DC Electrode Control of a Linear Paul Trap for Quantum Computing with Trapped Ions

Brennan MacDonald-de Neeve
Supervised by Vlad Negnevitsky

Supervising Professor
Prof. Jonathan Home
Trapped Ion Quantum Information Group
Department of physics, ETH Zürich

June 2016

Abstract

Improvements to the control system of a linear Paul trap's DC electrodes are presented. In particular, design modifications made to the Direct Ethernet Adjustable Transport Hardware (DEATH) PCB designed by former Ph.D. student Ludwig de Clercq are presented in detail. Potential pitfalls in manufacturing are included, and limitations of the control system are investigated and explained. The experimental system was updated with the addition of these new PCBs, thereby enabling control of individual DC trap electrodes, where previously opposing pairs would share a single voltage. This added capability permits better characterization of the ion trap, and better control and correction of experimental errors. Finally, an experiment undertaken to characterize the ion trap's performance by isolating the effect of individual electrodes is explained. The results of this experimental are presented, revealing the malfunction of a trap electrode, and a possible gain error in simulated results.

Contents

1 Introduction

One of the major challenges in trapped ion quantum information is to scale today's systems of several qubits to larger systems containing several thousand [1]. Due to technical difficulties, scaling of currently used ion traps is limited to tens of ions. This limitation has led to the conception of a quantum charge coupled device (QCCD) architecture in the last decade [2]. The ion trap used in the Trapped Ion Quantum Information (TIQI) group at ETH Zürich can also be thought of as a simple QCCD providing more flexibility than a traditional Paul trap. In this report, work done to improve the currently used ion trap control system is presented.

1.1 Trapped Ion Qubit

The Paul trap used by the segmented trap experiment in the TIQI group is capable of trapping both $^{40}Ca^{+}$ and $^{9}Be^{+}$ ions. For experiments performed in this semester project $^{40}Ca^{+}$ was used. The relevant levels and transition wavelengths for $^{40}Ca^{+}$ experiments are shown in figure 1. The ground state $|\downarrow\rangle = |S_{1/2}, m = +1/2\rangle$ and metastable excited state $|\uparrow\rangle = |D_{5/2}, m = +3/2\rangle$ are used for the so called optical qubit in the presence of a magnetic field of $\simeq 119$ Gauss. Initialization is performed by optical pumping using a 397 nm sigma laser to couple $|\mathrm{S}_{1/2}, m = -1/2\rangle$ to the rapidly decaying $|\mathrm{P}_{1/2}, m = +1/2\rangle$ state. If any population is in the $D_{5/2}$ manifold, a broadband laser at 854 nm is used to simultaneously couple the D states to the rapidly decaying $\mathrm{P}_{3/2}$ states. Similarly, if any population is in the $D_{3/2}$ manifold, an 866 nm laser is used to couple with the rapidly decaying $\mathrm{P}_{1/2}$ states. The 397 nm transition between the $|\downarrow\rangle$ state and the $\mathrm{P}_{1/2}$ states is used both for Doppler cooling and qubit state readout.

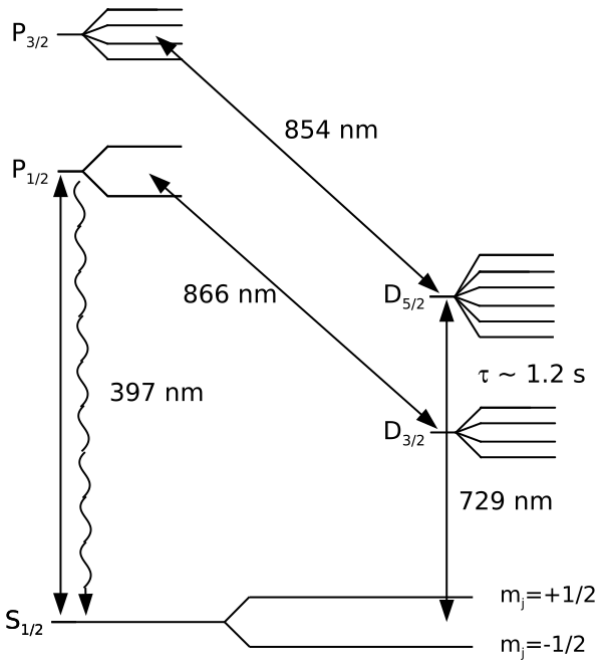

Figure 1: Relevant energy levels and transition wavelengths for the $^{40}Ca^{+}$ ion [3].

1.2 Ion Trap Physics

To use an ion as a qubit, it is isolated in a vacuum and trapped using electrostatic and electrodynamic forces. Since Laplace's equation in 3 dimensions admits no local maxima or minima, it is in fact not possible to trap a charged particle using only electrostatic forces [1]. It is however possible to trap ions using RF fields. A linear Paul trap geometry is used in the experimental set-up at ETH Zürich, for which the basic arrangement of the RF and DC electrodes is shown in figure 2.

Here I follow the analysis of D. Leibfried *et al.* [4]. Assuming a potential of approximately quadrupolar spatial shape at the centre of the trapping region, the potential at the location of the ion can be written to a good approximation as

$$\Phi(x,y,z,t) = U\frac{1}{2}(\alpha x^2 + \beta y^2 + \gamma z^2) + \tilde{U}\cos(\omega_{rf}t)\frac{1}{2}(\alpha' x^2 + \beta' y^2 + \gamma' z^2)$$

where the RF and DC contributions are separated to simplify the analysis. The Laplace equation $\nabla^2\Phi = 0$ imposes the restrictions

$$\alpha + \beta + \gamma = 0$$

$$\alpha' + \beta' + \gamma' = 0.$$

In the case of the linear Paul trap this is satisfied by

$$-(\alpha + \beta) = \gamma > 0,$$

$$\alpha' = -\beta',$$

such that the ions are confined dynamically in the x,y plane as will be explained below, and confined by a static potential along the z axis. The classical equation of motion for a $^{40}Ca^{+}$ ion in the x direction is given by

$$\ddot{x} = \frac{e}{m}\frac{\partial\Phi}{\partial x}.$$

Expanding the above results in the standard form of the Mathieu differential equation

$$\frac{d^2x}{d\xi^2} + [a_x - 2q_x\cos(2\xi)]x = 0$$

with

$$\xi = \frac{\omega_{rf}t}{2}, a_x = \frac{4eU\alpha}{m\omega_{rf}^2}, q_x = \frac{2e\tilde{U}\alpha'}{m\omega_{rf}^2}.$$

The lowest-order solution provides a good enough description for most of the physics of interest relevant for quantum information, and is given by

$$x(t) \approx 2A\cos\left(\beta_x\frac{\omega_{rf}}{2}t\right)\left[1 - \frac{q_x}{2}\cos(\omega_{rf}t)\right]$$

with

$$\beta_x \approx \sqrt{a_x + q_x^2/2},$$

[1] This fact formally goes by the name of Earnshaw's Theorem

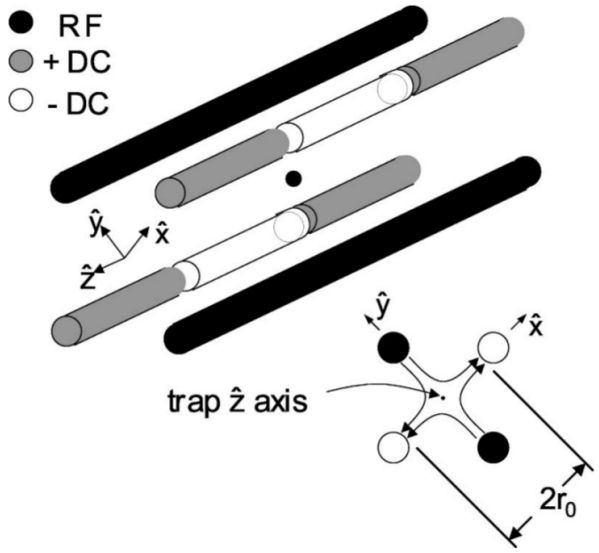

Figure 2: Paul trap geometry [4].

where A is an arbitrary constant to be determined from boundary conditions.

This solution features two distinct oscillatory motions of the ion, one at the secular frequency $\nu = \beta_x \omega_{rf}/2$ typically much less than the RF drive frequency, and another faster motion at the RF drive frequency but with much lower amplitude referred to as the micromotion. For the case of a Paul trap, an analogous derivation can be made for the y direction, while the analysis in the z direction is simplified by the absence of any dynamic potential. Thus the ion is confined dynamically in the x,y plane where to a good approximation it experiences effective harmonic potentials at the two secular frequencies, referred to collectively as the pseudopotential. Confinement can be obtained in the z direction by applying positive and negative (relative) voltages to the DC electrodes as illustrated in figure 2.

In typical trapped ion quantum information experiments, the harmonic motion of the ions is well described as a quantum harmonic oscillator. Under the appropriate conditions, this quantum oscillator can be coupled to the internal qubit states of the ion using laser light. Incident light which is red-detuned by an amount equal to one of the motional frequencies of the ion [2] from the qubit transition frequency leads to a qubit transition accompanied by the loss of a quantum of motion. Similarly, light blue-detuned from the qubit frequency by the appropriate amount leads to a qubit transition accompanied by an increase of one quantum of motion. These are referred to as the red and blue sideband transitions, respectively, and are illustrated in figure 3. In the experiments discussed in this report, the harmonic motion along the z axis (axial) is typically used for sideband transitions as opposed to those motions in the x,y plane (radial).

When multiple ions of the same mass are placed in the trap, their motions are coupled by the Coulomb interaction. In this case there exist several eigen-modes of the collective motion of the ions, each with a different energy. The "common mode", for which the ions all move in the same direction, is lowest in energy, while the "stretch mode", for which the ion motions are opposed relative to the centre of mass of the ion chain, has

[2] i.e. the energy difference between neighbouring quantum harmonic oscillator states of the ion motion.

