Compressed Sensing

Stefan Frick

University of BRISTOL

Compressed Sensing

Abstract: The steady growing number of quantum bits used in modern quantum information experiments gives rise to new problems. Especially if we want to determine the quantum state used in an experiment, i.e. ascertain the density matrix of the state, the number of needed measurement settings scales exponentially bad with $\Theta(4^n)$, where n is the number of qubits.
Compressed sensing is a technique developed to overcome this problem by using matrix completion methods to reconstruct a full density matrix of low rank states with fewer measurements.
This report explains the main ideas of compressed sensing to the reader and gives a (highly incomplete) overview of the work done in the field.

It is my experience that proofs involving matrices can be shortened by 50% if one throws the matrices out. - E. Artin

April 10, 2015

Contents

Chapter 1

Introduction

Modern experiments in the field of quantum information show a steady increase of quantum bits (qubits) [1--3]. This increase leads to an exponential growth of the required measurements needed in an experiment to fully characterise the quantum state.

This fact can be easily seen if we consider the entries in the density matrix of states with increasing numbers of qubits. For example the density operator ρ_1 of a 1-qubit state $|\psi_1\rangle = \alpha_0 |0\rangle + \alpha_1 |1\rangle$ is represented as a matrix as

$$\rho_1 = \begin{pmatrix} \alpha_0\alpha_0^* & \alpha_1\alpha_0^* \\ \alpha_0\alpha_1^* & \alpha_1\alpha_1^* \end{pmatrix}, \tag{1.1}$$

where * denotes the complex conjugate of the coefficient α_i. Since the Pauli matrices

$$\sigma_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \tag{1.2}$$

$$\sigma_2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \quad \text{and} \tag{1.3}$$

$$\sigma_3 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \tag{1.4}$$

together with the identity

$$I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} =: \sigma_0 \tag{1.5}$$

form a basis for density matrices we can reconstruct the density matrix ρ_1 as

$$\rho_1 = \sum_{i=0}^{4} \mathrm{tr}(\rho_1\sigma_i)\sigma_i, \tag{1.6}$$

where $\mathrm{tr}(\rho_1\sigma_i)$ is the Hilbert-Schmidt inner product. Meaning we have to measure all 4 expectation values for the Pauli measurements σ_i.

If we now increase the number of qubits, let's say by one, and consider the 2-qubit state $|\psi_2\rangle = \alpha_0 |00\rangle + \alpha_1 |01\rangle + \alpha_2 |10\rangle + \alpha_3 |11\rangle$, the number of entries in the matrix representation of the density operator ρ_2 increase exponentially.

$$\rho_2 = \begin{pmatrix} \alpha_0\alpha_0^* & \alpha_1\alpha_0^* & \alpha_2\alpha_0^* & \alpha_3\alpha_0^* \\ \alpha_0\alpha_1^* & \alpha_1\alpha_1^* & \alpha_2\alpha_1^* & \alpha_3\alpha_1^* \\ \alpha_0\alpha_2^* & \alpha_1\alpha_2^* & \alpha_2\alpha_2^* & \alpha_3\alpha_2^* \\ \alpha_0\alpha_3^* & \alpha_1\alpha_3^* & \alpha_2\alpha_3^* & \alpha_3\alpha_3^* \end{pmatrix}. \tag{1.7}$$

It is easy to see that we now would need to measure 16 observables to fully characterise the 2-qubit state. Namely the $4^2 = 16$ Pauli expectation values of the form $\sigma_i \otimes \sigma_j$ with $i, j \in \{0, 1, 2, 3\}$.

This method of determining the state produced in a quantum experiment is called full state tomography. The problem with this method is that we number of measurements needed scales as $\Theta(4^n)$ where n is that numbers of qubits involved in the experiment.

Compressed sensing[4] in a quantum sense is a method to overcome this problem in full state tomography by measuring only certain observables and utilising classical matrix completion methods to reconstruct the full information contained in the density matrix.

1.1 Compressed Sensing

Compressed sensing was first discussed in the context of classical signals. The Nyquist-Shannon sampling theorem states that if we sample a classical analogue signal with a certain rate r we can only reconstruct information from the signal which is contained in frequencies lower than $\frac{r}{2}$. This is the reason why audio CDs have a sampling rate of $\approx 44\,kHz$. Meaning that the highest frequency which can be reproduced is $\approx 22\,kHz$, a bit above the hearing threshold of the human ear ($\approx 20\,kHz$).

However it was shown that we can overcome the limitations of the Nyquist-Shannon theorem, if we have prior knowledge of the signal [4, 5]. More precisely the signal has to be sparse in some domain and we have to know about this sparsity.

In particular it is often possible to make certain conclusions about the signal in the frequency domain. If we say that the signal we want to reconstruct is periodic, what we are essentially meaning is that we know the signal is sparse in the frequency domain, i.e. it consist of only a few frequencies.

