

Nuclear Magnetic Resonance Quantum Information Processing

Samuel James Bader
 MIT Department of Physics
 (Dated: April 13, 2013)

We implement several classic algorithms to characterize the possibility for quantum computation on MIT Junior Lab’s Nuclear Magnetic Resonance Spectrometer: the CNOT logic gate, the Deutsch-Jozsa algorithm, and the Grover search. Results demonstrate that basic computations with fidelities on the order of 70% are possible.

I. INTRODUCTION

The idea of quantum computing is to store and manipulate information contained in quantum mechanical systems, rather than macroscale objects. This allows for fundamentally new means of computation, employing algorithms which make use of intrinsically quantum mechanical effects such as superposition (which allows for massive parallelism) and entanglement (which expands the computational space exponentially).

Among the many implementations proposed (optical circuits, trapped ions, nitrogen-vacancy centers, superconducting circuits, etc), here we probe the plausibility of Nuclear Magnetic Resonance quantum computing, which has the advantages of long coherence times (on the order of seconds) and requires only an already-developed platform (commercial NMR systems have long been in use in chemistry and medicine).

II. APPARATUS

This experiment is performed on a Bruker Avance 200 NMR Spectrometer provided by the MIT Junior Lab. This standard NMR configuration is depicted in Figure 1. The sample is contained within a strong magnetic field (defining \hat{z}), wrapped by a solenoid which allows the computer to (1) apply RF pulses of orthogonal magnetic field, and (2) read out the oscillating free induction decay (FID) of the nuclei after such a pulse sequence. The sample is a solution of $^{13}\text{CHCl}_3$, where the qubits are the spins of the hydrogen and carbon nuclei.

The FID is supplied to the software in a complex representation wherein the real part gives the magnetization along the x direction and the imaginary part gives the magnetization along y (both in the rotating reference frame of the nuclei [2]). The complex FID is Fourier transformed to produce the peak spectra depicted throughout this paper (see, for instance, Figure 3).

III. THEORY AND COMPUTATION

The first point to notice upon observing Figure 3 is that the Fourier transform of the FID contains two peaks, because the usual Lorentzian resonance of each nuclear magnetic spin at the Larmor frequency is split into two




FIG. 1: The computer gates the signal of an RF oscillator for each channel, which enables it to send pulses for each nuclei to the RF coil. The free induction decay is detected by the same coil and mixed down with the RF oscillator to be forwarded to the computer.

levels by coupling with the other nuclear spin. (That is, each nucleus provides some magnetic field at the location of the other, such that the energy of each nucleus depends on the state other.) The interaction is described by a scalar coupling constant J in addition to the standard NMR Hamiltonian.

$$H = \hbar\omega_1\sigma_z^1 + \hbar\omega_2\sigma_z^2 + \frac{\hbar}{4}J\sigma_z^1\sigma_z^2$$

This coupling will prove useful in that it allows for easy implementation of quantum operations which require one qubit to affect the other: free evolution for a time $\tau = 1/2J$ results in a conditional $\pi/2$ phase shift between the nuclei.

Now, in order to deal numerically with these peaks, we assign a magnitude to each by integration. Because the peak is theoretically a complex Lorentzian

$$g(\omega) = \frac{a\Gamma}{i(\omega - \omega_0) + \Gamma}$$

whose integral does not converge, we instead integrate the power spectrum (the norm squared of the spectrum), and then take the square root. This quantity, which we hereafter refer to as just the “peak integral” is proportional to the population difference of up versus down spins [2], so naturally it should have some expression in

terms of the diagonal terms of a two-spin density matrix. In fact, the correspondence is quite simple,

$$\begin{pmatrix} a & 0 & 0 & 0 \\ 0 & b & 0 & 0 \\ 0 & 0 & c & 0 \\ 0 & 0 & 0 & d \end{pmatrix} \Leftrightarrow \begin{pmatrix} a-c \\ b-d \\ a-b \\ c-d \end{pmatrix}$$

Where the left is the density matrix of the two qubits (convention: carbon is the second qubit and or “fast index”), and the right column is proportion to the peak integrals (the two for hydrogen, followed by the two for carbon). This correspondence means that converting between peak integrals and density matrices is just a system of linear equations (including the constraint that the trace of a density matrix is 1).

The ability to convert to density matrices is desirable as it enables us to use the tools of quantum information science. The device we will find most useful is the fidelity, which, ranging from 0 (orthogonal) to 1 (equal), quantifies how closely some state matches to another [1]. In the case of a pure basis state $|i\rangle$, the fidelity simplifies to the diagonal matrix element $\langle i|\rho|i\rangle$.

