Quantum information storage using tunable flux qubits

2010 J. Phys.: Condens. Matter 22 053201

(http://iopscience.iop.org/0953-8984/22/5/053201)

IOP PUBLISHING JOURNAL OF PHYSICS: CONDENSED MATTER

J. Phys.: Condens. Matter **22** (2010) 053201 (12pp) doi:10.1088/0953-8984/22/5/053201

TOPICAL REVIEW

Quantum information storage using tunable flux qubits

Matthias Steffen, Frederico Brito[1], David DiVincenzo, Matthew Farinelli, George Keefe, Mark Ketchen, Shwetank Kumar, Frank Milliken, Mary Beth Rothwell, Jim Rozen and Roger H Koch

IBM Watson Research Center, Yorktown Heights, NY 10598, USA

E-mail: msteffe@us.ibm.com

Received 14 September 2009, in final form 20 November 2009
Published 15 January 2010
Online at stacks.iop.org/JPhysCM/22/053201

Abstract

We present details and results for a superconducting quantum bit (qubit) design in which a tunable flux qubit is coupled strongly to a transmission line. Quantum information storage in the transmission line is demonstrated with a dephasing time of $T_2 \sim 2.5\ \mu$s. However, energy lifetimes of the qubit are found to be short ($\sim$10 ns) and not consistent with predictions. Several design and material changes do not affect qubit coherence times. In order to determine the cause of these short coherence times, we fabricated standard flux qubits based on a design which was previously successfully used by others. Initial results show significantly improved coherence times, possibly implicating losses associated with the large size of our qubit.

Contents

[1] Present address: Departamento de Física, Universidade Federal de Pernambuco, 50670-901, Recife, Pernambuco, Brazil.

1. Introduction

The quest to build a scalable quantum computer has taken several remarkable steps, particularly using superconducting circuits and Josephson junctions. Long coherence times with high measurement fidelities have been observed using several types of superconducting qubits [1–4], two-qubit interactions have been shown [5–8] and recently entanglement was demonstrated [9], as well as a demonstration of quantum algorithms [10]. Although it is unclear which approach ultimately leads to the development of a practical device, these experiments have provided invaluable insights into practical concerns of several different approaches. From these experiments it is possible to determine design features that help facilitate the fabrication of a scalable qubit.

After carefully considering various aspects including the ability to measure, address and couple qubits we believe the following is a highly desirable set of characteristics for a scalable design.

(1) Short and long term storage. It should be possible to store each qubit in such a way that it is insensitive to unavoidable cross-couplings and decoherence. This

significantly reduces the engineering requirements for selective excitation of any one of the qubits.

(2) Tunable qubits. The resonance frequency of each qubit should be adjustable by an external control parameter. This reduces a potentially large and expensive overhead in external microwave sources and helps reduce cross-couplings.

(3) Tunable coupling. The coupling between qubits should be controllable with a large ON/OFF ratio. This significantly improves the prospects of refocusing unwanted coupling between unselected qubits.

(4) Short and long range communication. The ability for qubits to communicate over short (e.g. nearest-neighbor) and long (e.g. many inter-qubit distances) ranges is a desirable feature for a scalable architecture. A design only employing nearest-neighbor interactions suffers from a large overhead, which we do not believe to be a fruitful avenue.

Here we present an adaptation of a tunable flux qubit coupled to a harmonic oscillator that our group previously described [11], and which satisfies our set of characteristics. We shall refer to this qubit as the 'IBM qubit' for the remainder of this review for simplicity. In this design, the first two levels of a harmonic oscillator are used as a means to store the qubit information. Large quality factors ($Q > 10^5$) of harmonic oscillators have been measured (e.g. [12, 13]) and we thus anticipate long coherence times for storage. Single-and two-qubit gates are implemented with microwave pulses, after adiabatically moving the quantum information from the harmonic oscillator into the flux qubit, or by quasi-dc pulses [14]. By using some qubits as communication ports, it is possible to use the harmonic oscillator for short and long range communication.

We demonstrate the basic operation of the IBM qubit by storing quantum information in the harmonic oscillator. The lifetime is found to be $T_2 \sim 2.5\ \mu$s, consistent with measured values of quality factors of coplanar waveguide resonators. However, lifetimes of the qubit itself are found to be quite short, of the order of 10 ns, and are much shorter than had been predicted. A range of design changes aimed at reducing or at least altering the coupling to external noise sources as well as using different substrates had no effect on qubit coherence times. In order to explain the short coherence times we fabricated and tested 'non-tunable' flux qubits based on a design which has been successfully used by others in the past [15] and found dramatically improved coherence times. In the following sections we shall discuss details of the IBM qubit including experimental results as well as recent results of the 'non-tunable' flux qubits.

