making a measurement on σ_x , the qubit must be rotated to the correct orientation; this is because in physical devices, and in this implementation, measurements can only be made along one axis. The matrices for rotation by an angle θ about the x- and y-axes are given in Ref. [1] as,

$$R_x(\theta) = \begin{pmatrix} \cos \frac{\theta}{2} & -i \sin \frac{\theta}{2} \\ -i \sin \frac{\theta}{2} & \cos \frac{\theta}{2} \end{pmatrix} \quad R_y(\theta) = \begin{pmatrix} \cos \frac{\theta}{2} & -\sin \frac{\theta}{2} \\ \sin \frac{\theta}{2} & \cos \frac{\theta}{2} \end{pmatrix}$$

All this while, we have considered our measurement device to be aligned with σ_z . Thus, we must rotate the updated state ρ_1 around the y-axis by $\theta = -90^\circ$ such that the qubit x-axis shifts to the position of the

z-axis. This is carried out by

$$\rho_1^{x\text{-basis}} = R_y(-90^\circ)\rho_1 R_y(-90^\circ)^\dagger \quad (17)$$

where \dagger is the conjugate transpose matrix operation. Using this state $\rho_1^{x\text{-basis}}$ we perform a measurement on σ_x through the same process used to measure σ_z . After generating a measurement outcome M_x , the state is updated, giving a new density matrix $\rho_2^{x\text{-basis}}$. Likewise, the number of valid measurements is noted and the expectation value is given by

$$\langle \sigma_x \rangle = \left(\frac{S_{\sigma_x}}{C_{\sigma_x}} \right) e^{\epsilon/2} \quad (18)$$

Here, a correction factor (as given in Ref. [18]) is introduced to compensate for the backaction on the system. This factor is dependent on measurement strength parameter ϵ . Finally, in order to make a σ_y measurement, the qubit must be rotated into the y-basis. First, the qubit state is rotated back into the original (z-) basis by $\rho_2 = R_y(-90^\circ)^\dagger \rho_2^{x\text{-basis}} R_y(-90^\circ)$. Next, the qubit state is rotated by $\theta = 90^\circ$ about the x-axis, so that

$$\rho_2^{y\text{-basis}} = R_x(90^\circ)\rho_2 R_x(90^\circ)^\dagger \quad (19)$$

By performing a COINTOSS operation with this state $\rho_2^{y\text{-basis}}$ we obtain either +1 or -1. These values are added to the sum S_{σ_y} . As there is no discard parameter for a projective measurement, the number of ‘valid’ measurements will simply be the number of qubits in the ensemble, n . Here, the expectation value is given by

$$\langle \sigma_y \rangle = \left[2 \left(\frac{S_{\sigma_y}}{n} \right) - 1 \right] e^\epsilon \quad (20)$$

Once again, a correction factor accounts for the backaction of the two prior measurements. A full proof of the calculation of $\langle \sigma_y \rangle$ is given in Appendix A, but the intuition behind this result comes from realizing that the ratio S_{σ_y}/n is simply $P(\sigma_y; +)$. These expectation values directly give us the coordinates of the Bloch sphere vector – that is, $\langle \sigma_x \rangle = x_{est}$, $\langle \sigma_y \rangle = y_{est}$, and $\langle \sigma_z \rangle = z_{est}$.

For the scheme of projective measurements, on the other hand, the ensemble is divided into three parts. Each part is first rotated into the correct basis, before a COINTOSS operation is performed. The results of these measurements are used to calculate the expectation values, using the calculation made in Eq. (20) for all three components, except without the correction factor. From here, we arrive at x_{est} , y_{est} , and z_{est} , and can compare the results for both measurement schemes. In order to quantify the comparison, we introduce the fidelity f as the distance between the actual and estimated Bloch sphere vectors. This measure is given by Das and Arvind as:

$$f = 1 - \left[(x - x_{est})^2 + (y - y_{est})^2 + (z - z_{est})^2 \right] \quad (21)$$

We repeat this process to obtain average trends of the tomographic process. In the interest of brevity, the algorithm summarizing the process above is described in Appendix B. Nonetheless, we list here the parameters for the simulation in Table 1.