Error propagation calculations are equally simple. Two different techniques are used for selecting the region of integration (*fixed-width*, 15MHz to either side of a peak, and *variable-width*, 5 FWHM to either side). Systematic errors in the procedure are estimated from the difference of these techniques, and directly propagated through the linear algebra. Where stated, statistical errors are estimated by obtaining the same spectra multiple times (nine times for CNOT and nine times for Deutsch-Jozsa).

As a final note, the notation for a $\pi/2$ pulse in use here is $R_{x,y,\bar{x},\bar{y}}^{H,C}$, where the letter in the superscript indicates which nuclei to pulse, and the subscript indicates the axis of rotation (determined by the phase of the pulse).

IV. CALIBRATION

First the pulse phase is calibrated. As discussed, pulses are applied with different phases to rotate the spins about the x and y axes. Since the x and y axes are already defined by the readout (the x -axis is that which produces the real component of the FID), the phases of the pulses should be referenced accordingly.

The procedure for this calibration is simply to apply what the computer imagines to be an R_y^H pulse (such that the spin should be rotated toward the x -axis), and examine the resulting hydrogen peak. The phase of this peak is then taken as a phase reference to be subtracted from all future hydrogen pulses (then the same is done for carbon).

Since our peak integration analysis for the actual experiment, as discussed, only takes a sign from the peaks, it is robust against any accidental phase offsets less than 90° (which easily achieved). Thus the precision of this calibration is thus not relevant to error analysis.

Next, the pulse widths are calibrated. The angle by which a spin rotates during a pulse is proportional to the pulse duration (as well as other factors, including pulse intensity at the sample, to which we do not have direct access), so pulse widths must be calibrated empirically.

The procedure is to apply pulses along the same axis (ie R_Y^H) of different durations, and examine the variation of the peak integral with duration. This plot should give a sinusoid whose first maxima is the $\pi/2$ pulse width (See Figure 2 for details). Since inaccurate pulse widths result in poor execution of computations, this vital calibration is performed at the beginning of every session.

Finally, the duration of free evolution for spin-coupling should be calibrated. As discussed, the J coupling constant provides the frequency difference between the peaks, so a typical double-peaked thermal spectrum such as that in Figure 3 provides us with a measurement of J .

Additionally, we conduct measurements of T_1 (using the spin-echo delay technique) and T_2 (by the width of the lorentzian peaks), to confirm that these decoherence processes are longer than the computational timescale. As they are sufficiently long and thus not important, and there is no space here, we will not discuss this further. (The companion paper by Zhou has more detail on this topic for those interested in decoherence rates.) It suffices to say that T_1 is $67\pm 1\text{ms}$ for hydrogen, $120.1\pm 4\text{ms}$ for carbon. T_2 is $18.6\pm .3\text{s}$ for hydrogen, $18.0\pm .4\text{s}$ for carbon.

V. EFFECTIVE BASIS STATES

In room temperature NMR, one always begins with a thermal distribution of spins rather than the pure state necessary for quantum computing. A technique known as temporal labelling is required to produce effective basis states for initializing NMR computations. While space prevents a lengthy discussion here, good explanations can be found in [1] or [2]. Before attempting computations upon the effective basis states, it is prudent to evaluate the fidelity of this procedure. The effective $|00\rangle$ is shown in Figure 4. (Note: From here on, all peak integrals are normalized to the sum of the magnitudes of the peak integrals in the effective $|00\rangle$.)

By generating each basis state, converting the peak integrals to density matrices, and evaluating the fidelities as explained in Computation, we arrive at Table I. Fidelities of the effective basis states limit how effective any computation carried out on the basis states will be.

State	$ 00\rangle$	$ 01\rangle$	$ 10\rangle$	$ 11\rangle$
Fidelity	$.953 \pm .005$	$0.800 \pm .004$	$.850 \pm .001$	$.690 \pm .014$

TABLE I: Fidelities of effective basis states. Uncertainties are only the systematics as discussed in Computation; since multiple time-consuming runs were not taken for the basis states.




FIG. 2: Pulse-width calibrations for hydrogen (left) and carbon (right). At each duration, the mean of the two peak integrals in the spectrum is plotted, and the difference between the two peaks integrals is taken as an uncertainty (error bars above). A sinusoid is fitted to the curve (χ^2_R values provided), and the location of the fitted $\pi/2$ pulse width is marked by the vertical line.




FIG. 3: (Left) We begin with merely a thermal distribution; shown is a hydrogen spectrum, providing a measurement of J via the distance between fitted peaks. (Right) However, by the techniques discussed here, we implement successively more complex algorithms, finally demonstrating the oscillatory behaviour of Grover's search. (Every three iterations, another maxima is reached, and we see this continue out to 60 iterations. (Different levels of iteration were probed over different overnight runs. More detail may be found in the companion paper by Leo Zhou.)