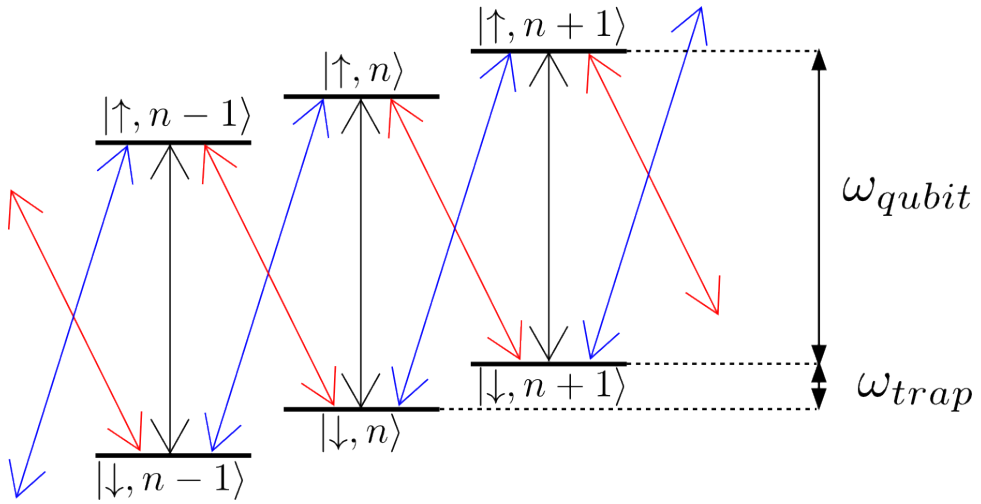

Figure 3: Red and blue sideband transitions. The energy levels shown are the combined internal $(\downarrow, \uparrow)$ + motional (quantum harmonic oscillator Fock states) states.

a higher energy. The motional couplings between multiple ions are a key element in trapped ion quantum information as they are used to perform multi-qubit gates. Thus additional technical requirements must be met in order to provide the motional control necessary for quantum information experiments; the potential must be as close as possible to being harmonic since anharmonicities can lead to coupling between eigen-modes, the frequency difference of the eigen-modes must be sufficiently great relative to the linewidth of the lasers used for the atomic transition to minimize coupling between different modes, and the system must be well isolated from the environment to achieve a low heating rate of motional states. The motional states can be cooled using Doppler cooling, EIT (Electromagnetically Induced Transparency) cooling, and also sideband cooling, where the red sideband transition is used to induce the $|\downarrow, n\rangle$ to $|\uparrow, n-1\rangle$ transition, while the 854 nm laser is also used to rapidly recycle the excited state back to the qubit ground state and progressively lower the motional energy to the motional ground state [5, 6].

1.3 Ion Trap and Control

Figure 4a shows the layout of one layer of electrodes in the ion trap used by the TIQI group at ETH Zürich [3]. Together, two such layers form the Paul trap with RF and DC electrodes diagonally opposed in an analogous way to the basic trap geometry shown in figure 2, and as seen in figure 4b. This gives a total of two RF electrodes and 30 DC electrodes [4]. Also labelled in figure 4a are two loading zones, as well as three experimental zones. The control system for the trap allows for a fast update rate of the DC electrode voltages in order to have precise control over the axial trap potential. This enables ions to be transported quickly through the trap with the appropriate temporal sequence of electrode voltages, henceforth referred to as waveforms.

To apply desired static voltages or dynamic waveforms to the DC electrodes of the

[3] This trap was designed and tested by Daniel Kienzler [7, p. 30–46].

[4] The physical implementation of the trap also contains an additional 28 shim electrodes to compensate for stray fields [7].

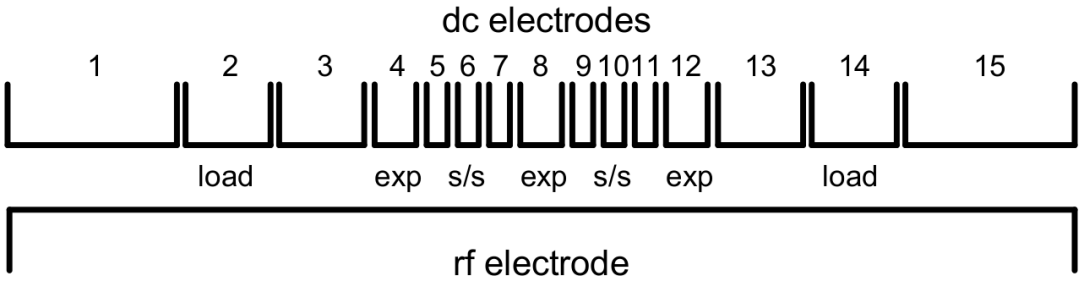

(a) Top view of one layer of trap electrodes showing the different trap zones [7].

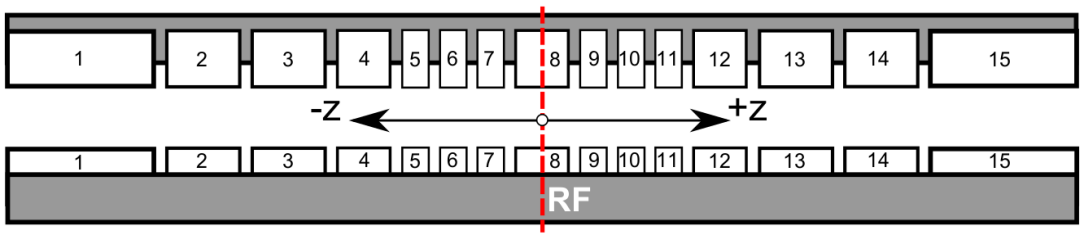

(b) Side view [8]. Note that the electrode numbering was in pairs before the addition of the new DEATHs.

Figure 4: Electrode layout of the segmented ion trap.

trap, a custom PCB solution designed by Ludwig de Clercq is used. The desired values for the voltages are sent via Ethernet to the unit, aptly named Direct Ethernet Adjustable Transport Hardware (DEATH), consisting of the Microzed 7010 FPGA and the custom PCB. The Microzed FPGA processes the received Ethernet data, and sends the results in the appropriate digital format to the custom PCB. This PCB, uses the MAX5898 digital to analogue converter (DAC) from Maxim to convert the digital signals received from the FPGA to analogue outputs between +10 and −10 volts for each of the DC trap electrodes [5]. Each DEATH contains two DACs, and each DAC can output two independent voltages, giving four independent channels per DEATH.

1.4 Original Scope of the Project

The original goal of the project was essentially twofold. Firstly, two observations had previously been made with respect to the linear Paul trap experiment. In particular, the heating rate of the stretch mode of the ion motion was significantly greater than that of the common mode, and it was also observed that the positions of the ions in the trap were rotated relative to their direction of motion along the trap [6]. These observations were unexpected; in principle the heating rate should be lower for a higher-energy mode, and the ions should lie on the trap axis along which they move when transported. In the experimental set-up used at the time, opposing electrodes in the trap were shorted

[5] For a more detailed description of the DEATH PCB and the original design considerations see Ludwig de Clercq's Ph.D. thesis [8, p. 53–77].

[6] i.e. if you were to join two ions by an imaginary straight line, then this line would be rotated relative to the direction in which the ions travel in the trap when they are transported.

together since ideally they should be at the same potential. In fact there are typically some stray fields which need to be compensated for, but this compensation was taken care of by shim electrodes designated to each of the trap zones [7]. However, being positioned further away from the ions than the DC electrodes, the shim electrodes have a limited correction capability. On the other hand, independently controlling the DC electrode voltages would provide significant added capability for correcting any anomalies in the field. Furthermore, individual electrode control enables operations such as swapping of ion positions which cannot be performed if opposing electrodes are shorted. Since the trap has 15 pairs of electrodes, the configuration with opposing electrodes shorted only required four DEATH boards (four channels per DEATH for a total of 16 channels, minus one spare channel which was used for monitoring). To obtain independent control of the 30 trap electrodes hence required four additional DEATHs.

The second goal of the project was to make some improvements to the design of the DEATH PCB. Since four new boards would be manufactured to enable individual electrode control, it was an appropriate time to make some modifications to their design. In particular, the original designer of the DEATHs, Ludwig de Clercq, had suggested several improvements that should be made to the DEATH design [8].

Firstly, the Serial Peripheral Interface (SPI) voltage lines used to configure the DEATH DACs were mismatched in voltage. That is, the DACs were meant to accept digital signals at 3.3 V (high) while the Microzed FPGAs used to configure them only output signals at 2.5 V (high). It had also been observed that, most likely due to the voltage mismatch, configuring the DACs typically had to be repeated several times before completing successfully [8]. It was thus desirable to add some chips to the DEATH PCB design to shift these voltages to the correct level.

Secondly, fluctuations of the DEATH output voltages with temperature were observed. This led to undesirable changes in the motional frequency of the ions. For this reason, potential sources of noise were investigated, leading to some additional modifications of the design.

Lastly, a header was added to the DEATH PCB to allow for fast branching in experiments. Typical trapped ion quantum information experiments involve the application of sequences of waveforms to the ions which are decided upon and programmed before the experiment takes place. Fast branching refers to a more flexible scenario, in which the next waveform or sequence of waveforms to be applied is decided during the experiment, conditioned on some information (e.g. the outcome of a measurement). The addition of a header to the DEATH PCB allows for direct communication between the Zedboard, which is the FPGA used for overall control of experiments [9], and the DEATHs. This direct and fast communication is an essential feature for fast branching experiments.

The rest of the report discusses the changes made to the DEATHs and the experimental set-up in more detail, and presents an experiment that was performed at the end of the project using the new capabilities of the ion trap control system.

Chapter 2 contains the details of the changes made to the DEATHs, as well as some other changes to the experimental set-up. An investigation on the temperature sensitivity of the DEATH PCB output voltages is also included.

[7] [7, p. 38, 39]
[8] [8, p. 77]
[9] [8, p. 28–30]

Chapter 3 explains experimental techniques used to investigate the ion trap system with the new capabilities explained in Chapter 2. Suggestions for future investigations are also mentioned.

Chapter 4 presents the results of an experiment to determine the input-output characteristics of the individual DC trap electrodes.

Appendix A contains a summary of challenges encountered in the manufacturing of the DEATH PCBs and provides recommendations for future manufacturing of the DEATHs or similar PCBs.

2 Changes to the Ion Trap Control System

2.1 Design Changes to the DEATH PCB

2.1.1 Level Shifting the SPI lines

The SPI lines are used to configure the DEATH DACs. Figure 5 shows a top view of the chip with numbered pins and the available areas on the PCB for placement of SPI level translators.

The SPI bus consists of five channels: DOUT, DIN, SCLK, $\overline{CS}$, and $\overline{RESET}$. All these signals are inputs to the DAC chip except for DOUT, which is the readout of the DAC registers sent back to the Microzed for debugging purposes. The choice of the level translators was mostly straightforward since Maxim has a family of chips designed specifically for shifting SPI line voltages. In particular, the MAX3390E is appropriate for shifting DOUT, DIN, SCLK, and $\overline{CS}$. This chip can translate signals at a maximum data rate of 16 Mbps for the voltage levels used in this case [10], which is well over the data rate of 10 Mbps (SCLK frequency of 10 MHz) at which the Microzed sends data to the DACs [9]. An additional level translator was needed for the $\overline{RESET}$ line. For this purpose the TXS0101 level translator from Texas Instruments was chosen [10] for a unique combination of low price and the option of a small package size (Package-Type: SOT, Orderable Device: TXS0101DRLR [11]) which was helpful given the limited space available on the DEATH design.