Table 1: List of parameters used in the simulation

Simulation Parameter	Physical Significance	Description of parameter
rho	ρ	Quantum state being estimated
ensemble	n	Number of qubits in the ensemble
epsilon	ϵ	Gives the easurement strength
sigma	σ	Standard Deviation of Gaussians ($1/\sqrt{\epsilon}$)
N	-	Number of reiterations of protocol
a	a	Discard parameter

IV. Results

A. Example of a randomly-generated state

We first test the scheme on a randomly generated state ρ_A , given by

$$\rho_A = \frac{1}{2} \begin{pmatrix} 1.399 & -0.385 + 0.042i \\ -0.385 - 0.042i & 0.601 \end{pmatrix} \quad (22)$$

This quantum state is used as a test case by Das and Arvind (see Ref. [18]). Thus, testing the protocol allows us to directly compare our results. As with Das and Arvind, we take an ensemble size of 30. However, we find that repeating the process only 10000 times (as done in Ref. [18]) gives rise to large statistical fluctuations. Thus, here, we instead use 100000 iterations to avoid such difficulties. In each run, the fidelity f is calculated. We then calculate the mean fidelity \bar{f} and standard deviation in fidelity $\sigma_{\bar{f}}$ for that particular measurement strength ϵ . By varying ϵ , we are able to reproduce the results noted by Das and Arvind for the state ρ_A .

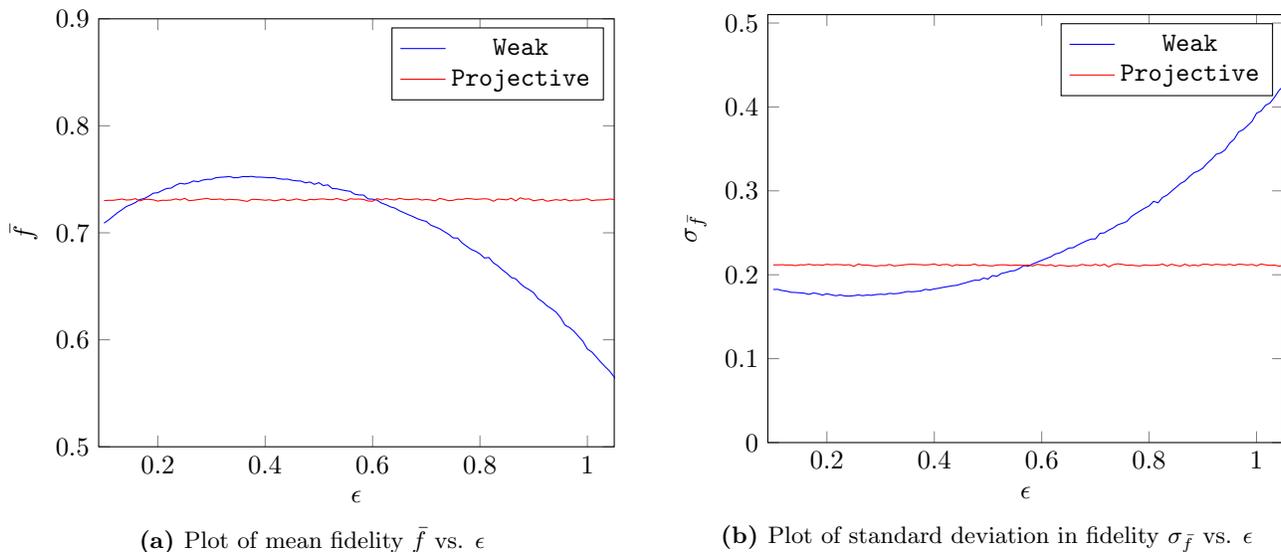


Fig. 7: Plots of mean fidelity (a) and standard deviation in fidelity (b) versus ϵ for weak (blue) and projective (red) measurement schemes. The ensemble of 30 qubits is prepared in the state ρ_A , and the discard parameter $a = 0$.