2. Background

Recently, a design for a tunable flux qubit was proposed [11], in which all operations are implemented using combinations of adiabatic and highly non-adiabatic pulses. While the approach contains novel ideas for generating quantum gates that are worthwhile in pursuing on their own, it also requires operating

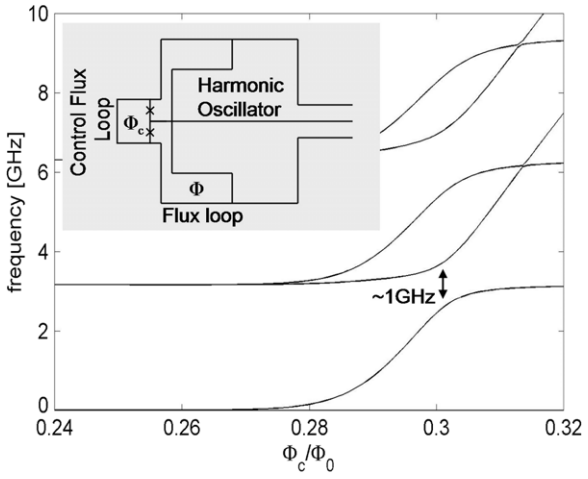

Figure 1. Energy levels of the tunable flux qubit coupled to a harmonic oscillator (see inset) for $\Phi = \Phi_0$. At large values of the control flux the qubit can be stored, or parked, in the lowest two energy levels of the harmonic oscillator ($\omega_{HO} \sim 3.1$ GHz). By adiabatically pulsing the control flux, the qubit can be transferred into the flux qubit where single-qubit operations are implemented via microwave pulses. The arrow indicates the location of the avoided level crossing where the qubit most strongly couples to the harmonic oscillator. Parameters used are the critical current $I_0 = 1\ \mu$A, the large loop inductance $L = 600$ pH, the small loop inductance $L_c = 32$ pH and effective junction capacitance $C = 50$ fF.

the qubit away from its degeneracy point where coherence times are strongly limited by $1/f$ flux noise. Such flux noise introduces significant dephasing [16]. Although a novel quantum error correction scheme has recently been shown to protect qubits from noise if it is almost entirely of the phase type [17], it may nonetheless prove advantageous to operate the device in a regime where it is less sensitive to these fluctuations. In such a scheme qubit gates are generated using microwave pulses, similar to the operating scheme of other superconducting qubits [1–8].

The tunable flux qubit design we introduce here is the same as that shown in [11] but we eliminate the junction that is shared between the two large loops. This reduces the sensitivity of the qubit's resonance frequency to control flux by a roughly a factor of two and also reduces the sensitivity to critical current fluctuations by a factor of a few. Otherwise there are no significant qualitative differences between the two- and three-junction approaches. But because the two-junction approach can be simulated using fewer computational resources and because of the reduced sensitivity of this qubit, we shall only consider its properties throughout this review.

The circuit of the qubit is shown in figure 1. It consists of three loops and is controlled via two external fluxes, which we shall call the *flux* Φ and the *control flux* Φ_c. The two large loops, each with inductance $L \sim 600$ pH, are arranged gradiometrically so a common flux applied to both loops has no net effect. The flux Φ threads one of the large loops. The small loop with inductance $L_c \sim 32$ pH is interrupted by two Josephson junctions, each with critical current I_0, and is threaded by the control flux Φ_c. The potential energy of the

system can be written as

$$U = -\frac{I_0\Phi_0}{2\pi}\cos(\delta_1) - \frac{I_0\Phi_0}{2\pi}\cos(\delta_2) + \frac{1}{2L}\left(\Phi + \frac{\delta_1\Phi_0}{2\pi}\right)^2 + \frac{1}{2L}\left(-\frac{\delta_2\Phi_0}{2\pi}\right)^2 + \frac{1}{2L_c}\left(\Phi_c + \frac{(\delta_1-\delta_2)\Phi_0}{2\pi}\right)^2, \quad (1)$$

where $\Phi_0 = h/2e$ is the quantum flux and δ_i represents the superconducting phase drop across the ith Josephson junction.