It is worth noting that the small and unusual package size chosen for the TXS0101 led to some later difficulties in the manufacturing by Gold Phoenix PCB Co.,Ltd. In particular, the pins of this level shifter and some power regulators were not soldered properly, leading to many hours of troubleshooting to find the problem and to resolder everything correctly. This should be avoided if more PCBs are ordered in the future. Some details on how the DEATHs were fixed and advice on how to avoid future problems are included in Appendix A.

2.1.2 Efforts to Reduce Temperature Fluctuations

Before the DEATHs were used to control the DC electrodes of the ion trap, the slower commercially available AD5371 DAC from Analog Devices was used instead. Although switching to the custom DEATH PCB hardware solution presented many advantages over the AD5371 [8], it was observed that the electrode voltages fluctuated more with temperature. For this reason efforts were made to find the source(s) of these fluctuations and to reduce them as much as possible.

One potential cause of fluctuations was the variation of the resistor values with temperature in the output stage op-amp circuit of the DACs, leading to changes of the gain with temperature. This was something that could be easily improved by replacing the resistors with new ones having lower resistance variation with temperature. Tables 1 and 2 list the old and new resistors of the DAC output stage, respectively. A few of the resistance values were changed slightly, leading to a different value for the expected gain

[10] Although any of the following chips could also have been good options: ADG3301 or ADG3231 from Analog Devices, PI4ULS5V201 from Pericom, NTS0101 from NXP Semiconductors, MAX13046E from Maxim Integrated, and NLSX4401 from ON Semiconductor.

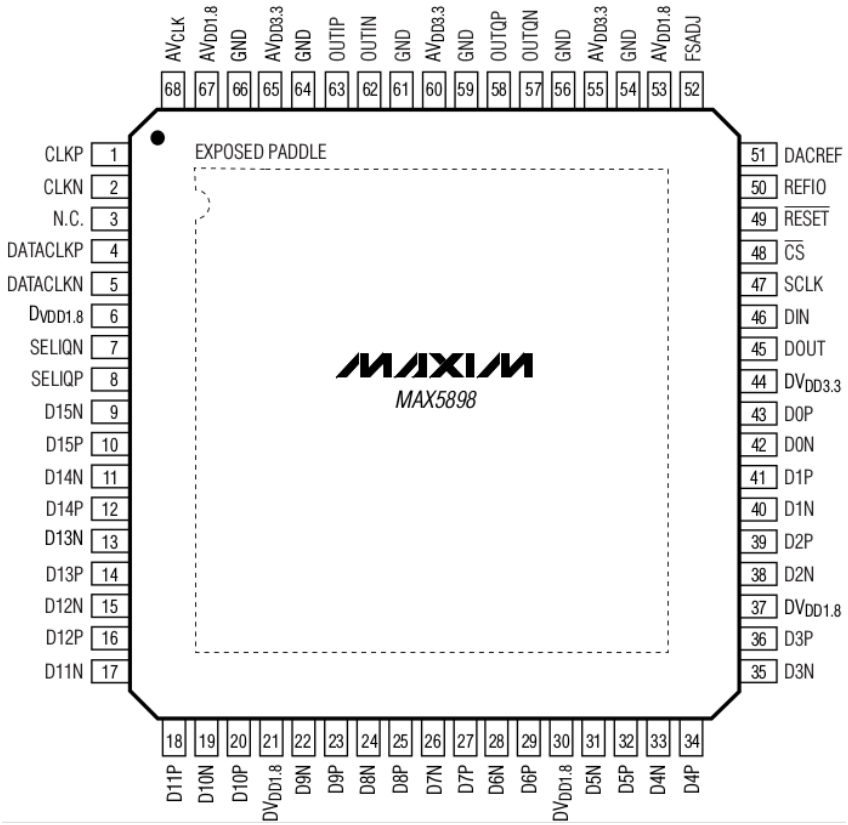

(a) Top view of MAX5898 DAC chip used in the DEATH PCB [9]. SPI lines are pins 45 to 49.

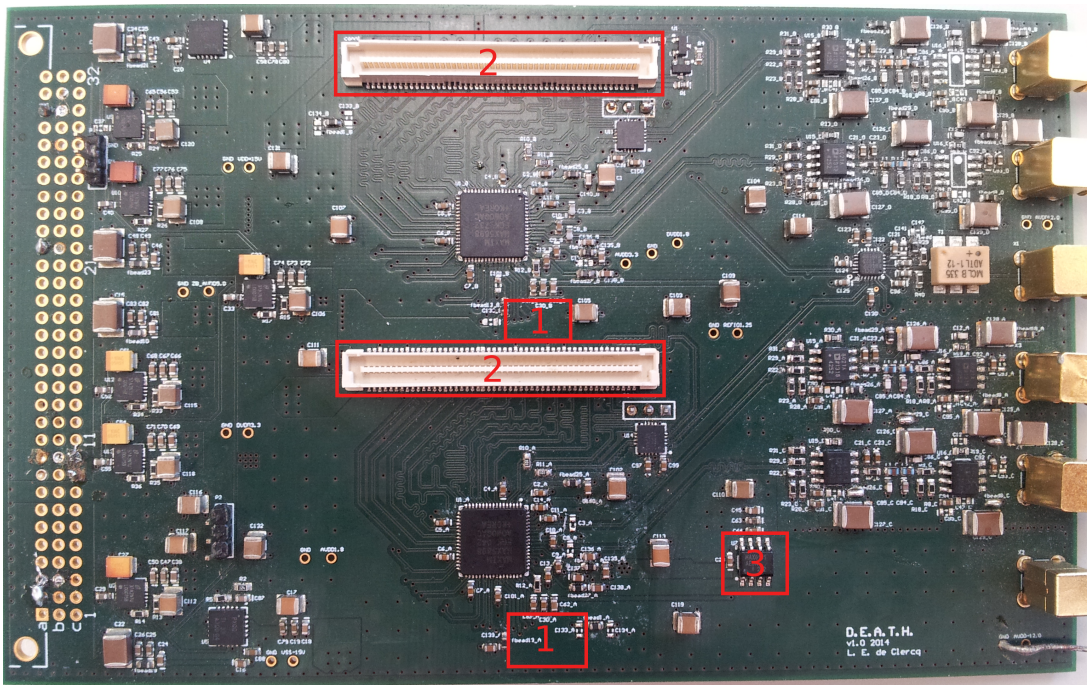

(b) Top view of the previous DEATH design. 1. Available spaces for level shifters. Note that the DACs are rotated 90° clockwise relative to the diagram of the DAC in (a). 2. Microzed headers. 3. MAX6161B voltage reference chip.

Figure 5: Level Shifter Placement

Schematic Designator	Resistance	Tolerance	Temperature Coefficient
R7	24.9Ω	±1%	±100ppm/°C
R8	49.9Ω	±1%	±100ppm/°C
R9	162Ω	±1%	±100ppm/°C
R18	71.5Ω	±1%	±100ppm/°C
R19	1.5kΩ	±0.1%	±25ppm/°C
R22	53.6 Ω	±1%	±100ppm/°C
R23	27Ω	±1%	±100ppm/°C
R28	287 Ω	±1%	±100ppm/°C
R29	340 Ω	±1%	±100ppm/°C
R30	1kΩ	±1%	±100ppm/°C
R31	28.7 Ω	±1%	±100ppm/°C

Table 1: Old DEATH DAC output stage resistors

of the new DEATH output stage compared to the old one; for a maximum output current of 20mA from the DACs [9], the old DEATH output stage leads to an expected maximum voltage of 9.70V, whereas the new output stage yields a value of 9.96V. The new values of resistance were chosen to impedance and gain match the DAC current outputs in order to reduce parasitic bandlimiting and increase linearity of the outputs [12].

In addition to noise resulting from the amplifying process discussed above, unwanted errors can be introduced further upstream by the DACs themselves. An initial suspect was the internal bandgap voltage reference of the DACs which has a typical temperature coefficient of ±50ppm/°C [9]. However these initial suspicions proved groundless since the DEATH design in fact already contained a MAX6161B external voltage reference for the DACs with a maximum temperature coefficient of ±15ppm/°C (typical: ±6ppm/°C) [13]. But the use of an external reference also requires a 2kΩ resistor placed across pins 51 and 52 of the DAC (see figure 5a). The expected full scale output current of the DACs can be calculated as

$$I_{outFS} = 32 \times (V_{refIO}/R_{set}),$$

where V_{refIO} is the external reference voltage, and R_{set} is the value of the resistance across pins 51 and 52 [9]. Therefore, we see that even if an extremely stable reference voltage is used, intrinsic fluctuations in resistance with temperature could lead to large changes in the DAC output. In particular, the resistance used on the old DEATHs had a temperature coefficient of ±100ppm/°C leading to fluctuations with temperature on the order of 0.1% in the DEATH's output voltage for a temperature change of 10°C. These resistors were replaced with new ones having a temperature coefficient of ±5ppm/°C.

Once the boards were manufactured, some measurements were performed to better characterize the temperature fluctuations. The general idea was to heat specific parts of the board with a heat gun [11] while measuring the response of the output voltages. The following procedure was used to try observing the effects as systematically as possible given the simple nature of the tests:

1. Turn on DEATH and run Microzed from desktop via USB (JTAG header) to set

[11] Hot Air Tool; Type PRO; 1800 W; LEISTER Process Technologies, CH-6060 Sarnen, Switzerland.

Schematic Designator	Resistance	Tolerance	Temperature Coefficient
R7	24.9Ω	±0.1%	±25ppm/°C
R8	49.9Ω	±0.1%	±25ppm/°C
R9	162Ω	±0.1%	±25ppm/°C
R18	71.5Ω	±0.1%	±25ppm/°C
R19	1.5kΩ	±0.05%	±10ppm/°C
R22	53.6 Ω	±0.1%	±25ppm/°C
R23	26.7Ω	±0.1%	±25ppm/°C
R28	316 Ω	±0.1%	±25ppm/°C
R29	332 Ω	±0.1%	±25ppm/°C
R30	1kΩ	±0.05%	±10ppm/°C
R31	28.7 Ω	±0.01%	±25ppm/°C

Table 2: New DEATH DAC output stage resistors

the target output voltage (t = 0).

2. Record voltage at the moment when the run is complete.

3. Apply heat for 10 seconds to a specific area of the DEATH board using the heat gun on heat setting 1, air flow setting at I, and distance of about 3 cm from the surface of the board (start this at t $\sim$ 1 min).

4. Record voltage right at the end of the 10 seconds of heating in step 3.

5. Turn off to let cool for $\sim$ 1 min before doing more measurements.

The first set of measurements aimed to investigate three effects: those of heating the DACs in comparison to heating the output stage, those of replacing the resistor R_{set} discussed above, and also whether the observed effects were consistent across different channels. The results of the first few measurements are plotted in figure 6.

From these first results the main observations were that the voltage fluctuations caused by heating the DACs were consistently larger than those caused by heating the output stage. Also, these were of opposite sign. Since the results were consistent across different channels, and reproducible upon repeated measurements, only a single channel was used for later measurements. The second significant result was that the observed fluctuations were not affected much by the new temperature stable resistors.