Through an analysis of Fig. 7, we see that the mean fidelity for the weak measurement protocol peaks at $\epsilon \approx 0.4083$. In this case, the weak measurement scheme offers an improvement over the projective scheme. We clarify that the projective measurement scheme, based on a series of biased coin tosses, results in a roughly constant fidelity. The reason for the weak measurement scheme outperforming the projective scheme, in this case, is not obvious; the former consists of a series of projective σ_y measurements, while the latter makes projective measurements on σ_x , σ_y , and σ_z . Likewise, it is also unclear why the weak measurement scheme displays the trend that it does. The value of ϵ for which the data peaks, in the case of ρ_A , corresponds to a measurement strength that is neither extremely weak nor extremely strong. One hypothesis offered by Das and Arvind is that for very weak measurements, the Gaussians are too wide to be able to determine which eigenstate the pointer value belongs to, whereas for very strong measurements, the state collapses after only one measurement. This thinking is in accordance with Lundeen et. al., who claim that weak measurements must be carried out “in a gentle way” [32]. Nonetheless, having reproduced the results of Das and Arvind, we are encouraged to continue the analysis.

We next repeat the protocol on the state ρ_A , with an ensemble size of 60 qubits. As seen in Fig. 8, this increases the fidelity of the projective scheme. The reason for this is not difficult to see: more coin tosses will give a better estimate of the associated probabilities, and thus a more accurate reconstruction of the state. However, we see that on average the weak measurement scheme does *not* outperform the projective scheme, for a larger ensemble. We can, however, compare the results of Fig. 8 to those in 7 to see that increasing the

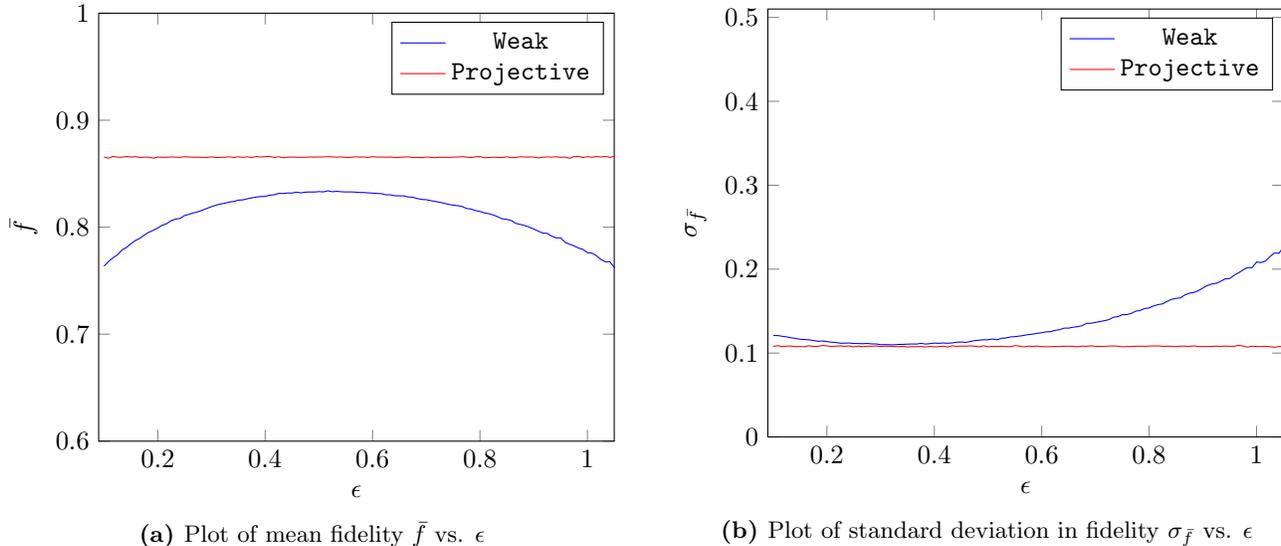


Fig. 8: Plots of mean fidelity (a) and standard deviation in fidelity (b) versus ϵ for weak (blue) and projective (red) measurement schemes. The ensemble of 60 qubits is prepared in the state ρ_A , and the discard parameter $a = 0$.

ensemble size does improve the fidelity for both schemes. However, for the projective measurement scheme this improvement is greater; Fig. 8 shows that the weak scheme peaks at $\bar{f} \approx 0.82$, whereas for the projective scheme $\bar{f} \approx 0.86$ consistently. Das and Arvind also acknowledge that their weak measurement scheme is only effective for small ensemble sizes, and thus this result is not surprising. However, it is interesting to note that by changing the ensemble size, the peak fidelity appears to occur at a different value of ϵ .