FIG. 4: Hydrogen (left) and carbon (right) spectra in the effective $|00\rangle$ state generated by temporal averaging. Ideally all of the power spectrum should be in the left peak, however the temporal averaging appears to be off by about 6-7%. This will be discussed quantitatively via the use of fidelities.

VI. CONTROLLED-NOT

Controlled-NOT (CNOT) is the most basic quantum logic gate. On the computational basis, it will flip the second qubit if and only if the first qubit is 1, so it has the following representation as a unitary operator:

$$\text{CNOT} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} \quad (1)$$

In terms of the pulses discussed thus far, CNOT can be written $R_x^C - R_y^C - R_x^C - R_x^H - R_y^H - \tau - R_y^C$.

Demonstration of a high-fidelity CNOT is a common early step for many different quantum computation schemes, so it is natural to implement this in our experiment. We apply the above CNOT pulse sequence to each of the basis states, and record the peak structures.

Three groups of data are plotted for the CNOT analysis in Figure 5.

1. The first is the ideal CNOT output. These were obtained by applying the CNOT gate (Eq 1) to the computational basis. Assuming the basis states are pure, and the CNOT is implemented perfectly, we would expect to see the black bars.
2. The white bars are the actual measured outputs. As we see, they are qualitatively similar to the ideal outputs, except that the large peaks are reduced and the peaks that should be zero do show some noise.
3. The grey bars are what we would expect from an ideal, noiseless CNOT implementation upon our imperfect basis states. That is, these were obtained by taking the effective basis density matrices from Section V, applying the CNOT gate (Eq 1) to them, and converting the resulting density matrices to peaks. They provide a visual way of deciding how much error in the calculation can be attributed to the effective basis and how much to the CNOT implementation.

As a quantitative means of tracking the errors, we tabulate the fidelities below:

Begin	00⟩	01⟩	10⟩	11⟩
End Fidelity	.766	.671	.722	.644
$\pm\sigma_{sys}$	$\pm .004$	$\pm .004$	$\pm .002$	$\pm .002$
$\pm\sigma_{stat}$	$\pm .0008$	$\pm .0003$	$\pm .0009$	$\pm .0004$

TABLE II: Fidelities for the CNOT computation. These follow the trend set by the fidelities of the basis states (Table I), as can also be seen in Figure 5.

VII. DEUTSCH-JOZSA

The first non-trivial algorithm we test is the classic Deutsch-Jozsa (DJ) algorithm, whose description follows.

There is some mystery function $f(x) : \mathbb{Z}_2 \rightarrow \mathbb{Z}_2$. We are given an ‘‘oracle,’’ U , which can compute the mysterious f in this way:

$$U|x, y\rangle = |x, y + f(x)\rangle$$

for any $x, y \in 0, 1$. The goal is to determine whether $f(0) = f(1)$ (f is constant), or $f(0) \neq f(1)$ (f is balanced). Classically this takes two evaluations of f ; to determine whether $f(0) = f(1)$, one must evaluate $f(0)$ and $f(1)$. The Deutsch-Jozsa algorithm exploits superposition to do it with only one evaluation [1].

For our experiment, we employ the following DJ pulse sequence: $R_y^C - R_y^H - U_f - R_y^C - R_y^H$. Moreover, there are four possible functions $f(x) : \mathbb{Z}_2 \rightarrow \mathbb{Z}_2$, and we test the Deutsch-Jozsa algorithm on each. We use the following pulse sequences for the oracles:

$$U_f = \begin{cases} (f(x) = x) & R_x^C R_y^C R_x^C R_x^H R_y^H \tau R_y^C \\ (f(x) = 1x) & R_x^C R_y^C R_x^C R_x^H R_y^H \tau R_y^C R_x^C R_x^C \\ (f(x) = 0) & I \\ (f(x) = 1) & R_x^C R_x^C \end{cases}$$

Results are presented in Figure 6, which is similar in nature and interpretation to Figure 5, and in Table III.

$f(x) =$	x	$1 - x$	0	1
End fidelity	.819	0.653	.867	.845
$\pm\sigma_{sys}$	$\pm .01$	$\pm .003$	$\pm .004$	$\pm .004$
$\pm\sigma_{stat}$	$\pm .002$	$\pm .0006$	$\pm .004$	$\pm .0007$

TABLE III: All Deutsch-Jozsa fidelities are above 80% but for the $f(x) = 1 - x$, as discussed in Conclusion.