Following this first set of measurements, the fluctuations of the REFIO voltage, which is the external reference voltage for the DACs, were investigated. Upon measuring the REFIO voltage itself (nominal 1.25V) it was found to be incredibly stable and never fluctuated by more than 0.1 mV even after prolonged heating of either the DAC, the output stage, or the MAX6161B reference chip for the REFIO voltage itself. This was a strong indication that the reference voltage was not the root of the problem; so far the data had suggested that the DACs were in fact the main source of fluctuations. But since these did not seem to stem from the reference voltage or the R_{set} resistor, they must be caused by internal fluctuations in the DAC chips. The only remaining possibility was the gain error drift of the DACs of ±110ppm/°C [9]. To verify this, a second set of

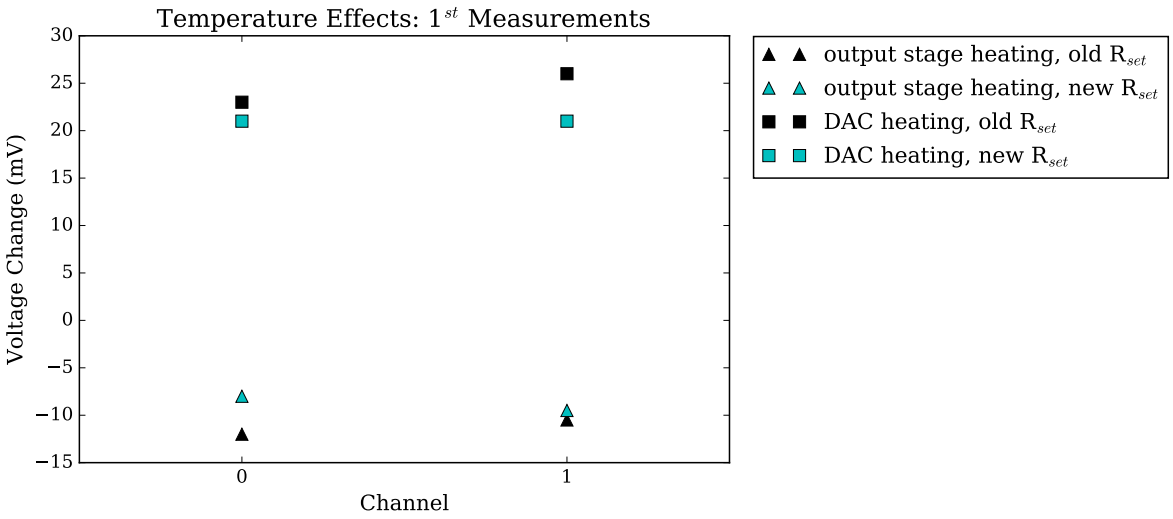

Figure 6: First set of temperature tests performed on the DEATHs. Three effects are investigated: the effect of heating the output stage versus heating the DAC chip, the effect of the resistor R_{set}, and possible differences between channels.

measurements was performed to investigate the effects of changing the output voltage. The results are plotted in figure 7. We see that upon changing the sign and magnitude of the output, the sign and magnitude of the fluctuations also change. This indicates that the inherent fluctuations of the MAX5898 gain are indeed the main cause. Some data for the effects of heating the MAX6161B are also presented in figure 7, indicating again that this is not a significant source of fluctuation. And in fact, the voltage change seen for the MAX6161B heating could even be due to residual heating of the DAC by the heat gun since the MAX6161B is reasonably close to the DAC (about 3cm).

A third set of measurements were made to explicitly isolate the effects of the DAC from the output stage by measuring the differential voltage directly output from the DACs: $SIG_P - SIG_N$. Since the input impedance of the op-amp is very high, almost all the fluctuations observed should be due to upstream effects (i.e. the DACs). These results are plotted in figure 8. The effects of heating the output stage show mainly two things. One, that the fluctuations due to heating of the output stage are greatly reduced, and two, that they are of the same sign as those due to heating the DAC, which is in contrast to previous measurements of the final outputs. One possible explanation is that the observed fluctuations upon heating the output stage are due to residual heating of the DAC or other relevant upstream components (e.g. the MAX6161B is positioned between DAC A and the output stages for channels 0 and 1 on the physical PCB, 3. in figure 5b). The results of heating the DACs agree quite well with previous results for the output voltage. For instance the fluctuations in the last measurement are of 998 μV which corresponds to roughly 20 mV fluctuations of the output – i.e. most of the previously observed fluctuation range (e.g. 21 mV). The fact that only a 480 μV change was observed when the offset was set to 50500 may only be due to a lower amount of heat applied compared to other runs.

Thus, from this investigation we see that the main source of temperature fluctuations is the inherent fluctuation of the DAC gain. The second most important contributor is

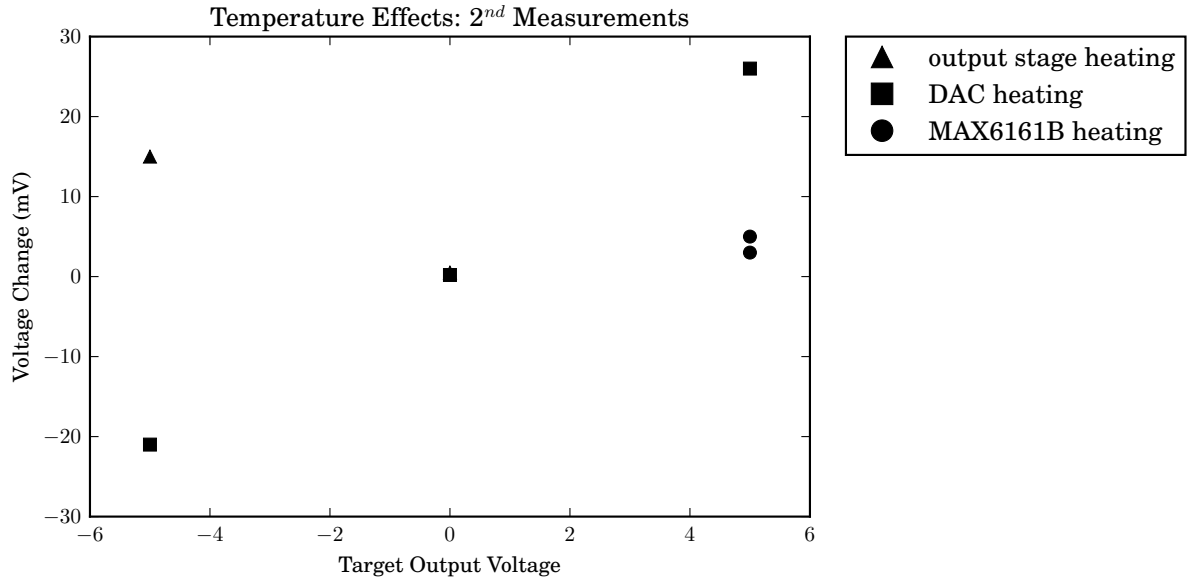

Figure 7: Second set of temperature tests performed on the DEATHs. Here the effects of target output voltage are investigated. Note that for an output voltage of 0V there are in fact two data points overlapping – one for output stage heating, and one for DAC heating. The output stage heating was not measured for +5 V.

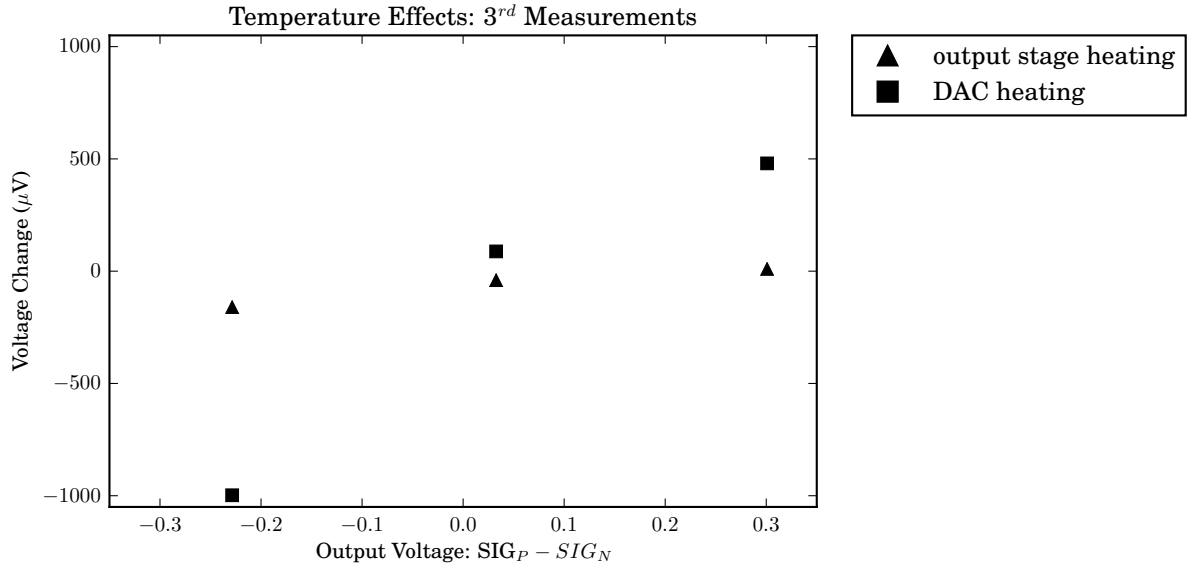

Figure 8: Third set of temperature tests performed on the DEATHs. In this case, the DAC outputs were measured directly to isolate the DAC fluctuations. In order of increasing voltage, the three voltages plotted correspond to digital offsets of the DAC channel of 15500, 32768, and 50500.

the output stage which also exhibits gain fluctuations. It may also be that to some extent the fluctuations of these two components cancel each other as they were found to be of opposite sign. If significant improvements are needed in the future it may be necessary to change the DACs, something the original designer had also anticipated for different reasons: *"In hindsight we would have chosen a DAC with lower update rate and higher output voltage if possible to avoid amplification on the final stage minimizing amplification of drifts and pre-amplifier noise"* [8]. However, some improvement could also be achieved by placing thick metal plates in contact with the DEATHs via thermal pads, ideally on both sides, in order to stabilize operating temperature.

2.1.3 Adding a Header for Fast Branching

Fast branching of experiments is achieved by preparing several options or branches from which to choose for the next stage of the experiment. Since the qubit states can decohere on time-scales of milliseconds or less, the time between the moment when the next branch is chosen and the moment when the DEATH begins outputting the next waveform should be minimized. With the old design of the DEATHs, such branching is possible by having the Zedboard communicate the next branch choice to the control computer, which in turn sends information to the Microzed via Ethernet, thereby initiating the next waveform output from the DEATHs. This process takes on the order of 20 μs at best. It is possible to achieve much faster communication by allowing the Zedboard to communicate directly with the Microzed. To this end, a new header was placed on the DEATH PCB and routed to unused pins of the Microzed header (labelled 2 in figure 5b), potentially allowing for communication from the Zedboard to the Microzed in under 200 ns. The choice of location for the new header was considered carefully to maximize the number of previously unused pins that could be routed to the header for fast branching. It was possible to make use of eight free pins, giving up to one byte of branching options.