B. Measurement Fidelity Depends on the State

We now extend the protocol to test different quantum states. We consider the simple case ρ_B , where

$$\rho_B = \begin{pmatrix} 0.5 & 0.5 \\ 0.5 & 0.5 \end{pmatrix} \quad (23)$$

The parameters of this simulation are identical to those used to measure ρ_A in Fig. 7. That is, the protocol is run for 30 qubits and repeated 100000 times. We set the discard parameter to 0.

On Fig. 9, we see that the weak measurement scheme performs quite poorly compared to the projective scheme. Let us emphasize that the simulation parameters here are identical to those used to measure state ρ_A in Fig. 7. Here, the projective measurement scheme has a mean fidelity $\bar{f} \approx 0.80$, whereas the weak measurement scheme peaks at $\bar{f} \approx 0.62$. Likewise, it is interesting that this peak should occur at $\epsilon = 0.625$, as compared to $\epsilon = 0.408$ for the state ρ_A .

V. Discussion and Summary

In this study, we have been able to reproduce the results of Das and Arvind. In particular, we see that for certain states, the Das-Arvind weak measurement protocol offers slight advantages in mean measurement fidelity. We further verify that the phenomenon observed is dependent on ensemble-size. By running the simulation for larger ensembles, we see that both the Das-Arvind scheme and projective measurement scheme perform better with more qubits – however, the improvement is more pronounced for the projective measurement case. Clearly, as Das and Arvind themselves take note of, the Das-Arvind protocol only performs better for small ensembles. Finally, we see that the protocol does not offer improvements for all states, and that the fidelity of the process is state dependent.

At this stage it is worth calling into question the ‘binning’ process, by which meter readings to the left of the discard region are interpreted to be -1 and those to the right are interpreted as +1. We note, interestingly,

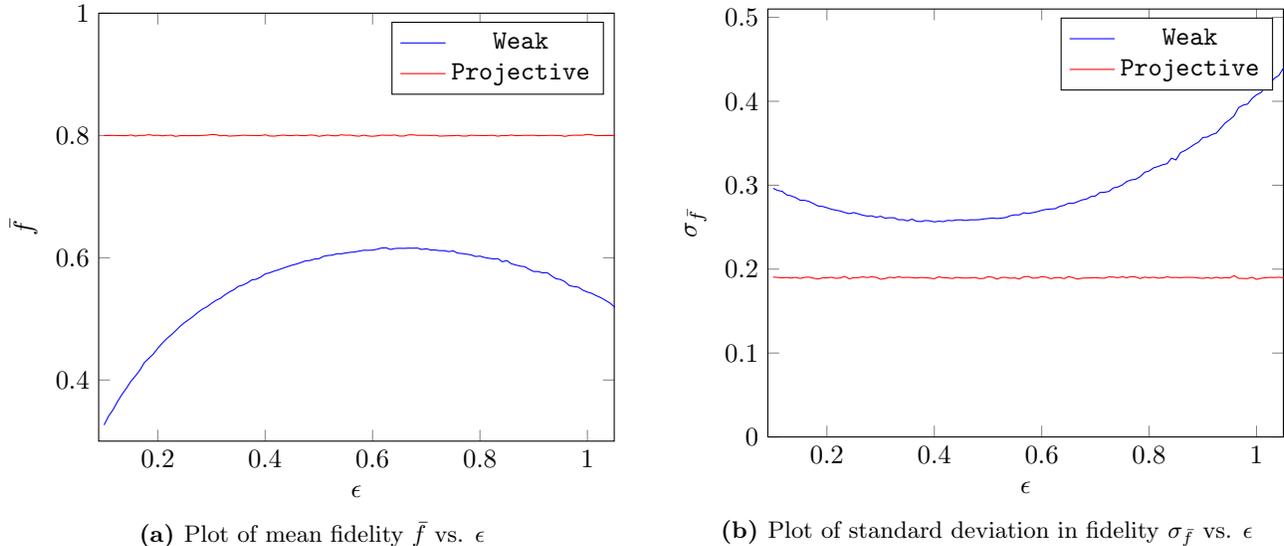


Fig. 9: Plots of mean fidelity (a) and standard deviation in fidelity (b) versus ϵ for weak (blue) and projective (red) measurement schemes. The ensemble of 30 qubits is prepared in the state ρ_B , and the discard parameter $a = 0$.

that if this interpretation is not made (and therefore if the pointer value is added when calculating the sum instead of simply +1 or -1) then the Das-Arvind protocol performs better. Indeed as per Ref. [33], this binning process “degrades tomographic performance.”