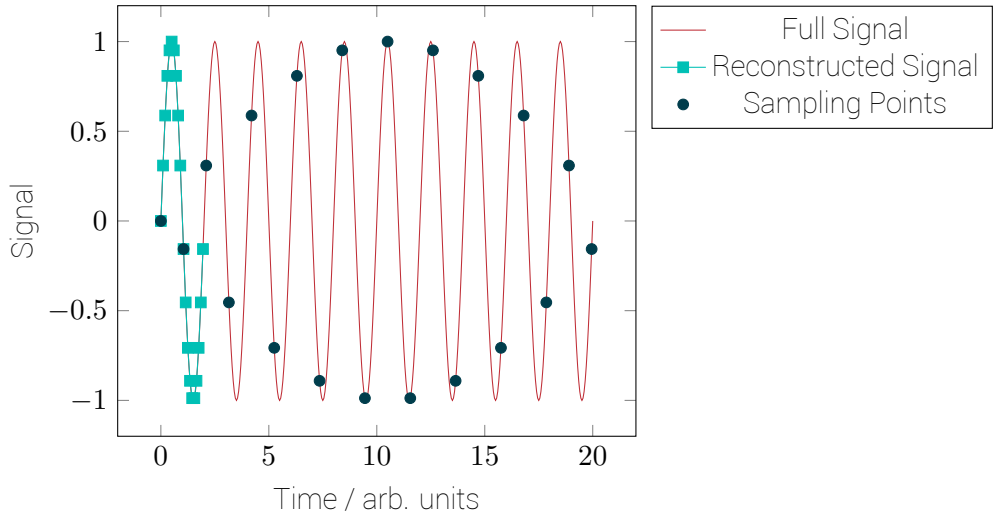

Figure 1.1: An example for a reconstruction of an signal wich is sparse in the frequency domain (here the signal consists of only one frequency). It can be seen that the signal (red) is sampled with a rate slower then half its frequency. The sampling points are shown in dark blue. However, since we know that the signal is sparse in the frequency domain, we can reconstruct the signal (bright blue) by apply a modulo operation on the x-values of the sampling points.

An example for such a reconstruction is shown in figure 1.1. There the blue points represented samples of the signal taken with a rate slower than double the frequency of the signal. Yet, it is possible to reconstruct a full period of the signal by applying a modulo operation on the x-values of the sampling points.

Chapter 2

State Tomography via Compressed Sensing

If we are talking about compressed sensing in a quantum setting [6], we are talking about the possibility to gain full information about the density matrix of an unknown state without having to measure the full set of observables which would be needed in full state tomography.

The application of compressed sensing in quantum mechanics is taking the ideas from classical signalling theory [4] and extends them to a ``signal" being a matrix, i.e. generalising compressed sensing from vectors to matrices. The main idea is that in general any matrix that is sparse in the sense that it is of low rank can be fully reconstructed if we know some entries of the matrix. Meaning that if we have measured some entries of our density matrix, we can use classical post-processing of this data to fully reconstruct the density matrix. The problem of reconstruction matrices from few entries is subject to the studies of Matrix completion [7, 8].

Intuitively it should be clear, that, since the (column) rank of a matrix is defined by the number of individually linear independent column vectors, rd elements of the matrix should be enough to uniquely describe it. Here r is the rank of the matrix and $d = 2^n$ is the number of rows in the matrix. Note that this technique wont overcome the exponential growth in measurement complexity, since rd still scales $O(2^n)$. It is provably impossible for any protocol that can handle generic pure states to overcome this exponential growth [6]. Also note that restricting a density matrix ρ to be of low rank is acceptable, since a lot of interesting quantum states are either pure or inhibit certain symmetries leading to a reduction of the rank of ρ. Pure states are always of rank 1, since they are defined by d coefficients if written in ket notation. Even with the presence of noise in the measurements a low rank density matrix is still a good approximation for a nearly pure state.

2.1 Gross, 2009

This section is dedicated to the work presented work presented by Gross et al. in [6] and [9]. It will give the reader an overview of the techniques involved in the proof and show how the method reduces the measurements necessary to determine a density matrix uniquely. The letter proves that $O(rd \log^2 d)$ measurements are enough to determine a density matrix ρ of rank r uniquely.

2.1.1 A Protocol for efficient Tomography

The protocol Gross et al. describe in [6] goes as follows. Consider the case of a n-qubit system. Any n-qubit Pauli matrix can be written as $w = \bigotimes_{i=1}^{n} w_i$ where $w_i \in \{\sigma_0, \sigma_1, \sigma_2, \sigma_3\}$. Since every w_i has 4 values it can take, there exist 4^n such matrices. Intuitively it should be clear that, a matrix-basis for $d \times d$ matrices needs at least that many elements since $d^2 = 2^{n2} = 4^n$. Indeed those matrices w form a basis for $d \times d$ density matrices. If we now draw m integers $A_1, \ldots, A_m$ from $[1, 4^n]$ we can measure the corresponding expectation value of the observable $w(A_i)$, i.e. $\mathrm{tr}(w(A_i)\rho)$, where $w(A_i)$ labels the A_ith matrix constructed as mentioned above. This information is then used to solve the convex optimisation problem

$$\begin{aligned} \min \quad & \|\sigma\|_{tr} \\ \text{subject to} \quad & \mathrm{tr}\,\sigma = 1 \\ \text{and} \quad & \mathrm{tr}(w(A_i)\sigma) = \mathrm{tr}(w(A_i)\rho), \end{aligned} \tag{2.1}$$

where σ is the best guess for ρ. Line 2 states that the trace of σ must be one, which is important for any density matrix and line 3 guarantees the overlap between the information we have about the density matrix ρ and the reconstruction σ.

Also, instead of minimising the trace norm $\|\sigma\|_{tr}$ one should actually minimise the rank of σ. However rank minimisation is a NP-hard problem, hence the trace norm acts as a surrogate to approximate low-rank in a matrix. Only the restriction on low rank matrices allows to set up the convex optimisation problem likes this.

The convex optimisation problem (2.1) can then be efficiently solved on a classical computer. Gross et al. claim that their reconstruction algorithm takes about one minute on a laptop computer for a 8-qubit state. Remember that this means calculating $4^8 = 65536$ matrix entries.