The new DEATH PCB design is shown in figure 9.

2.2 Adding the New DEATHs to the Experimental Set-up

2.2.1 Hardware

In addition to making changes to the DEATH design, it was also necessary to add several other components to the experimental set-up to accompany the DEATHs. And since twice as many channels were to be used in order to control each trap electrode independently, a few changes to the previous set-up were also needed.

To accompany the new DEATHs, the same components were purchased and assembled as for the original DEATHs, including a box with power supplies and fans to house the DEATHs, Ethernet cables, and XLR audio cables [12]. The changes necessary to accommodate double the number of DEATH channels were straightforward thanks to convenient design of the existing system. These included increasing the size of the connector box, and removing the shorts between electrodes in the backplane PCB of the filter box.

[12]The XLR cables used here were soldered manually, which took a significant amount of time. If more similar cables are needed in the future, it might be worth directly purchasing audio cables with XLR connectors already mounted.

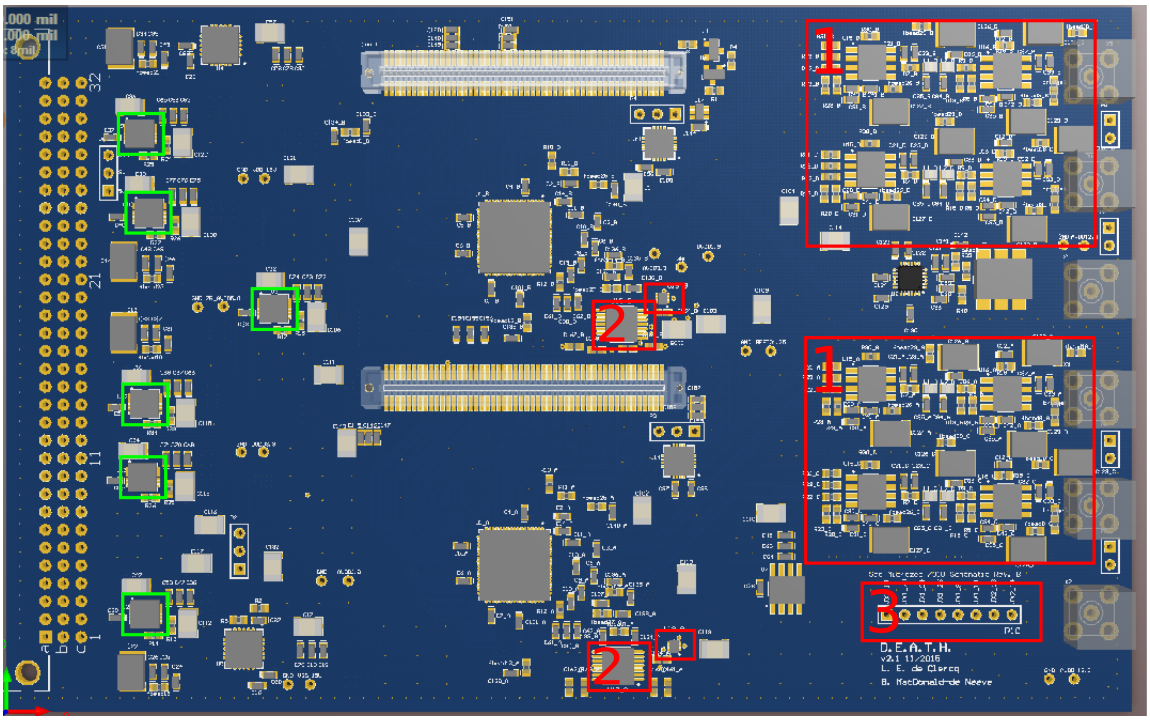

Figure 9: The new DEATH PCB. Components that were added or changed are indicated in red and numbered. Region 1 indicated the output stage where the resistors were replaced. 2 labels the added level shifters – the larger one is the MAX3390E, the smaller one the TXS0101. 3 labels the header added for fast branching. Also indicated in green are the LP38798 power regulators which led to inadequate soldering (see Appendix A).

2.2.2 Software

Before the new DEATHs were ready to be added to the experiment, the Microzeds had to be programmed with the appropriate firmware. For the most part this was essentially the same as for the old DEATHs, but required some debugging along the way before everything was working correctly. The debugging process can be summarized as:

1. Programming the Microzeds: generate bitstream, and program flash
2. Configure DEATH DACs from the Microzed
3. Test DEATH outputs using serial (USB) communication
4. Test DEATH outputs using Ethernet communication

The bitstream and firmware in step 1 was already developed for the old DEATHs and the only new feature that had to be added was four new configurations for the PMOD jumper in order to assign unique Ethernet MAC addresses to the new DEATHs. Step 2 involved programming the DAC registers via the SPI bus which should have been straightforward and more effective than for the old DEATHs due to the level translation of the SPI lines. In step 3, individual DEATHs were tested by programming simple waveforms from within the SDK (Software Development Kit). The last step was somewhat more involved since the Zedboard must be used to control the DEATHs via Ethernet; however this is how they are operated in real experiments. The success of this final step ensured that the likelihood of encountering any problems once the DEATHs were added to the experimental set-up was quite low.

A significant amount of debugging was required before the last step was successful. Initially, step 2 failed due to the manufacturing problems with the new DEATHs mentioned in section 2.1.1. Once these were fixed, step 3 was successful for all four of the new boards. In particular, a simple ramp waveform was programmed and all of the DEATH outputs were tested [13]. In the last step, one other problem was found with two channels on one of the DEATHs [14], namely, the lower DAC was outputting large sinusoidal oscillations with a peak to peak amplitude of several volts.

The fact that this did not occur while testing the simple ramps via serial gave the clue needed to fix the problem. One difference between running the DEATHs via serial versus running them via Ethernet is that, for the former, the DEATHs (including Microzeds) are powered on for several seconds or minutes before the user programs the PS (Programming System) of the Microzed. Only then does the Microzed configure the DACs via SPI. With Ethernet, the Microzed is run from its previously programmed flash memory which means that the DACs are configured almost immediately when they are powered on. The problem was that the lower DAC on board 2 was somewhat slower to power on, so that the Microzed would try to configure it before it was ready.

To solve this problem, the Microzed flash was reprogrammed so that it would reconfigure the DACs several times with a wait time in between trials until the DAC responded

[13] The data was saved to the J Drive `/Temp/Brennan/deaths2_ramp_data/` for future reference when testing new DEATHs

[14] This was DEATH 2. They are numbered on the back of the main header that provides the DEATH power.

correctly. To determine whether this fix worked initially, it was useful to reproduce the effect using serial communication and Teraterm (in which case the Microzed still runs from flash and programs the DACs when powered on). This way the number of iterations required, and the DAC SPI outputs, could be obtained. In principle this could also be done with Ethernet instead of serial, but was not possible with the existing SDK code.

With this last fix, all the new DEATHs were seen to operate correctly in step 4. If ever these channels are found to fail, re-powering the system should typically fix the problem; once they are working they should not stop working unless the system is powered off and then back on again. It may be that over time the DAC takes more time to power on, in which case reprogramming the flash memory to iterate longer should fix the problem.

3 Characterizing the Ion Trap

Once the new DEATHs were installed and working, the added capability could be used for several new experiments. As mentioned above, an unexpectedly high heating rate was observed for the stretch mode of the ions in the trap, and the ion positions appeared to be misaligned relative to the axial direction of the trap.

Some experiments were devised in order to better characterize the trap performance, and perhaps also reveal the causes of these unexpected observations in the process. The first of these experiments verified the operation of each of the electrodes individually, and was implemented at the end of this project. The idea was to trap a single calcium ion and initially apply inputs to the electrode channels in order to generate an axial harmonic well in the centre of the trap with a desired frequency. Then offset voltages were applied to each of the electrodes individually leading to a shift in the frequency of the potential well. This frequency shift was predicted from simulation, and also measured directly using the ion. To measure the frequency of the trap with the ion, an RF voltage oscillation, typically referred to as a "tickle", was applied to the trap using one of the DC electrodes. We expected a maximal increase in the energy (i.e. number of quanta of motion) of the the ion's motional state, when the "tickle" was resonant with the trap frequency. The ion was first cooled close to the motional ground state, then a "tickle" was applied, and the motional state measured again. By repeating for different "tickle" frequencies, a peak in motional excitation was observed on resonance with the trap frequency.

To make the measurement of the trap frequency very sensitive, we cooled very close to the ground state using first Doppler cooling, followed by EIT cooling, followed finally by sideband cooling, each allowing to cool to a lower temperature of the ion. Then the "tickle" was applied, followed by a red sideband pulse to measure motional excitation. If the ion was in the motional ground state, the red sideband process was inhibited, since the ion cannot lose any motional quanta. Otherwise, the internal state of the ion was excited. This provided a sensitive measure of motional excitation since only small deviations from the motional ground state led to population in the excited state of the ion qubit.

Thus, by applying voltage shifts to each of the electrodes individually, and by comparing the measured axial frequency shifts to those predicted from simulation, we were able to spot any anomalies in DC electrode performance. The results are presented in the next chapter. In order to predict the trap frequencies, a short python script tailored to this experiment was written using the `pytrans` python module written by Vlad Negnevitsky.

Since the unexpected properties of the trap were measured only for the central experimental zone of the trap, it would also be useful to verify what the alignment of the ions is along the entire trap. In addition, it would be useful to repeat the experiment explained above for the ion positioned in different trap regions. Since the potential well in the above experiment is located in the centre of the trap the frequency shifts caused by electrodes at the edges of the trap are very small and difficult to measure. For this reason, repeating a similar experiment for cases where the potential well is located closer to the edges of the trap would yield more accurate characterization of input-output responses for all trap electrodes.

4 Results

As discussed in the previous section an experiment was performed to better characterize the performance of the ion trap. A set of voltages was applied in order to generate a harmonic well in the centre of the trap at a desired frequency – in this case roughly 1.5 MHz. Offset voltages were then applied to each of the electrodes in turn and the resulting frequency shifts of the harmonic well were measured.

The measured shifts compared with predictions from simulation are presented in figure 10. Each value for the measured frequency was calculated by fitting a Lorentzian to the calcium "tickle" experiment results. An example plot of typical data is shown in figure 11. Each data point results from a "tickle" applied at the specified frequency [15]. The error bars in figure 10 were calculated from the uncertainty in the Lorentzian fit, but in general the error is relatively small and can hardly be seen in the figure.