The results of Das and Arvind, reproduced here, raise a number of questions for which states the Das-Arvind protocol offers improvements, and the optimal ensemble sizes in each case. Das and Arvind believe that the two weak measurements followed by a projective measurement allow for more information to be extracted from a qubit than from a projective measurement alone. However, we know from Busch’s theorem that any information gained from a quantum system necessarily disturbs the state [34]. The limit on the amount of information that can be extracted is given in Ref. [35].

Instead, we try to explain the trend observed here by placing the Das-Arvind protocol in the context of a more recent work by Caves et. al. [33] In their analysis of the Das-Arvind protocol, the authors clarify that the projective measurement scheme used here and in [18] is a type of a mutually unbiased basis (MUB) scheme, where three orthogonal measurement bases (i.e x- y- and z-axes) are chosen. Caves et. al. show, however, that the MUB scheme is not the most efficient scheme for projective measurements. Indeed, the most efficient scheme involves choosing a random measurement basis from a uniformly distributed set of axes. This is known as a Haar-invariant measurement, and it falls under the general classification of one-dimensional orthogonal projective measurements (ODOPs).

Intuitively, it is not difficult to see why ODOPs provide higher fidelity measurements. Imagine a pure state pointing in the $|\sigma_z; +\rangle$ direction; then the σ_z measurements will be highly accurate, whereas the σ_x and σ_y measurements give the correct result only 50% of the time (i.e relatively inaccurate). Likewise, generally, for a fixed-basis measurement scheme, the quality of state reconstruction will depend on the choice of measurement basis. This makes sense given that we observed the Das-Arvind protocol to be state dependent – the choice of basis is optimal for some states, but not for others. ODOPs on the other hand would overcome this by randomly sampling bases.

Caves et. al. show that the Das-Arvind protocol does not outperform their ODOP-based protocol. Through their analysis it is shown that the first two weak measurements cause random backaction on the state, and result in a distribution around the Bloch sphere from which the final projective measurement is sampled. For the optimum value of ϵ , this results in a uniform distribution. Under these conditions the final projective measurement in the Das-Arvind scheme is essentially random, approximating an ODOP. Thus, although more information is extracted using the Das-Arvind protocol than simply a σ_y projective measurement alone, this is not due to any improvements offered by weak measurements. Instead, the weak measurements in this scheme only serve to randomly rotate the state, an effect that could be achieved

by using a randomly-selected unitary (as is used in ODOPs). Caves et. al. thus conclude that the Das-Arvind protocol does not offer any improvements over an ODOP based scheme. At the present stage in our understanding, our data supports this claim, as the Das-Arvind protocol offers only marginal improvements under certain specific conditions.

VI. Future Work

In the light of [33], we hope to extend this analysis of weak and projective measurement schemes to include other types of generalized measurements. In doing so, we hope to verify the claims made by Caves et. al., and test the efficiency of a Haar-invariant ODOP-based scheme. We will also work towards comparing the efficacy of these various tomographic processes experimentally.

Despite some dispute about the claims of several weak measurement tomography based schemes, the notion of the weak measurement maintains its theoretical and practical significance in quantum computing. From Refs. [34] and [35] we know that there is no way around the uncertainty-disturbance relations, for both weak and strong measurements. Nonetheless, weak measurements may find use in so-called adaptive measurements, where the measurement bases are continuously updated based on the results of a series of weak measurements. Thus, testing such protocols and comparing to the non-adaptive schemes presented here is a further goal for us.