The main result of [6] is given as

Theorem 1. *Let ρ be an arbitrary state of rank r. If $m = cdr \log^2 d$ randomly chosen Pauli expectations are known, then ρ can be uniquely reconstructed by solving the convex optimisation problem* (2.1) *with probability of failure exponentially small in c.*

To proof this Gross et al. utilise a ``sampling operator''

$$\mathcal{R} : \rho \to \frac{d}{m} \sum_{i=1}^{m} w(A_i) \;\; \mathrm{tr}\left(\rho w(A_i)\right). \tag{2.2}$$

Note that the sampling operator is essentially a matrix valued projection. It is exactly a matrix valued projection if $A_i \neq A_j \; \forall i \neq j$ (compare equation (1.6)). With that we can rewrite the convex optimisation problem (2.1) to be

$$\min \quad \|\sigma\|_{tr} \tag{2.3}$$

$$\text{subject to} \quad \mathcal{R}\sigma = \mathcal{R}\rho. \tag{2.4}$$

It can be easily seen that σ is a unique solution for ρ if for all deviations $\Delta := \sigma - \rho$, Δ is either infeasible

$$\mathcal{R}\Delta \neq 0, \tag{2.5}$$

meaning the coefficients $\mathrm{tr}((A_i)\sigma)$ do not agree with the information we have learned from the measurements of ρ, or

$$\|\rho + \Delta\|_{tr} > \|\rho\|_{tr}, \tag{2.6}$$

meaning σ is not a solution to the convex optimisation problem, since the deviation increases the trace norm (which we are minimising).

A tool from convex optimisation becomes handy in the proof of theorem 1: the construction of a strict subgradient Y for the norm serves as a so called dual certificate [10]. Y is called a strict subgradient if

$$\|\rho + \Delta\|_{tr} > \|\rho\|_{tr} + \mathrm{tr}(Y\Delta) \quad \forall \Delta \neq 0. \tag{2.7}$$

Meaning that if we find such a Y which is also in the range of $\mathcal{R}$ and since $\mathcal{R}\Delta = 0$, it follows that $\mathrm{tr}(Y\Delta) = 0$. This fact is illustrated in figure 2.1. If we can find a matrix

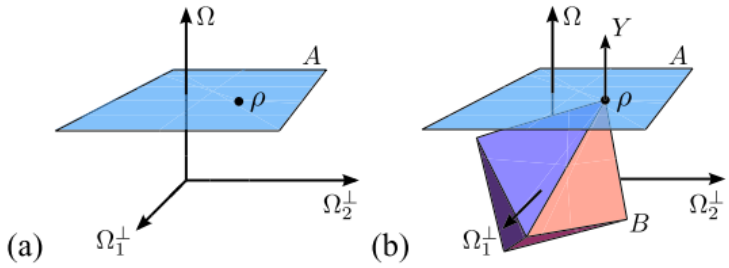

Figure 2.1: The matrix space our optimisation problem is living on. In (a) we can see an affine hyperplane A which is defined by the expectation values tr $(\rho w(A_i))$ obtained by measurements. The axis Ω defines the m known parameters. (b) Illustrates the fact that we recover the density matrix ρ uniquely, if A is an supporting hyperplane of the space $B = \{\sigma | \|\sigma\|_{tr} \leq \|\rho\|_{tr}\}$ at ρ. In other words we must find a matrix Y normal to a supporting hyperplane of B and A. Picture taken from [9].

Y that is normal to a supporting hyperplane of the space of matrices with lesser trace norm than ρ and is also normal to an affine space given by the m measured expectation values. We find that tr $\Delta Y = 0$ since Δ lies in the affine space.

For the further proceeding of the proof it is useful to define the vector space $U = \mathrm{range}(\rho)$, i.e. the column space of ρ and further the space $T = \{\sigma \mid (I - P_U)\sigma(I - P_U) = 0\}$, where P_U is the projector onto space U. Then the matrix valued projector $\mathcal{P}_T$ (projection into T) can be used to decompose Δ into $\Delta_T + \Delta_T^{\perp}$, denoting the parts of Δ living in T and in the space orthogonal to T, respectively. This is useful since as shown

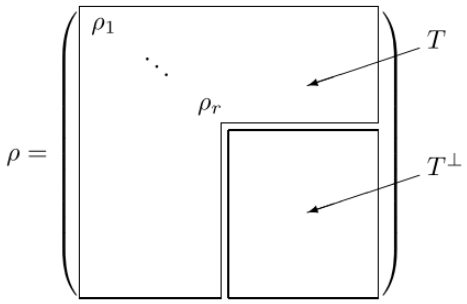

Figure 2.2: The decomposition into the spaces T and $T^{\perp}$. From the construction of T it can be seen, that if we write ρ with respect to its eigenbasis, all r eigenvalues ρ_i are living on the matrix space T meaning that any deviation $\Delta_T^{\perp}$ is decoupled from ρ. Picture taken from [9].

in figure 2.2 this construction helps to separate the space on which the r eigenvalues of our density operator are defined on, i.e. the space on which its diagonal entries are non-zero if represented with respect to its eigenbasis, from the space orthogonal to that. In particular we can distinguish two cases for Δ: case (i) where $\|\Delta_T\|_2 > d^2\|\Delta_T^{\perp}\|_2$ and case (ii) where $\|\Delta_T\|_2 \leq d^2\|\Delta_T^{\perp}\|_2$.