The numbering of the electrodes used here is 1–15 for the top 15 electrodes of the trap, and 16–30 for the bottom electrodes. The mapping of the DEATH channels to the electrodes that was used at the time of the experiment is shown in table 3 [16].

Top Electrodes	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**
DEATH Channels	1	5	9	3	7	11	15	19	23	27	31	17	21	25	14
Bottom Electrodes	**16**	**17**	**18**	**19**	**20**	**21**	**22**	**23**	**24**	**25**	**26**	**27**	**28**	**29**	**30**
DEATH Channels	2	6	10	4	8	12	16	20	24	28	32	18	22	26	30

Table 3: Channels used to drive each DEATH electrode at the time of the experiment. Channels 1–16 correspond to the old DEATHs, channels 17–32 to the new DEATHs. Channels 13 and 29 were used for monitoring.

Both negative and positive shifts were applied and are shown in separate plots. In general the applied offsets led to a shift of the same sign, however for the centre electrodes this is not the case. This can be understood simply as a reduced (increased) confinement due to increasing (decreasing) the potential in the centre of the potential well. Unfortunately, a one of the centre electrodes of the trap, electrode 23, was not measured, as can be seen in the bottom two plots of figure 10. But this should not affect the results or conclusions of the experiment as will be discussed in more detail below.

The electrodes at the edge of the trap have less effect on the frequency shift since they are furthest from the potential well at the centre of the trap. For this reason, offsets were applied to three electrodes simultaneously for the edge electrodes. As discussed in the previous section, the edge electrodes could be measured more accurately using potential wells closer to the trap edges.

Electrode 11 was seen to not work at all since no frequency shifts were observed for either positive or negative voltage offsets.

Although the relative frequency shifts shown in figure 10 matched the predicted values quite well (except for electrode 11), the absolute frequency values measured were consistently higher by roughly 45 kHz as seen in figure 12a. By assuming a voltage offset

[15] See section 3 for an explanation of the experimental steps performed for each "tickle".

[16] This configuration was chosen to place lower DAC channels (of each DEATH) in the central region of the trap ,where greater precision is desired, since non-linearities were observed to be an order of magnitude larger for the upper DACs.

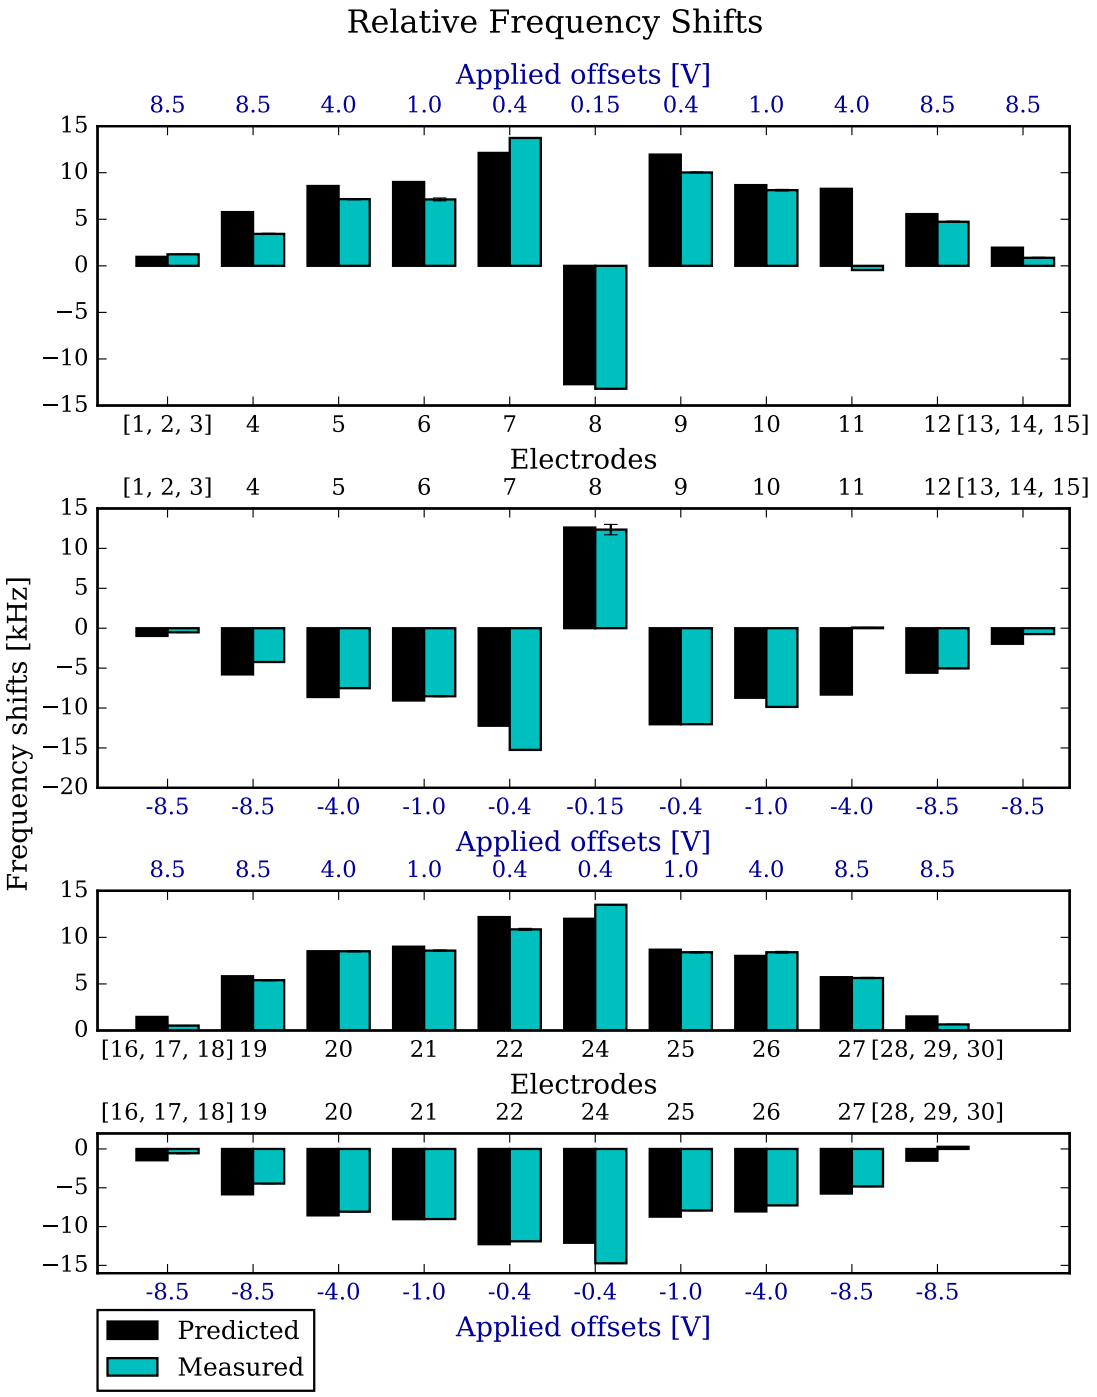

Figure 10: Results for the frequency shifts due to offset voltages applied to the trap electrodes one-at-a-time. Error bars from fitting the results are generally too small to be seen on the plot. The top two plots show the shifts which resulted from applying positive and negative voltage offsets respectively to electrodes 1 to 15. The bottom two plots show the results for electrodes 16 to 30 (note that electrode 23 was not measured). For the electrodes near the edge of the trap, shifts were applied for three electrodes simultaneously to try increasing the effect which was weaker due to the greater distance from the centre of the potential well. Electrode 11 did not cause any shifts.

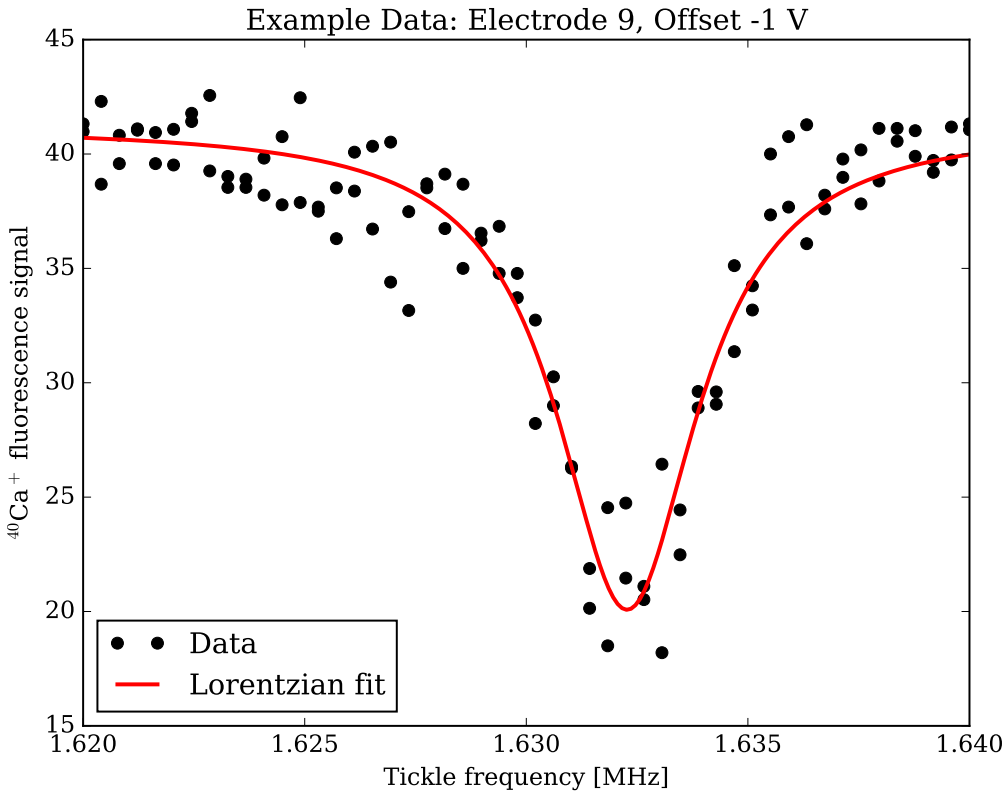

Figure 11: Data obtained for electrode 9 with an offset of -1V shows the typical results for a series of tickle experiments. We see an increase in excitation of the internal qubit state when the tickle is resonant with the motional frequency. Error bars in figure 10 are calculated from the fit uncertainty.

of 0.0 V for electrode 11, and an overall gain correction of +5.7% [17] for all electrodes, it was possible to obtain much better agreement between predicted and measured values, as shown in figure 12b. Whether this agreement reveals error physically due to gain error or some other form of errors is not clear. On the other hand the good agreement of the data with predictions at least shows that correcting for the gain is an effective way of increasing the accuracy of the simulations. In figure 13 the differential error (measured - predicted) is also shown.