VII. Acknowledgements

The author would like to thank Dr. Rajamani Vijayaraghavan of the Quantum Measurement and Control Laboratory at the Tata Institute of Fundamental Research (TIFR) for his mentorship and support, as well as Malay Singh and Soumen Das for their invaluable guidance and advice.

Appendices

A. How to calculate $\langle \sigma_y \rangle$

Here, we continue from Section III, offering a short proof for the result in Eq. (20). We first imagine trying to find $\langle \sigma_z \rangle$ (yes, in the z-direction) for an arbitrary qubit state ρ . This can be calculated by the trace of ρ times the σ_z Pauli matrix, or

$$\langle \sigma_z \rangle = Tr(\rho \sigma_z) = \rho_{00} - \rho_{11} = \rho_{00} - (1 - \rho_{00}) \quad (24)$$

As here $\rho_{00} = P(|0\rangle)$, we can rewrite this as

$$\langle \sigma_z \rangle = 2P(|0\rangle) - 1 \quad (25)$$

In Section III, with Eq. (20), we rotated the qubit y-axis to the position of the z-axis. Thus, the process used here can be effectively used to calculate $\langle \sigma_y \rangle$ (as we are measuring along the z-axis of the rotated state, corresponding to the y-axis of the original state). Replacing $\langle \sigma_z \rangle$ with $\langle \sigma_y \rangle$ and $P(|0\rangle) = P(\sigma_z; +)$ with $P(\sigma_y; +)$ in Eq. (25) yields Eq. (20). The ‘correction factor’ can then be added by hand.

B. An algorithm to implement the Das-Arvind protocol

We present parts of the algorithm (written using C) that was used to generate our results. In order to run this code using the gcc compiler, the following command-line argument must be provided: `gcc -o FILENAME FILENAME.c -lm -l`. These linker commands are essential, as we make use of the following libraries: `<math.h>`, `<complex.h>`, and `<stdlib.h>`

In order to work with the operators of quantum mechanics, we define a new data type to handle 2×2 matrices. This is done implemented using a `struct`:

```
typedef struct {
    double complex p00, p01, p10, p11;
} matrix;
```

All rotation matrices, and quantum states are defined using this data type. For the latter, note that the numbers `p00`, `p01`, `p10` and `p11` correspond to the matrix elements ρ_{00} , ρ_{01} , ρ_{10} , and ρ_{11} . Using our new data type, we also define a simple `multiply()` function to multiply the corresponding matrix elements. This is necessary to perform the necessary qubit rotations.

For any given run, we perform σ_x , σ_y , and σ_z measurements; in our implementation, we simply loop the process up till the number of qubits, `ensemble`. For an input state `rho`, we first compute the Bloch sphere coordinates using:

```
z_actual = rho.p00 - rho.p11;
y_actual = (1I)*(rho.p01 - rho.p10);
x_actual = rho.p01 + rho.p10;
```

where the constant `I` is imported from `<complex.h>` and represents $i = \sqrt{-1}$. Next we simulate a weak measurement. As mentioned in Section II, this is done using a combination of a biased coin toss and a Gaussian random number to generate a measurement outcome `M`.

```
double coin = cointoss(creal(rho.p00));
if (coin == 0)
    M = gaussrand(sigma, mean0);
else if (coin == 1)
    M = gaussrand(sigma, mean1);
```

The `cointoss()` function is defined exactly like `COINTOSS` from Section II. However, we must convert from a `complex double` data type to a `double`. To generate uniform random numbers, we use the `arc4random()` function from `<stdlib.h>`. It produces random numbers of sufficient quality for our simulation:

```
double cointoss(double p) {
    double x = arc4random() / (RAND_MAX + 1.0);
    //uniform random number between -1 and 1
    x = .5*(x + 1);
    //adjusts the distribution to be over [0,1]

    if (x <= p)
        return 0;
    else if (x > p)
        return 1;
}
```