In case (i) the proof is relatively easy written down, but involves a very powerful and complex technique, a large deviation bound, whose introduction lies beyond the scope of this report. The interested reader may refer to section II-D in [9]. In this case Δ is mostly dominated by Δ_T and all we need to show is that the restriction $\mathcal{A} := \mathcal{P}_T\mathcal{R}\mathcal{P}_T$ of $\mathcal{R}$ on T is invertible [9]. Using the large deviation bound we can show that

$$\Pr\left(\|\mathcal{A} - I_T\| > t\right) < 4dre^{-t^2\kappa/8}, \tag{2.8}$$

where $\kappa = m/(dr)$. We are only giving a short outline on how the proof for the above statement works. If we define a family of operators

$$Z(A_i) = \frac{d^2}{m}\mathcal{P}_T\mathcal{P}_{w(A_i)}\mathcal{P}_T, \tag{2.9}$$

where $\mathcal{P}_{w(A_i)}$ is the matrix valued projection on to $w(A_i)$, it is easy to see that

$$\mathcal{P}_T\mathcal{R}\mathcal{P}_T = \sum_{i=1}^{m} Z(A_i), \tag{2.10}$$

and $\mathbb{E}(Z(A_i)) = I_T$. Hence we can write

$$\mathcal{P}_T\mathcal{R}\mathcal{P}_T - I_T = \sum_{i=1}^{m} Z(A_i) - \mathbb{E}(Z(A_i)) \tag{2.11}$$

and derive an operator valued Bernstein inequality for the random variables $X(A_i) = Z(A_i) - \mathbb{E}(Z(A_i))$.

Equation (2.8) states that the probability of $\|\mathcal{A}-I_T\| > \frac{1}{2}$ is smaller than $4dre^{-\kappa/32} := p_1$. It is thus with high probability that $\|\mathcal{R}\Delta\|_2 > 0$, and hence Δ is infeasible.

The proof for case (ii) is more elaborate and uses the dual certificate mentioned before. In [9] it shown that

$$\|\Delta_T\|_2 < d^2\|\Delta_T^{\perp}\|_2 \tag{2.12}$$

together with

$$\Delta \in \text{range}(\mathcal{R}^{\perp}) \tag{2.13}$$

yields $\|\rho + \Delta\|_{tr} > \|\rho\|_{tr}$ (see equation (2.7)) and hence Δ does not fullfill the minimisation condition. This however does only hold if we can find such a dual certificate Y. So what remains is to prove the existence of a dual certificate Y.

We call the matrix $Y \in \text{span}(w(A_1), \ldots, w(A_m))$ an almost subgradient if

$$\|\mathcal{P}_T Y - P_U\|_2 \leq 1/(2d^2), \qquad \|\mathcal{P}_T^{\perp} Y\| < 1/2. \tag{2.14}$$

Gross et al. developed a recursive process, the ``golfing scheme", which converges exponentially fast to such a dual certificate and hence completes the proof of theorem 1.

If we draw l sets of $\kappa_0 rd$ Pauli observables, the process is written down as follows. Define $X_0 = P_U$,

$$Y_i = \sum_{j=1}^{i} \mathcal{R}_j X_{j-1}, \qquad X_i = P_U - \mathcal{P}_T Y_i, \tag{2.15}$$

$Y = Y_l$, where $\mathcal{R}_j$ is associated to the jth batch of Pauli matrices drawn. As before we assume again that $\mathcal{A}_j$ is the restriction of $\mathcal{R}_j$ on T and $\|\mathcal{A}_j - I_T\|_2 < \frac{1}{2}$. We also denote the probability that the latter doesn't hold as p_2. We find

$$\begin{aligned} \|X_i\|_2 &= \|X_{i-1} - \mathcal{P}_T\mathcal{R}_i X_{i-1}\|_2 & (2.16)\\ &= \|(I_T - \mathcal{A}_i)X_{i-1}\|_2 \leq 1/2\|X_{i-1}\|_2, & (2.17) \end{aligned}$$

it follows from the first and last term above that $\|X_i\|_2 \leq 2^{-i}\|X_0\|_2 = 2^{-i}\sqrt{r}$. If $l \geq \log_2(2d^2\sqrt{r})$, we fullfill the first part of equation (2.14). Again using the large-deviation technique it is shown in [6] that $\|\mathcal{P}_T^{\perp}\mathcal{R}_{\rangle}X_{i-1}\| \leq 1/(4\sqrt{r})\|X_{i-1}\|_2$ with high probability $(1-p_3)$. Hence,

$$\|\mathcal{P}_T^{\perp}Y_l\| \leq \sum_{j=1}^{l}\|\mathcal{P}_T^{\perp}\mathcal{R}_j X_{j-1}\| \leq \frac{1}{4}\sum_{j=0}^{\infty}2^{-l} < \frac{1}{2}, \tag{2.18}$$

which is the second half of equation (2.14).

In a last step Gross et al. estimate their probability of failure $p_f \leq p_1+p_2+p_3$. If we want a exponential small failure probability $p_f e^{-\mu}$ we have to set $\kappa_0 = 64\mu(1+ln(8dl))$, which means that $m = dr(ln\ d)^2 O(1)$ Pauli expectation values have to be sampled. □

2.1.2 Further Remarks

In [6] Gross et al. make further observations on their results. They state that methods are robust against noisy measurements and the possibility that the true state ρ_t is only approximated by a low rank matrix ρ ($\|\rho_t - \rho\| < \varepsilon_1$). Noisy measurements would manifest in a way such that the expectation values are not $\mathrm{tr}(\rho_t w(A_i))$ but $\mathrm{tr}(\omega w(A_i))$ for some matrix ω. Their altered protocol reads as follows: choose some $\lambda \geq 1$ and $\varepsilon \geq \lambda(\sqrt{d^2/m}\varepsilon_1+\varepsilon_2)$, where $\varepsilon_2 \geq \|\mathcal{R}\omega-\mathcal{R}\rho_t\|_2$ is the deviation of the measurements from the true state. Note that ε_2 can be estimated from statistics on the measurement process. Now solve the convex optimisation problem

$$\min \|\sigma\|_{tr} \quad \text{subject to} \|\mathcal{R}\sigma - \mathcal{R}\omega\|_2 \leq \varepsilon. \tag{2.19}$$

Furthermore the authors state that their method of tomography can be certified in the sense that one can find a certificate (assuming the unknown state is pure) the reconstruction error from $O(cd\ \log^2 d)$ Pauli expectation values, where the probability of failure is exponentially small in c.