Although electrode 23 was not measured in this experiment, it is unlikely that it is malfunctioning for two reasons. Firstly, being a centre electrode it has a large effect on many experiments since the centre experimental zone is often used. Therefore a malfunction of this electrode would likely have had a noticeable effect on these experiments, which was not the case. Secondly, since the potential in the experiment above is itself centred relative to the trap, malfunction of electrode 23 would have had a significant effect on the measured frequencies. In particular, assuming a voltage shift of 0.0 V for electrode 23 in the simulation leads to significant deviation of all frequency values from those measured.

[17] Using numerical least squares minimization a value of 5.6616% was obtained; however, since the applied offsets can bring some of the voltages close their maximum capability, increasing the gain can lead to non-physical voltage values in the simulation. Therefore, the calculation was also repeated excluding all values for electrode offsets of ±8.5V, and in this case, a value of 5.6983% was found.

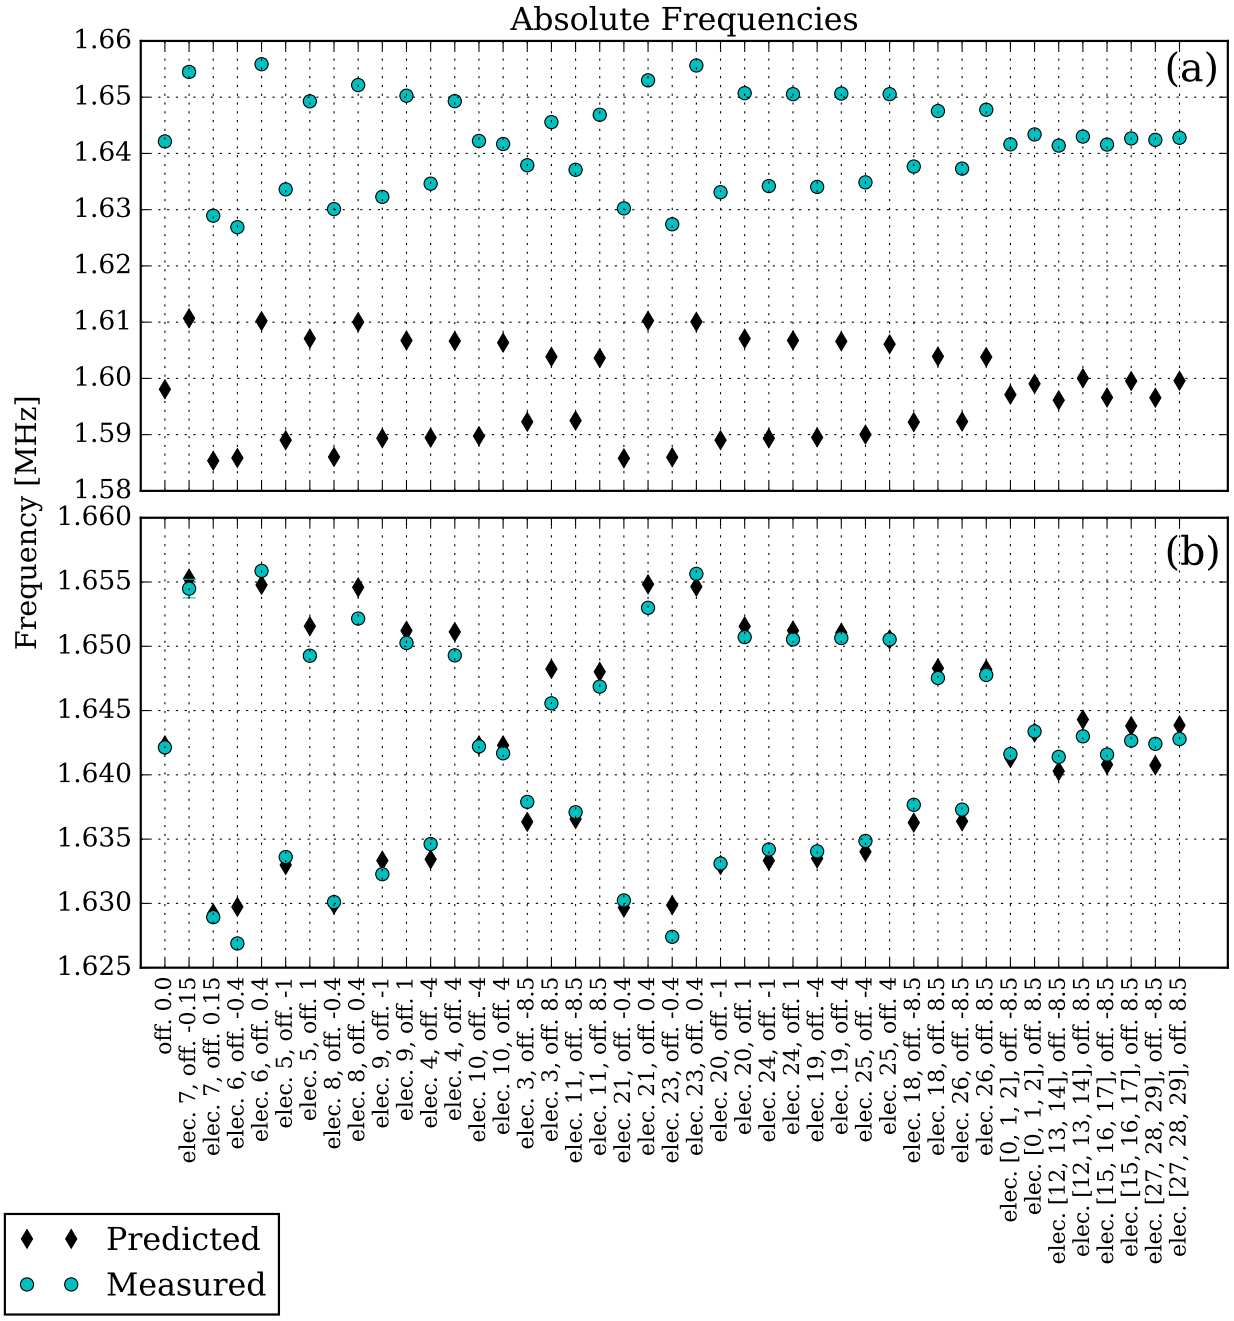

Figure 12: Absolute frequency shifts predicted and measured. (a) Initial predictions from the simulation plotted with measured data. (b) Adjusted predictions obtained by setting the voltage of electrode 11 to 0.0 V, and adding a gain correction of +5.7% to all electrodes.

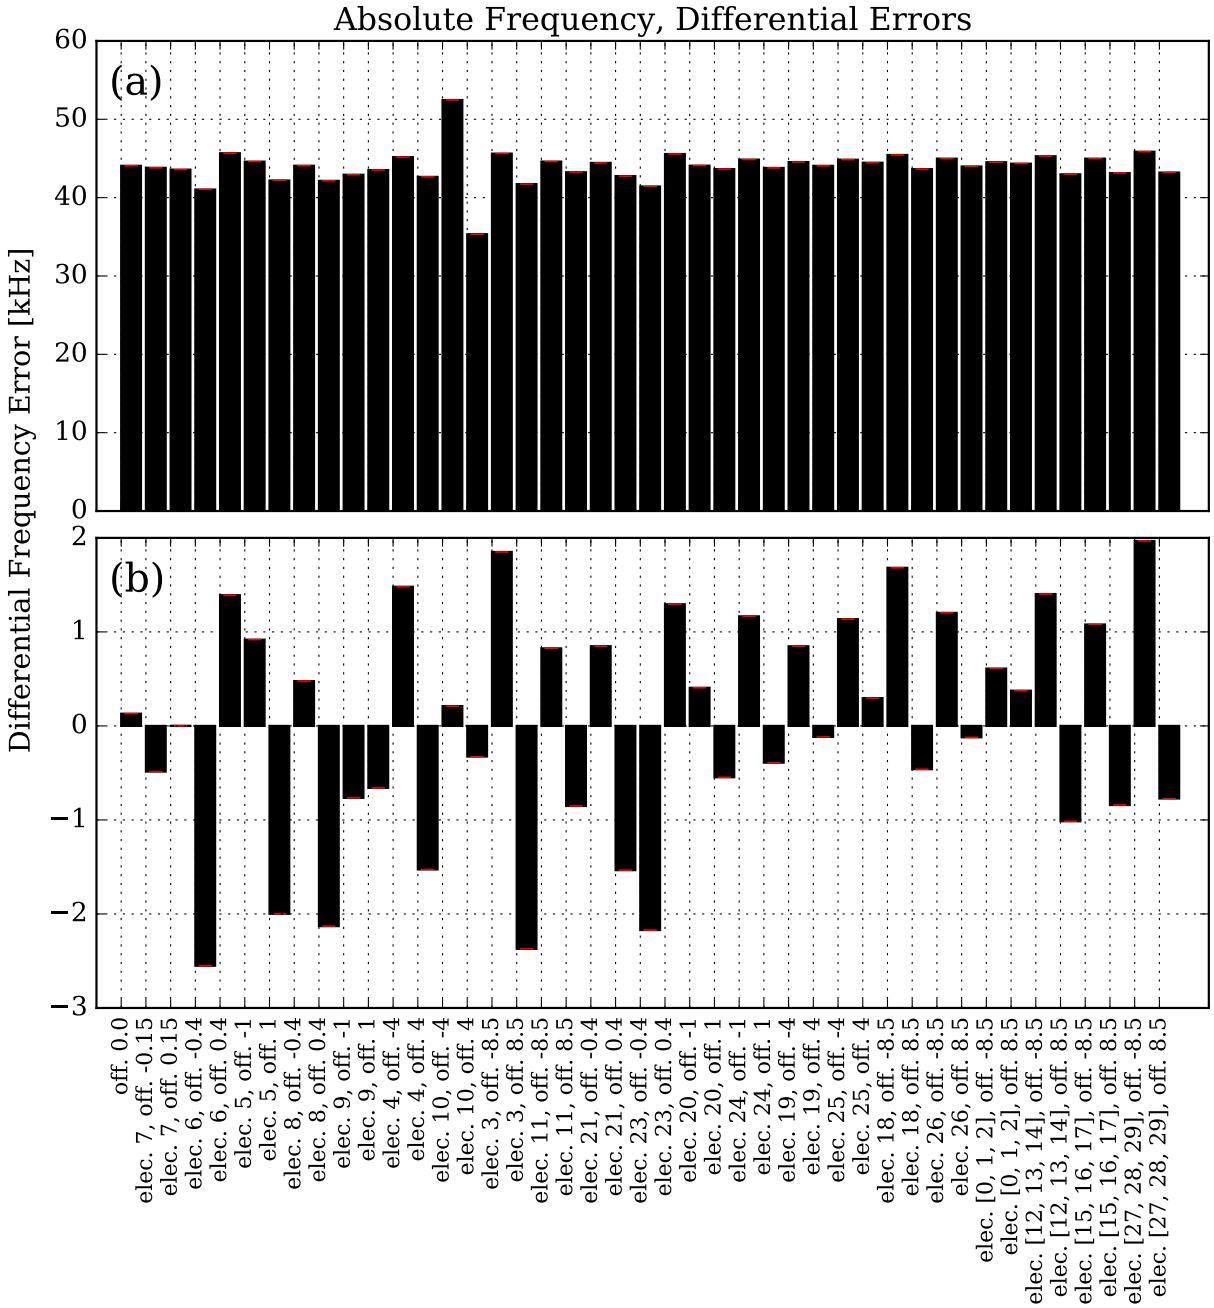

Figure 13: Differential error (measured - predicted) of the absolute frequency values. Error bars in red are mostly too small to be seen. (a) Initial predictions from the simulation. (b) Adjusted predictions obtained by setting the voltage of electrode 11 to 0.0 V, and adding a gain correction of +5.7% to all electrodes.