The wavefunctions for this two-dimensional potential can be solved by numerical integration of the Schrödinger equation. However, because $L_c \ll L$, it is possible to approximate this potential with a one-dimensional form [18], which gives a more intuitive understanding of the energy landscape. Assuming that $L_c \ll L$, the last term in the potential energy dominates unless $\Phi_c = -\frac{(\delta_1-\delta_2)\Phi_0}{2\pi}$, which constrains the phase differences across the two Josephson junctions:

$$\delta_2 = \frac{\Phi_c 2\pi}{\Phi_0} + \delta_1. \quad (2)$$

Given this constraint, the approximate one-dimensional potential can be written as

$$U = -\frac{I_0\Phi_0}{\pi}\cos\left(\delta_1 + \frac{\Phi_c\pi}{\Phi_0}\right)\cos\left(\frac{\Phi_c\pi}{\Phi_0}\right) + \frac{1}{L}\left(\frac{\Phi+\Phi_c}{2} + \frac{\delta_1\Phi_0}{2\pi}\right)^2. \quad (3)$$

Note that, in order to compute wavefunctions using this approximated potential, one must rescale the mass [18]. A more formal and more accurate computation, which reduces the two-dimensional potential, follows the description of [18]. However, we find that the most significant difference between the approximation and the more detailed calculation is a small rescaling in Φ_c. Only minor qualitative and quantitative differences exist, and we hence use the simple form (equation (3)) for the remainder of the review, for simplicity.

It is clear from the reduced potential that it is equivalent to that of a flux qubit except that the Josephson energy is tunable via Φ_c and that the flux must be biased at $\Phi = \Phi_0$. This is, in fact, quite similar to the tunable flux qubit design used in reference [19]. We now couple the main loops to a coplanar transmission line, which can be regarded as a harmonic oscillator (HO) of frequency ω_{HO} and is used to store the quantum information when the qubit is not addressed. The coupling term is equal to $M_{qho}\hat{I}_q\hat{I}_{HO}$, where M_{qho} is the effective mutual inductance between the qubit and the harmonic oscillator, $\hat{I}_q$ is the circulating current operator of the qubit and $\hat{I}_{HO}$ is a current operator of the harmonic oscillator. The wavefunctions and energies for this two-dimensional potential can be solved numerically; such calculations are described in [20].

Thereby, the computational *logical states* $|0\rangle$ and $|1\rangle$ are encoded using the lowest two eigenstates of the qubit–HO system which, as we shall describe below, can be tuned, through changes in the control flux Φ_c, to have either qubit or HO state characteristics.

In order to analyze the dynamics of the energy levels as a function of Φ_c, let us focus on the case $\Phi = \Phi_0$, illustrated in figure 1. For that condition (indeed, whenever $\Phi = n\Phi_0$, where n is an integer) one can show that the qubit potential equation (3) presents a perfectly symmetric structure. Furthermore, for small values of $\Phi_c \sim 0$, we observe that the qubit potential has a double-well form with a high barrier, which leads the qubit states to be nearly degenerate. Since the qubit spectral gap is expected to be much smaller than $\hbar\omega_{HO}$, we find that the 0–1 logical states are determined solely by the qubit states in this regime. As the tunneling time between the left and right wells is long (virtually infinite), one may think of the system as being classical with its state being either in the left $|L\rangle$ or right $|R\rangle$ well of the potential, corresponding to two macroscopically different circulating currents which can be sensed by a magnetometer.

As we increase Φ_c, the barrier height is reduced, leading to coupling between the left and right wells and thus an exponential growing of the tunneling amplitude between the wells, so that the qubit eigenfunctions become the symmetric and anti-symmetric combinations of the left and right states. Because the qubit is coupled to a harmonic oscillator, when the qubit and HO energy scales are in resonance, energy can be effectively exchanged between the two systems and gives rise to an avoided level crossing in the energy diagram. The coupling strength is equal to the size of the avoided level crossing and is determined by a mutual inductance between the qubit and the harmonic oscillator, as discussed earlier. Upon further increasing Φ_c, the lowest two energy levels become those of the harmonic oscillator.

We note that differences in the critical currents of the two Josephson junctions lead to an asymmetric potential for $\Phi = \Phi_0$, which can always be changed to a symmetric one by adjusting the flux. Although this offset is a function of the applied control flux, we believe this effect not to have detrimental consequences.