On a last remark [6] improves the classical post processing by choosing the elements of the Pauli basis, which we measure the expectation values for, in a more structured way. This speeds up the classical processing time, however for the cost of precision in the reconstruction. This approach is called a hybrid method because it is equivalent to a certain structured matrix completion problem. A comparison between both methods (random sampling and hybrid) can be seen in figure 2.3. It is obviously that both methods approach good values for the fidelity of the reconstruction very quickly with increasing sparsity. However, especially in the plot for trace norm the decreased accuracy for the hybrid methods becomes visible.

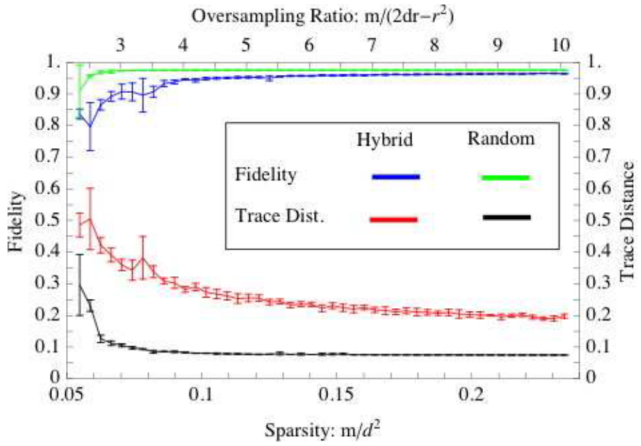

Figure 2.3: Numerical simulation of the reconstruction process presented in [6]. It can be seen that the fidelity of the reconstruction quickly approaches high values of fidelity ($> 95\%$) with increasing sparsity, while the trace distance decreases. In the plot for the trace norm the lesser accuracy for the hybrid method can be clearly seen. The data used to test the methods was produced numerically from a simulated 8-qubit state, where depolarising and statistical noise were added.

2.2 Flammia, 2012

In 2012 Flammia together with Gross et al. refined their previous result by utilising a restricted isometry property (RIP) for low rank matrices in the analysis of the tomography protocol [11]. The construction of the matrix basis $w(A_i)$ used to determined the m expectation values is the same as before in [6], however this letter aims to make better predictions about the precision needed for each measurement. Hence in the analysis it is assumed that t identical copies of the quantum state are available. So every measurement setting can be carried out t/m times. This approach gives better estimators for the total number of needed measurement, contrary to before were every expectation value was assumed to be known with perfection (except for during the remarks on noisy measurements at the end of [6]).

Also the definition of the ``sampling operator'' is somewhat different, $\mathcal{A}$ is given as a linear map $\mathcal{A} : \mathbb{H}^d \to \mathbb{R}^m$ defined for all $i \in [1, m]$

$$(\mathcal{A}(\rho))_i = \sqrt{\frac{d}{m}}\, tr\,(w(A_i)\rho). \tag{2.20}$$

This now allows us to write the output of the sampling operation as a vector

$$y = \mathcal{A}(\rho) + z, \tag{2.21}$$

where z contains all the statistical noise in the measurement and the vector y contains all the coefficients as a list. Note that those ``known'' coefficients are exactly the same as before.

The paper goes on and analyses two different estimators for reconstructing ρ. The first estimator is as before a trace norm minimisation called Dantzig selector

$$\hat{\rho}_{DS} = \arg\min_X \|X\|_{tr}, \quad \text{subject to } \|\mathcal{A}^*(\mathcal{A}(X) - y)\| \leq \lambda, \tag{2.22}$$

where $\mathcal{A}^*$ is the adjoint of $\mathcal{A}$ and λ has to be chosen according to the noise in the system, similar to ε in equation (2.19). The other estimator is obtained by a least-squares

linear regression with regularisation called matrix Lasso (least absolute shrinkage and selection operator)

$$\hat{\rho}_{Lasso} = \arg\min_X \frac{1}{2}\|\mathcal{A}(X) - y\|_2^2 + \mu\|X\|_{tr}, \tag{2.23}$$

where again μ is chosen according to the noise present in the measurements.

As mentioned before the analysis of the protocols in this thesis makes use of the RIP. We say that the sampling operator complies with the RIP if there exists a constant $0 \leq \delta_r < 1$ for all $X \in \mathbb{C}^{d\times d}$ with

$$(1-\delta_r)\|X\|_F \leq \|\mathcal{A}(X)\|_2 \leq (1+\delta_r)\|X\|_F. \tag{2.24}$$

This is to say that the space of low-rank matrices can be embedded in a lower dimensional manifold, which is the action of $\mathcal{A}$, in a way such that the norm is only slightly distorted. In a different article it is shown that the sampling operator satisfies the RIP [12].

Combining the RIP of the sampling operator together with results from other sources [8, 13] Flammia et al. are able to construct a proof for the following theorem, which wont be presented in this report. In the following C, C_0, C_1, C_0', and C_1' are fixed constants. ρ_c is part of the decomposition of $\rho = \rho_r + \rho_c$, where ρ_r is the best rank r approximation of ρ and ρ_c is the residual.