Using the result of the `cointoss()` function, we randomly select a measurement outcome from a Gaussian distribution. This is done with the `gaussrand()` function, which uses the Marsaglia method [29] to generate normally distributed variables with mean 0 and standard deviation 1. In our definition `gaussrand()` takes two arguments: `sigma` and the mean of the distribution (`mean0 = +1` and `mean1 = -1`). It adjusts the Gaussian distribution by multiplying by `sigma` and adding the mean.

```

double gaussrand(double sigma, double mean) {
    static double V1, V2, S;
    static int phase = 0;
    double X;
    double U1, U2;

    if (phase == 0) {
        do {
            U1 = arc4random() / (RAND_MAX + 1.0);
            U2 = arc4random() / (RAND_MAX + 1.0);

            V1 = 2 * U1 - 1;
            V2 = 2 * U2 - 1;
            S = V1 * V1 + V2 * V2;
        } while(S >= 1 || S == 0);

        X = V1 * sqrt(-2 * log(S) / S);
    } else
        X = V2 * sqrt(-2 * log(S) / S);

    phase = 1 - phase;
    return (double) ((X * sigma) + mean);
}

```

Once we have generated a measurement outcome M , we can update the state using Bayesian inference. The procedure for this is given in the `TRAJECTORY` function in Section II. However, we add the following condition which bins the results according to whether M lies inside the discard region of width $2a$:

```

if (M >= a) {
    sum_x += 1.0;
    count_x += 1.0;
} else if (M <= -a) {
    sum_x += -1.0;
    count_x += 1.0;
}

```

As listed in Section III, the expectation values can be calculated using these counts. We implement this as follows:

```

ex_z[j] = (sum_z / count_z);
ex_x[j] = (sum_x / count_x)*pow(E, (epsilon/2));
ex_y[j] = (2*(count_y / ensemble) - 1)
          *pow(E, epsilon);
double delz = z_actual - ex_z[j];
double delx = x_actual - ex_x[j];
double dely = y_actual - ex_y[j];
fidelity[j] = 1 - (pow(delx, 2) + pow(dely, 2)
                 + pow(delz, 2));

```

Here we are calculating the fidelity for the j th iteration. We run an outer `for` loop up till N (the number of reiterations), we can calculate the mean fidelity by taking the average value of the array `fidelity[]`. Another `for` loop can then be used to repeat this process for different values of `epsilon` (and hence of `sigma`). If we then plot mean fidelity against `epsilon`, we obtain the figures seen in Section IV.

References

- [1] M. A. Nielsen, and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, Cambridge, 2010.
- [2] D. F. V. James, P. G. Kwiat, W. J. Munro, and A. G. White. *Measurement of qubits*. Physical Review A, vol. 64, no. 5, Oct. 2001.
- [3] A. Barchielli. *Measurement theory and stochastic differential equations in quantum mechanics*. Physical Review A, vol. 34, no. 3, 1986.
- [4] A. Barchielli, L. Lanz, and G. M. Prosperi. *A model for the macroscopic description and continual observations in quantum mechanics*. Nuovo Cim, 1982. doi:10.1007/BF02894935
- [5] C. M. Caves and G. J. Milburn. *Quantum-mechanical model for continuous position measurements*. Physical Review A 36, no. 12, 1987.
- [6] T. A. Brun. *A simple model of quantum trajectories*. American Journal of Physics, vol. 70, no. 7, 2002.
- [7] E. Prugovecki. *Quantum two-particle scattering in fuzzy phase space*. Journal of Mathematical Physics, vol. 17, no. 9, 1976.
- [8] S. T. Ali, and E. Prugovecki. *Systems of imprimitivity and representations of quantum mechanics on fuzzy phase spaces*. Journal of Mathematical Physics, vol. 18, no. 2, 1977.
- [9] L. Diósi. *Single qubit estimation from repeated unsharp measurements*. Fortschritte der Physik, vol. 51, no. 2-3, Mar. 2003.
- [10] S. T. Ali, and G. G. Emch. *Fuzzy observables in quantum mechanics*. Journal of Mathematical Physics, vol. 15, no. 2, 1974.
- [11] M. Ueda, and M. Kitagawa. *Reversibility in quantum measurement processes*. Physical Review Letters, vol. 68, no. 23, 1992.
- [12] H. F. Hofmann. *Complete characterization of post-selected quantum statistics using weak measurement tomography*. Physical Review A, vol. 81, no. 1, 2010.
- [13] S. Wu. *State tomography via weak measurements*. Scientific Reports 3, 2013.
- [14] Y. Shikano, and S. Takana. *Estimation of spin-spin interaction by weak measurement scheme*. European Letters Association, vol. 96, no. 4, 2011.
- [15] R. Vijay, C. Macklin, D. H. Slichter, S. J. Weber, K. W. Murch, R. Naik, A. N. Korotkov, and I. Siddiqi. *Stabilizing Rabi oscillations in a superconducting qubit using quantum feedback*. Nature, vol. 490, no. 7418, 2012.
- [16] C. Ahn, H. M. Wiseman, and G. J. Milburn. *Quantum error correction for continuously detected errors*. Physical Review A, vol. 67, no. 5, 2003.
- [17] K. W. Murch, R. Vijay, and I. Siddiqi. *Weak Measurement and Feedback in Superconducting Quantum Circuits*. Superconducting Devices in Quantum Optics. Springer International Publishing, 2016.
- [18] D. Das, and Arvind. *Estimation of quantum states by weak and projective measurements*. Physical Review A, vol. 89, no. 6, 2014.
- [19] I. M. Duck, P. M. Stevenson, E. C. G. Sudarshan. *The sense in which a "weak measurement" of a spin-1/2 particle's spin component yields a value 100*. Physical Review D, vol. 40, no. 6, 1989.
- [20] M. B. Mensky. *Continuous quantum measurements and path integrals*. CRC Press. Vancouver. pp. 32-33
- [21] R. Schack, T. A. Brun, C. M. Caves. *Quantum bayes rule*. Physical Review A, vol. 65, 2001.