3. Operation of the qubit: single qubits

3.1. Parking

The qubit is stored in the first two levels of the harmonic oscillator for large values of the control flux ($>0.32\Phi_0$). We believe this is a good method for information storage. Considering the usual trade-off between the ability to manipulate (and measure) the qubit and coherence times, the best approach, in our opinion, is to move the qubit to a point in parameter space where it is not straightforward to manipulate or measure the system. We refer to storing the quantum information as 'parking' the qubit [20]. We believe parking is particularly useful in a scalable device because it is likely that qubits will generally spend a significant amount of time idle.

Large quality factors in excess of 10^5, corresponding to coherence times of longer than 10 μs for our operating frequency, have been measured in harmonic oscillators of various kinds (e.g. [3, 12, 13]), and are therefore suitable for information storage. We note that quality factors beyond 10^6 appear to be prevented by phase noise [12] believed to be due to two-level fluctuators. Although this phase noise is unlikely to be a limitation in simple experiments, a scalable device should ultimately address this problem.

The qubit is parked and unparked via adiabatic pulses applied to the control flux. The avoided level crossing around $\Phi_c = 0.3\Phi_0$ determines the speed of these pulses. Assuming a splitting of about 1 GHz and linear pulse shapes, pulse durations of 10 ns give fidelities of about 0.99 with larger fidelities accessible with slower or shaped adiabatic pulses. During unparking, the quantum information is transferred back into the first two energy levels of the tunable flux qubit. Any unwanted phase accumulation during the parking and unparking process can be refocused using standard spin–echo type sequences. Parking is useful not just for storing quantum information but also as a means to facilitate communication between qubits.

3.2. Single qubit gates

Once unparked, single qubit gates can be applied via microwave pulses tuned to the difference of the lowest two energies of the system, and applied on the flux line. The fidelity of the microwave single qubit gates is limited by leakage to neighboring energy levels, Bloch–Siegert oscillations [21] and off-resonance effects due to $1/f$ flux noise. These are concerns that are always present in varying degrees whenever microwave pulses are used to implement quantum gates. We shall not detail the implications for the IBM qubit here but solutions include techniques such as pulse shape optimization [22, 23]. Single-qubit operations can also be implemented by using quasi-dc pulses as detailed in [14].

3.2.1. Energy relaxation. Energy relaxation, characterized by a time T_1, is the result of high frequency fluctuations. In general, both the control flux and the flux biases will have high frequency fluctuations. The limit to energy lifetimes due to fluctuations in flux can be computed using [24]. In general, the T_1 times are a function of control flux. Given a mutual inductance to the flux line of about 1 pH and a 50 Ω environment, we estimate a lifetime $T_1 \approx 0.25\ \mu$s when $\Phi_c \sim 0.3\Phi_0$, which should be sufficient for initial tests. The matrix elements relevant for fluctuations in the control flux can also be computed, and for the lowest two energy levels is essentially zero when the qubit potential is symmetric, so that energy relaxation is dominated by fluctuations in the flux line. However, for the first two *excited* states the matrix element due to control flux fluctuations is non-zero around $\Phi_c \sim 0.3\Phi_0$, i.e. when the qubit is coupled to the harmonic oscillator. Therefore, care must be taken to avoid thermal excitations, because the frequency difference is only about 1 GHz.

It is also convenient at this point to express the qubit lifetime as an RC time constant. Such an approximation is not commonly used for flux qubits: however, it produces the correct result to within a factor of a few. It is therefore quite useful when estimating coherence times without employing detailed calculations and also ties energy relaxation to classical electromagnetic concepts.

As explained earlier, the IBM qubit can be approximated by a one-dimensional potential energy that is identical to that of an RF SQUID except for the critical current is tunable via Φ_c. The energy lifetime for an RF SQUID flux qubit is given by (see, e.g., [25])

$$T_1 = \left(\frac{2\pi}{\Phi_0}\right)\frac{\hbar}{2\omega_{01}}\frac{\coth(\hbar\omega_{01}/2k_BT)}{|\langle 0|\delta|1\rangle|^2 \mathrm{Re}\{Y(\omega_{01})\}} \tag{4}$$