Theorem 2. *Let $\mathcal{A}$ be the random Pauli sampling operator (2.20) with $m \geq Crd \log^6 d$. Then, with high probability, the follwing holds:*

Let $\hat{\rho}_{DS}$ be the matrix Danzig selector (2.22), and choose λ so that $\|\mathcal{A}^*(z)\| \leq \lambda$. Then

$$\|\hat{\rho}_{DS} - \rho\|_{tr} \leq C_0 r\lambda + C_1\|\rho_c\|_{tr}. \tag{2.25}$$

Alternatively, let $\hat{\rho}_{Lasso}$ be the matrix Lasso (2.23), and choose μ so that $\|\mathcal{A}^*(z)\| \leq \mu/2$. Then

$$\|\hat{\rho}_{Lasso} - \rho\|_{tr} \leq C_0' r\mu + C_1'\|\rho_c\|_{tr}. \tag{2.26}$$

Note that the first terms on the right side in both inequalities in the theorem above depend on the noise z which depends on the number of measurements available for each Pauli expectation value. Also note the lower bound for m given above is a poly $\log d$ factor bigger then for the result using a dual certificate, however the error bounds given in equations (2.25) and (2.26) are tighter.

2.2.1 Sample Complexity

The more interesting result of the paper, which discussed in this section, are the consideration of the total sample complexity t. Since it is an import figure of merit in an actual experimental setting. Because t represents the number of total measurements needed to reconstruct the density matrix, it relates directly to the time needed to carry out the tomography protocol.

Theorem 3. *Given $t = O((\frac{rd}{\varepsilon})^2 \log d)$ copies of ρ and used to measure different Pauli expectation values, then the following holds with high probability over the measurement outcomes:*

Let $\hat{\rho}_{DS}$ be the matrix Dantzig selector (2.22), and set $\lambda = \varepsilon/(C_0 r)$ for some $\varepsilon > 0$. Then

$$\|\hat{\rho}_{DS} - \rho\|_{tr} \leq \varepsilon + C_1\|\rho_c\|_{tr}. \tag{2.27}$$

Alternatively, let $\hat{\rho}_{Lasso}$ be the matrix lasso (2.23), and set $\mu = \varepsilon/(C_0' r)$. Then

$$\|\hat{\rho}_{Lasso} - \rho\|_{tr} \leq \varepsilon + C_1'\|\rho_c\|_{tr}. \tag{2.28}$$

Proof. Let's assume we have t copies of a quantum state and fix our choice of Pauli observables we want to measure, i. e. we choose $A_1, \dots, A_m$ numbers. Now we define B_{ij} to be the outcome of the jth measurement of the ith observable, hence $j \in [1, \frac{t}{m}]$, $i \in [1, m]$, $B_{ij} \in \{1, -1\}$ and $\mathbb{E}B_{ij} = \mathrm{tr}(w(A_i)\rho)$. The vector y can then be written as

$$y_i = \sqrt{\frac{d}{m}} \frac{m}{t} \sum_{j=1}^{t/m} B_{ij}, \quad \mathbb{E}y = \mathcal{A}(\rho). \tag{2.29}$$

If we want to bound $\|\mathcal{A}^*(y - \mathcal{A}(\rho))\|$, we can use the Bernstein inequality. First write

$$\mathcal{A}^*(y) = \sqrt{\frac{d}{m}} \sum_{i=1}^{m} w(A_i) y_i = \frac{d}{t} \sum_{i=1}^{m} \sum_{j=1}^{t/m} w(A_i) B_{ij} \tag{2.30}$$

and

$$\mathcal{A}^*\mathcal{A}(\rho) = \frac{d}{m} \sum_{i=1}^{m} w(A_i)\, \mathrm{tr}(w(A_i)\rho). \tag{2.31}$$

Then express $\mathcal{A}^*(y - \mathcal{A}(\rho))$ as sum of independent random variables

$$\mathcal{A}^*(y - \mathcal{A}(\rho)) = \sum_{i=1}^{m} \sum_{j=1}^{t/m} X_{ij}, \quad X_{ij} = \frac{d}{t} w(A_i)[B_{ij} - \mathrm{tr}(w(A_i)\rho)]. \tag{2.32}$$

It is notable that $\mathbb{E}X_{ij} = 0$ and $\|X_{ij}\| \leq 2d/t =: R$. For the second moment it holds

$$\mathbb{E}(X_{ij}^2) = \mathbb{E}(\frac{d^2}{t^2} I[B_{ij} - \mathrm{tr}(w(A_i)\rho)]^2) \tag{2.33}$$

and hence

$$\sigma^2 := \left\| \sum_{ij} \mathbb{E}(X_{ij}^2) \right\| = \sum_{ij} \frac{d^2}{t^2} [1 - \mathrm{tr}(w(A_i)\rho)^2]) \tag{2.34}$$

$$\leq t \cdot \frac{d^2}{t^2} = \frac{d^2}{t}. \tag{2.35}$$

Now using the matrix-valued Bernstein inequality we find

$$\Pr\left[\|\mathcal{A}^*(y - \mathcal{A}(\rho)\| \geq \lambda\right] \leq d \cdot \exp\left(-\frac{\lambda^2/2}{\sigma^2 + (R\lambda/3)}\right) \leq d \cdot \exp\left(-\frac{t\lambda^2/2}{d(d+1)}\right). \tag{2.36}$$

Set $\lambda = \varepsilon/(C_0 r)$ for the Dantzig selector and we conclude for any $t \geq 2C_4\lambda^{-2}d(d+1)\log d = 2C_4(C_0 r/\varepsilon)^2 d(d+1)\log d$,

$$\Pr\left[\|\mathcal{A}^*(y - \mathcal{A}(\rho)\| \geq \frac{\varepsilon}{C_0 r}\right] \leq d \cdot \exp\left(-C_4 \log d\right) = d^{1-C_4}. \tag{2.37}$$

Applying theorem 2 completes the proof! □

The paper then goes a long way to proof the tightness of the bound shown above, which wont be covered here.