5 Conclusion

The initial scope of the project was to make some improvements to the control of the DC electrodes of the Paul trap, enabling control of the *individual* electrodes, and thereby giving the necessary means to better characterize the trap and correct for undesirable characteristics. To achieve individual control of the trap electrodes, it was necessary to manufacture four new DEATH PCBs. At the same time, the DEATH design had some known problems, namely, the SPI voltage lines did not operate at the correct voltage, some unknown problems, i.e. undesirable fluctuations of their output voltages had been observed during experiments, and the design could be improved with additional features, i.e. a header could be added for fast branching experiments.

Therefore, some components were added to correct the voltage of the SPI lines, a new header was added to the design for fast branching, some changes were made to try reducing output voltage fluctuations, and investigations were made to better understand the origins of these fluctuations. Other than a few difficulties encountered in the manufacturing of the new DEATHs, the design changes were successful. Investigations into the output voltage fluctuations revealed that the largest contributors were the DAC chips, and therefore it would not be possible to improve the stability of the output much further by design changes not involving replacement of the DACs. If significantly lower fluctuations are required in the future, it should however be possible to improve the stability substantially by adding thick metal plates placed in direct contact with thermal pads ideally on both surfaces of the DEATH PCBs to stabilize the temperature. Another option would be to redesign the DEATHs or a similar custom PCB to use different DAC chips with lower update rate and larger output voltage.

Some changes were made to the experimental system and some components were added to allow the new DEATHs to be used alongside the old ones to control the trap electrodes individually instead of in pairs. This added capability allows for better control and correction of experimental errors in order to achieve better conditions for quantum computing. It should also enable operations which were not possible with paired electrode control, such as controlled swapping of the positions of different ion species in the trap.

Using the new capabilities of the ion trap, an experiment was performed to characterize the performance of the individual DC trap electrodes. It was found that electrode 11 of the trap was not outputting the desired voltage – a problem which can now be investigated further and fixed if necessary. The results of this experiment also allow for better simulation of the trap since differences between simulated and measured results can be fed back into simulations to achieve more accurate control over the trap potential. In this experiment it was found that increasing the voltage gain in the simulation by 4.7% led to much better agreement with measured absolute frequency values.

Thus, improvements in the control of the trap and better characterization of the trap and its control systems were achieved, and also possible avenues for improved performance found. The unexpected observations of rotated ions and stretch mode heating have yet to be fully understood, but the new capabilities of the trap will hopefully help to reveal their cause and enable correction of the undesirable characteristics.

Appendix

A DEATH Manufacturing

As mentioned in section 2.1.1, some of the chips on the DEATHs were not soldered properly during their manufacturing by Gold Phoenix, meaning that the DEATHs did not work initially after they were received – they required re-soldering of several chips before they were able to operate correctly. This section presents a short summary of the specific problems that were found, how they were fixed, as well as suggestions on how to avoid them in the future, and is intended to serve as a reference for future manufacturing of the DEATHs, or other PCBs.

A.1 Where the Problems Occurred and Where to Expect Them in the Future

Most of the inadequately soldered chips were the TXS0101 level shifters of which only three were soldered correctly out of the eight chips contained on the four new DEATHs. In addition to this, many of the LP38798 power regulators from Texas Instruments were not soldered adequately – one of which completely lacked solder on the ground pin, leading to the digital 3.3V on that board being at 1.4V instead. None of the other chips had any problems. Both the TXS0101 (TXS0101DRLR, Package-Type SOT) and the LP38798 have the smallest spacing between their pads of any chips on the DEATH PCB, namely 0.2 mm spacing. In addition, both these chips feature packages with unconventional pins as shown in figure 14.

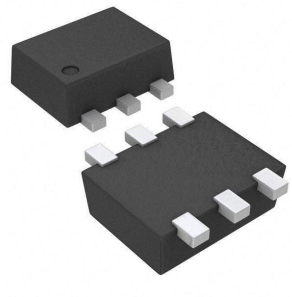

(a) Package of the TXS0101 level shifters used in the new DEATHs.

(b) Package of the LP38798 power regulators.

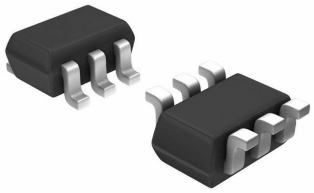

(c) Alternative package of the TXS0101 level translator with more conventional pins.

Figure 14: Different Pin Types

To prevent this problem in the future, it should be specified to Gold Phoenix that these two chips require special attention with respect to soldering. Another option would be to change the package of the TXS0101 to SC70 (see Addendum-Page 1 in the TXS0101 datasheet [11]), which would require updating the PCB design slightly. Note that if the package size is changed to SC70, the chip NTS0101 from NXP Semiconductors is also a good option.

A.2 Soldering Level Shifters

The TXS0101 level shifters were particularly difficult to resolder. For this reason it was very helpful to check the voltages on the pins of the level shifter to verify whether the soldering was optimal. Figure 15 shows the pins of the TXS0101 level shifter. To check whether the pins were soldered correctly a scope or multimeter was used to measure the voltages at each of these pins when the DEATH was powered on, but *without* the Microzed in place. A properly soldered and undamaged level shifter had the voltages listed in Table 4. This measurement was especially useful since sometimes I found that poorly-soldered level shifters could be enough to get a DEATH board operating properly, but upon repeated use would stop working at random intervals. I always found this undesirable behaviour to be correlated with a failure to achieve the voltages listed in Table 4.

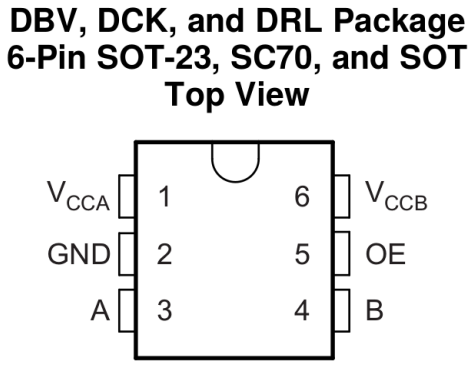

Figure 15: Pins of the TXS0101 level translator [11]

Pin	Voltage
V_{CCA}	2.5V
GND	0V
A	2.5V
V_{CCB}	3.3V
OE	2.5V
B	3.3V

Table 4: Expected pin voltages when the level shifter is correctly soldered. This can be measured when the DEATH is powered on, and *without* the Microzed connected.

Acknowledgements

Thanks to Jonathan for hosting the project and to everyone in the group that helped and encouraged me along the way; in particular Christa for helping me understand better the how the ion trap works and the problems that were observed, Jonathan for some nice discussions about ion trap physics, and Robin for helping me get through the difficulties encountered with the new DEATHs. And special thanks to my supervisor Vlad Negnevitsky for all his help, guidance, and encouragement throughout the project, for encouraging me in my transition to Linux, and for introducing me to Altium, Git, Emacs, and Python!

References

[1] C. Monroe and J. Kim. Scaling the ion trap quantum processor. *Science*, 339(6124):1164–1169, 2013.

[2] D. Kielpinski, C. Monroe, and D.J. Wineland. Architecture for a large-scale ion-trap quantum computer. *Nature*, 417:709–711, 2002.

[3] H. Häffner, C. F. Roos, and R. Blatt. Quantum computing with trapped ions. *Physics Reports*, 469(4):155–203, 2008.

[4] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland. Quantum dynamics of single trapped ions. *Rev. Mod. Phys.*, 75:281–324, Mar 2003.

[5] Michael Fleischhauer, Fachbereich Physik, Technische Universität Kaiserslautern, D-Kaiserslautern, Atac Imamoglu, and Jonathan P. Marangos. Electromagnetically induced transparency: Optics in coherent media. *Rev. Mod. Phys*, pages 633–673, 2005.

[6] Serge Haroche and Jean-Michel Raimond. *Exploring the Quantum*. Oxford University Press, 2006.

[7] Daniel Kienzler. *Quantum Harmonic Oscillator State Synthesis by Reservoir State Engineering*. PhD thesis, ETH Zürich, 2015.

[8] Ludwig Erasmus de Clercq. *Transport Quantum Logic Gates for Trapped Ions*. PhD thesis, ETH Zürich, 2015.

[9] Maxim Integrated Products, 120 San Gabriel Drive, Sunnyvale, CA 94086. *16-Bit, 500Msps Interpolating and Modulating Dual DAC with Interleaved LVDS Inputs*, 2010. MAX5898 datasheet, 19-3756; Rev 2; 8/10.

[10] Maxim Integrated Products, 160 Rio Robles, San Jose, CA 95134 USA. *MAX3372E-MAX3379E/MAX3390E-MAX3393E* $\pm 15kV$ *ESD-Protected,* $1\mu A$*, 16Mbps, Dual/Quad Low-Voltage Level Translators in UCSP*, 2013. MAX3372E-MAX3379E/MAX3390E-MAX3393E datasheet, 19-2328; Rev 3; 1/13.

[11] Texas Instruments, Post Office Box 655303, Dallas, Texas 75265. *TXS0101 1-Bit Bidirectional Level-Shifting, Voltage-Level Translator With Auto-Direction-Sensing for Open-Drain and Push-Pull Applications*, 2015. TXS0101 datasheet, SCES638C– October 2007–REVISED DECEMBER 2015.

[12] Michael Steffes. *Wideband Complementary Current Output DAC to Single-Ended Interface: Improved Matching for the Gain and Compliance Voltage Swing*. Texas Instruments, Application Report, Post Office Box 655303 Dallas, Texas 75265, 2005.

[13] Maxim Integrated Products, 120 San Gabriel Drive, Sunnyvale, CA 94086. *Precision, Micropower, Low-Dropout, High-Output-Current, SO-8 Voltage References*, 2005. MAX6161-MAX6168 datasheet, 19-1650; Rev 3; 8/05.