VIII. CONCLUSION

Having now reported on all of the basic algorithms, one observation worth note is precisely when the fidelities drop the most sharply. After the basis state rotation from effective |00⟩ (which has 95% fidelity) to any other state, we see much of the fidelity has already been lost, (despite this step being only one or two pulses). Additionally, the second DJ function shows a much lower fidelity than any other DJ. What both the basis rotation and that specific DJ oracle have in common is the presence of a 180° flip, which is implemented by instructing the apparatus to supply one long pulse (rather than two $\pi/2$ pulses). Perhaps future experiments may double-check the validity of the long pulse option.

Nonetheless, we find that basic computations with fidelities of 70% are possible on MIT Junior Lab’s spectrometer. This is nowhere near the precision required

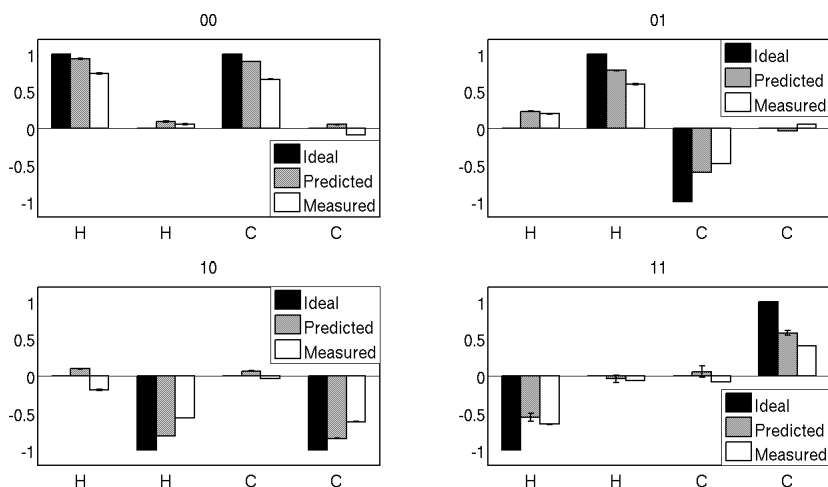


FIG. 5: Evaluation of our CNOT gate.

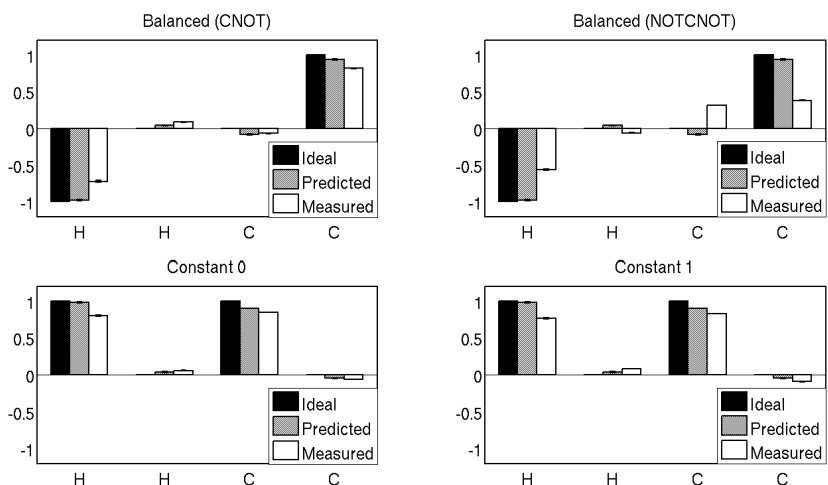


FIG. 6: Evaluation of our Deutsch-Jozsa implementation.

for arbitrarily long, fault-tolerant computation, but it is nonetheless a vital demonstration of several basic components of quantum computation.

In fact, to test the limits of the apparatus, we attempt a more demanding computation, Grover’s algorithm, at various levels of iteration (see Figure 3). Though space does not permit a full analysis here, we note that oscil-

lations continue beyond $k = 60$ iterations, which may perhaps be a record k in Junior Lab.

In summary we have characterized the performance of two basic algorithms on an NMR quantum computer (CNOT and Deutsch-Jozsa), demonstrated the possibility of these basic computations at low-fidelity, and stretched to the limits of what is possible with Grover’s algorithm on the Junior Lab spectrometer.

[1] *Quantum Computation & Quantum Information*. Michael Nielsen, Isaac Chuang. Cambridge University Press, 2000.
 [2] MIT Junior Lab Experimental Guide. <http://web.mit.edu/8.13/www/JLEperiments/JLExp49.pdf>

[3] Gernmecker, Gerd. “Basic Principles of FT NMR.” <http://www.chem.wisc.edu/areas/organic/courses/nmr-gernmecker/chem843-2.PDF>, 1999.