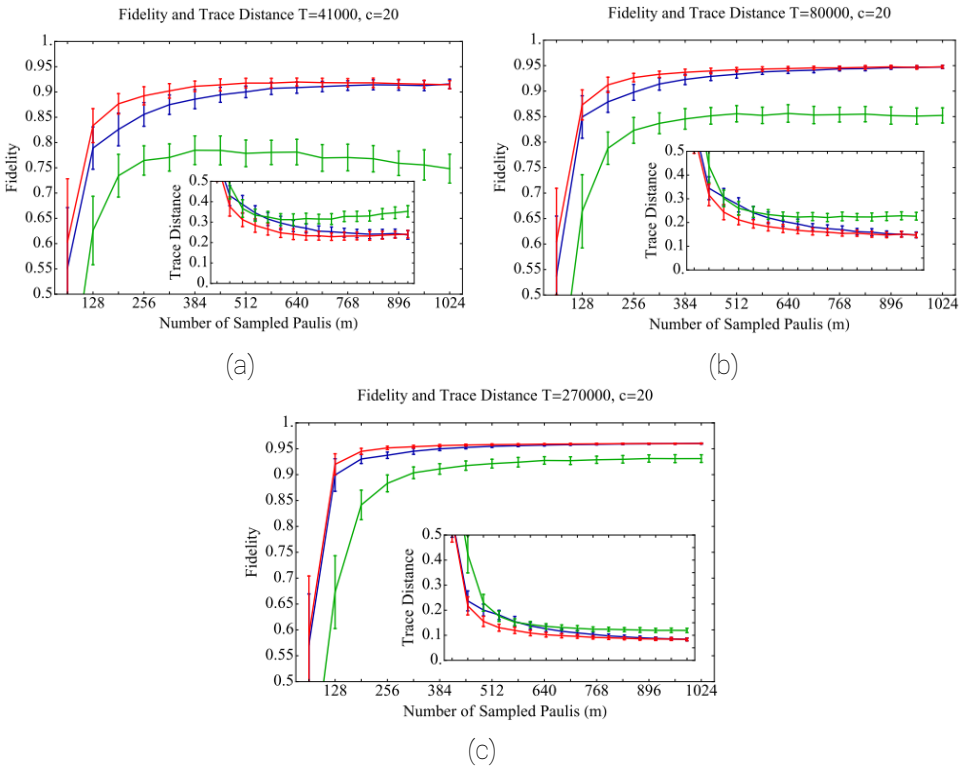

Figure 2.4: The result for the numerical simulations in [11]. The total time T on which the simulated measurements are taken is increased between sub-figures 2.4a and 2.4c. The time needed to switch between two measurement settings c was assumed to be 20, i.e. 20 times longer then one measurement itself. In every plot the matrix Lasso (red) and the Dantzig selector (blue) outperform the maximum likelihood estimator (green) given as a benchmark. More surprising is that as soon as a high fidelity is reached a further increasing of m results in approximately the same fidelity. As a consequence measuring fewer Pauli expectation values with higher precision is the favourable technique, since the classical algorithms scale with m. Picture taken from [11].

2.2.2 Numerical Simulations

As before in [6] Flammia et al. tested their results with numerical simulations. However to simulate a realistic lab situation they assumed a time c it takes to change the measurement setting. Hence they were able to address the question whether it is more favourable to measure one Pauli expectation value many times, resulting in lesser noise, or whether it is better to have noisier results for more Pauli expectation values. In both cases the process was restricted to a time T. The results can be found in figure 2.4.

Finally the paper ends by mentioning the application of compressed sensing methods in quantum process tomography, which is another interesting field where resources typically scale exponentially.

2.3 Schwemmer, 2014

A good example for the usefulness of compressed sensing is given in [14]. There the authors compare different types of tomography protocols considering a permutationally invariant 6-photon Dicke state. It's only thanks to the extremely high pump power ($8.4\ W$) used in this cavity-enhanced spontaneous parametric down conversion (SPDC) experiment that a full state tomography could be carried out and used as a reference for different state reconstruction methods.

Note that the number of parameters needed to be determined is $6^4 - 1 = 4095$. However due to the symmetries in the Dicke states

$$|D_N^{(n)}\rangle = \binom{N}{n}^{-1/2} \sum_i \mathcal{P}_i(|H^{\otimes(N-n)}\rangle \otimes |V^{\otimes n}\rangle), \tag{2.38}$$

where $\mathcal{P}_i$ denotes all possible permutations of the qubits, $|H/V\rangle_i$ denotes the polarisation of the ith qubit, N is the total photon number and n is the number of photons in the $|V\rangle$ state. The authors go beyond compressed sensing schemes, in particular using the permutationally invariance (PI) of the states to develop even more effective methods reaching polynomial scaling in the number of qubits. This is not surprising since compressed sensing is capable of dealing with any generic quantum states whereas the restriction on Dicke states naturally allows for further reductions of the parameters needed to describe the state uniquely. A PI state is defined by $\binom{N+3}{N} = O(N^3)$ parameters.

In the experiment all $3^6 = 729$ Pauli settings where measured over 50 hours. Note that the identity observable doesn't have to be measured since it drops out of any other Pauli measurement. Randomly chosen subsets of this data was then used to evaluate the different tomography methods.

A comparison of the different methods tested can be seen in table 2.1. Firstly, note that the state produced in the 6-photon experiment mainly is a mixture states $|D_6^{(2)}\rangle$, $|D_6^{(3)}\rangle$ and $|D_6^{(4)}\rangle$, hence is suitable for compressed sensing (low rank). The first column in table 2.1 represents the overlaps of the measured data with the respective Dicke states obtained by full tomography and serves as a reference. Taking the permutational symmetry into account the second column of data was obtained by only 28 measurement settings, still giving good agreement with the full state tomography. The third column is the result of a compressed sensing scheme in which only 270 expectation values from the full state tomography were taken into account. Finally the last row combined both techniques and needed only 16 measurement settings to reconstruct the quantum state with good fidelity.