where Y is the admittance looking out from the qubit and δ is the superconducting phase operator. For a harmonic system the matrix element can be computed with the result $\langle 0|\delta|1\rangle = (2\pi/\Phi_0)\sqrt{\hbar/2\omega C}$. Substituting this into equation (4), and assuming $k_BT \ll \hbar\omega_{01}$, gives $T_1 = C/\mathrm{Re}\{Y\}$ which reduces to RC when viewing the real part of the admittance as an effective resistance at the qubit frequency ω_{01}. We note that, although flux qubits have a quartic potential, the harmonic approximation remains valid to within a factor of a few over a wide range of values for critical currents and capacitances that give rise to frequencies ω_{01} of the order of a few GHz. This remarkable fact allows the use of a simple equation to calculate energy relaxation times to within an order of magnitude.

As a specific example, consider energy relaxation due to flux coupling to external bias coils. For a three-junction flux qubit T_1 has been derived quantum mechanically as [26]

$$T_1 = \frac{Z_0\hbar^3\omega}{2(\Delta h)^2(MI)^2}. \tag{5}$$

Assuming the qubit is operated at its sweet spot ($\Delta = \omega$), converting the critical current to the Josephson inductance ($I = \Phi_0/2\pi L$), and making a rough harmonic approximation ($\Phi_0/2\pi = \sqrt{\hbar/2\omega C}$), equivalent to setting the matrix element $|\langle 0|\delta|1\rangle|$ to unity, gives the following analytical equation for T_1:

$$T_1 = Z_0C(L/M)^2. \tag{6}$$

This equation can be derived in a straightforward manner by noting that the inverse of the admittance seen by the qubit is $Z_0(L/M)^2$. Multiplying by C then recovers equation (6). This result highlights the known fact that the three-junction flux qubit is not so different from the RF SQUID flux qubit except for the two Josephson junctions acting as classical inductors.

We can use equation (6) to approximate T_1 for the IBM qubit and compare to the quantum mechanically derived number. The designed loop inductance is $L = 600$ pH and the capacitance of a shadow-evaporated Josephson junction is about $C = 10$ fF. Assuming $Z_0 = 50\ \Omega$ and $M = 1$ pH gives $T_1 \sim 180$ ns, compared to $T_1 \sim 0.25\ \mu$s when computed quantum mechanically. To within an order of magnitude they agree with each other.

In a later section we make use of the connection of relaxation times to classical electromagnetics because it highlights just how important classical electromagnetic considerations are for flux qubits.

3.2.2. Measurement. The qubit is measured by an external SQUID whose critical current is dependent on the flux generated by one of the large loops. Decoherence due to the SQUID during the qubit operation is minimized by biasing the SQUID so that small fluctuations in the SQUID bias do not generate circulating currents which couple to the qubit.

The measurement procedure consists of adiabatically setting the control flux to zero. By further changing the flux bias slightly, it is possible to adiabatically change the symmetric and antisymmetric states to the macroscopically distinct states 'left' and 'right'. Note that, because of undesired parasitic mutual inductances, a slight bias change is already present during the control flux pulse and is sufficient for a high fidelity measurement without intentionally applying a flux pulse.

At zero control flux the barrier between the left and right well is maximized. By choosing the critical currents of the two junctions to be more than about 1 μA, the tunneling times can be made virtually zero. Fluctuations in the flux bias current destroy any coherence, collapsing the wavefunction either into the left or the right well, corresponding to two different macroscopic fluxes generated by the circulating current in the large loops.

The SQUID generally contributes to energy dissipation but its effect can be made negligible by arranging it geometrically such that current noise in the SQUID does not generate a flux which couples to the qubit [27]. We designed the SQUID in this manner and have experimentally confirmed the isolation of the qubit from SQUID bias currents. Furthermore, the SQUID is designed to have a high self-resonance frequency (>20 GHz) such that the qubit states are sufficiently decoupled from the SQUID states.