- [22] J. von Neumann. *Wahrscheinlichkeitstheoretischer Aufbau der Quantenmechanik*. Göttinger Nachrichten, no. 1, pp 245-272, 1927.
- [23] H. Carmichael. *An Open Systems Approach to Quantum Optics* Springer Verlag, 1993.
- [24] C. W. Gardiner, A. S. Parkins, and P. Zoller. *Wave-function quantum stochastic differential equations and quantum-jump simulation methods*. Physical Review A vol. 46, no. 7, 1992.
- [25] J. Dalibard, Y. Castin, and K. Mølmer. *Wave-function approach to dissipative processes in quantum optics* Physical Review Letters, vol. 68, no. 5, 1992.
- [26] D. Tan, S. J. Weber, I. Siddiqi, K. Mølmer, and K. W. Murch. *Prediction and retrodiction for a continuously monitored superconducting qubit*. Physical review letters, vol. 114, no. 9, 2015.
- [27] K. W. Murch, S. J. Weber, C. Macklin and I. Siddiqi. *Observing single quantum trajectories of a superconducting quantum bit*. Nature, 2013. doi:10.1038/nature12539
- [28] A. N. Korotkov. *Selective quantum evolution of a qubit state due to continuous measurement*. Physical Review B, vol. 63, 2001.
- [29] G. Marsaglia, and T. A. Bray. *A convenient method for generating normal variables*. SIAM Review, vol. 6, no. 3, 1964. pp.260-264.
- [30] O. Oreshkov, and T. A. Brun. *Weak measurements are universal*. Physical review letters vol. 95, no. 11, 2005
- [31] G. M. D'Ariano, M. De Laurentis, M. G. A. Paris, A. Porzio, and S. Solimeno. *Quantum tomography as a tool for the characterization of optical devices*. Journal of Optics B: Quantum and Semiclassical Optics vol. 4, no. 3, 2002.
- [32] J. S. Lundeen, B. Sutherland, A. Patel, C. Stewart, and C. Bamber. *Direct Measurement of the Quantum Wavefunction*. Nature, vol. 474, 2011.
- [33] J. A. Gross, N. Dangniam, C. Ferrie, and C. M. Caves. *Novelty, efficacy, and significance of weak measurements for quantum tomography*. Physical Review A, vol. 92, no. 6, 2015.
- [34] P. Busch. *No Information Without Disturbance": Quantum Limitations of Measurement*. In Quantum Reality, Relativistic Causality, and Closing the Epistemic Circle, Springer Netherlands, 2009.
- [35] S. Massar and S. Popescu. *Optimal extraction of information from finite quantum ensembles*. Physical review letters vol. 74, no. 8, 1995.