State	Full	PI	CS	PI, CS
$\|D_6^{(0)}\rangle$	0.001	0.001	0.001	0.002
$\|D_6^{(1)}\rangle$	0.005	0.008	0.011	0.006
$\|D_6^{(2)}\rangle$	0.197	0.222	0.181	0.207
$\|D_6^{(3)}\rangle$	0.604	0.590	0.615	0.592
$\|D_6^{(4)}\rangle$	0.122	0.127	0.118	0.119
$\|D_6^{(5)}\rangle$	0.003	0.004	0.003	0.005
$\|D_6^{(6)}\rangle$	0.000	0.003	0.001	0.004
$\sum$	0.933	0.954	0.929	0.935

Table 2.1: The fidelities compared to the different Dicke states measured with full tomography, PI tomography, compressed sensing from $\mathbf{270}$ measurement settings and PI tomography combined with compressed sensing, respectively.

The paper is a good demonstration of the power of compressed sensing schemes but also a reminder that if we can assume more information about the state we want to determine (on top of low rank), it is possible to find even more efficient schemes then compressed sensing. Especially in state preparation problems, where we want to certify the quality of a source of known states, additional assumption about the output state are often valid.

Chapter 3

Conclusion

In this report we summarised the work done in the field of (quantum) compressed sensing in the group of David Gross. They were able to show the feasibility of compressed sensing methods in the framework of quantum information. Although compressed sensing can never overcome the exponential scaling in terms of measurement effort, the square root improvement for low rank states is a very powerful tool to reduce measurement costs and enables experimentalists to carry out more complex studies with more qubits, which were not possible before.

The second paper [11] presented in this report gives especially good bounds on the accuracy of their protocol and makes precise predictions on the actual time needed to carry out the measurements needed for a successful reconstruction of the density matrix.

The last paper presented here [14] is a good example of the application of compressed sensing in a real experimental setting, yet, it also reminds the reader that it is often feasible to go beyond compressed sensing and make further assumption on the state which we want to determine. By doing so Schwemmer et al. showed a dramatic decrease in the number of needed measurements and even achieved polynomial scaling. The original work on permutationally invariant quantum tomography can be found in [15], where the polynomial scaling for such states was presented for the first time.

We want to conclude by mentioning an idea suggested by the author. The work carried out on compressed sensing so far mainly focuses on efficient tomography schemes. Yet, intuitively, compressed sensing could find a much broader application. Consider an algorithm simulating quantum mechanics on a classical computer. As we know, such a algorithm also scales exponentially bad comparable to the measurement effort in full tomography. However if we could design an algorithm working on a quadratically smaller space of dimension rd a quadratic speed up might be achieved while maintaining a good fidelity, enabling classical simulations with many more qubits than possible today.

Bibliography

[1] X.-C. Yao, T.-X. Wang, P. Xu, H. Lu, G.-S. Pan, X.-H. Bao, C.-Z. Peng, C.-Y. Lu, Y.-A. Chen & J.-W. Pan. Observation of eight-photon entanglement. *Nature Photonics* **6**, 225 (2012)

[2] Y.-F. Huang, B.-H. Liu, L. Peng, Y.-H. Li, L. Li, C.-F. Li & G.-C. Guo. Experimental generation of an eight-photon Greenberger-Horne-Zeilinger state. *Nature communications* **2**, 546 (2011)

[3] T. Monz, P. Schindler, J. T. Barreiro, M. Chwalla, D. Nigg, W. a. Coish, M. Harlander, W. Hänsel, M. Hennrich & R. Blatt. 14-Qubit Entanglement: Creation and Coherence. *Physical Review Letters* **106**, 130506 (2011)

[4] D. Donoho. Compressed sensing. *IEEE Transactions on Information Theory* **52**, 1289 (2006)

[5] E. J. Candès, J. K. Romberg & T. Tao. Stable Signal Recovery from Incomplete and Inaccurate Measurements. *Communications on Pure and Applied Mathematics* **LIX**, 1207 (2006)

[6] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker & J. Eisert. Quantum state tomography via compressed sensing. *arXiv.org* **quant-ph**, 4 (2009)

[7] E. J. Candès & B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics* **9**, 717 (2009)

[8] E. J. Candes & Y. Plan. Tight oracle bounds for low-rank matrices recovery from a minimal number of random measurements. *IEE TRANS. Inform. Theory* **57**, 2342 (2011)

[9] D. Gross. Recovering low-rank matrices from few coefficients in any basis. *arXiv.org* pages 1--19 (2009)

[10] D. P. Bertsekas, A. Nedi & A. E. Ozdaglar. *Convex analysis and optimization*. Athena Scientific (2003)

[11] S. T. Flammia, D. Gross, Y.-K. Liu & J. Eisert. Quantum Tomography via Compressed Sensing: Error Bounds, Sample Complexity, and Efficient Estimators. *arXiv.org* **quant-oh**, 16 (2012)

[12] Y.-K. Liu. Universal low-rank matrix recovery from Pauli measurements page 14 (2011)

[13] M. Fazel, E. Candès, B. Recht & P. Parrilo. Compressed sensing and robust recovery of low rank matrices. *Conference Record - Asilomar Conference on Signals, Systems and Computers* pages 1043--1047 (2008)

[14] C. Schwemmer, G. Tóth, A. Niggebaum, T. Moroder, D. Gross, O. Gühne & H. Weinfurter. Experimental Comparison of Efficient Tomography Schemes for a Six-Qubit State. *Physical Review Letters* **113**, 040503 (2014)

[15] G. Tóth, W. Wieczorek, D. Gross, R. Krischek, C. Schwemmer & H. Weinfurter. Permutationally Invariant Quantum Tomography. *Physical Review Letters* **105**, 250403 (2010)