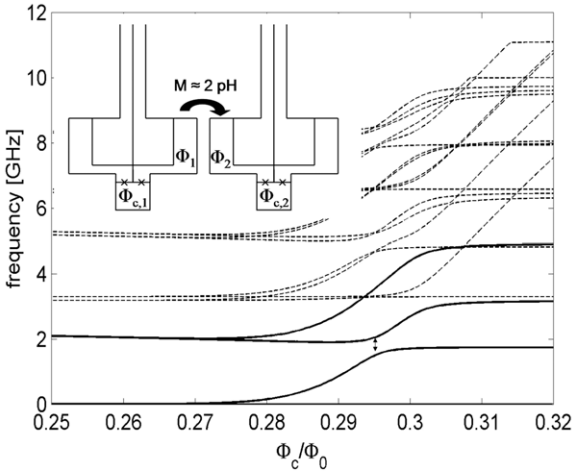

Figure 2. Energy spectrum of two coupled qubits as a function of the control flux of qubit 1 and a control flux bias of $\Phi_c = 0.295\Phi_0$ for qubit 2. The avoided level crossing at $\Phi_c = 0.295\Phi_0$ indicates coupling between the qubits, indicated by the arrow. A microwave pulse tuned to the size of the splitting, and applied to either qubit's control flux, excites a transition between these two energy levels and corresponds to a two-qubit gate. The inset depicts the physical set-up for coupling qubits. The solid lines correspond to the two-qubit energy levels. Both qubits and harmonic oscillators are nominally identical with $\omega_{HO} \sim 3.1$ GHz as before.

4. Operation of the qubit: two coupled qubits

4.1. Qubit–qubit interaction

Two tunable flux qubits can be coupled via a mutual inductance of a few pH between any two of the large loops, as depicted in figure 2. Qualitatively, this set-up leads to a flux bias of each qubit that is dependent on the external flux bias Φ, *and* the circulating current in the large loop of the second qubit. Previous work used such a set-up to generate coupled qubit gates using adiabatic pulses [14], but here we shall introduce gates that are based on pulsed microwave operation.

We shall focus on the case where the flux of both qubits is set to Φ_0 while Φ_c is set such that both qubits nominally have the same resonance frequency. It can be shown that this set-up leads to a coupling of the form $\sigma_x \otimes \sigma_x$, where σ_x is the x Pauli matrix. When the resonance frequency of both qubits is much larger than the interaction strength, this coupling reduces to an interaction between the $|01\rangle$ and $|10\rangle$ of the system, resulting in an avoided level crossing of the energy levels.

More formally, the interaction Hamiltonian is of the form $M_{12}\hat{I}_1\hat{I}_2$, where $\hat{I}_1$ and $\hat{I}_2$ are the circulating current operators for the two qubits. A full simulation of the system becomes computationally expensive. Therefore, we shall make a few approximations and attempt to write the Hamiltonian in Pauli matrix formulation. A single qubit coupled to a harmonic oscillator can be written in a Pauli matrix form by approximating the harmonic oscillator as a two, leading to a four-level Hilbert space [14]. If the magnitude of the qubit–qubit coupling can be computed, the Hamiltonian of the system can be computed in a straightforward manner, thereby allowing time-dependent simulations without significant computational resources. The magnitude of the qubit–qubit coupling can be estimated by computing $M_{12}|\langle I_l 1\rangle - \langle I_r 1\rangle||\langle I_l 2\rangle - \langle I_r 2\rangle|$, where $\langle I_l 1\rangle$ and $\langle I_r 2\rangle$ are the expectation values of the circulating currents for the states $|L\rangle$ in qubit 1 and $|R\rangle$ in qubit 2. These are easily computed using the simulation code for a single qubit.

The energy levels of the coupled system are shown in figure 2 for $M_{12} = 2$ pH as a function of control flux for qubit 2. The solid lines correspond to the four energy levels of the two qubits. The control flux of qubit 1 was chosen to be $0.295\Phi_0$. Both qubits had to be slightly detuned from the flux condition $\Phi = \Phi_0$ [14], which is of little experimental consequence. The predicted avoided level crossing is clearly visible at the control flux value of $0.295\Phi_0$ for the second qubit. For simplicity, we chose the two qubits and harmonic oscillators to be identical. In practice, the harmonic oscillators will have a slightly different resonance frequency and, in fact, this is needed to avoid cross-coupling between neighboring parked qubits. However, this detail is not necessary to understand the basics of qubit–qubit interaction when unparked.

4.2. Two-qubit gates

Two-qubit gates are generated as follows. Both qubits are initially parked and do not interact. We first unpark one of the qubits to some resonance frequency, and then adiabatically unpark the other qubit to the same resonance frequency. This corresponds to adiabatically moving along the energy levels shown in figure 2. A microwave pulse on either qubit's control flux tuned to the frequency difference between the second and third energy level then excites transitions between these two

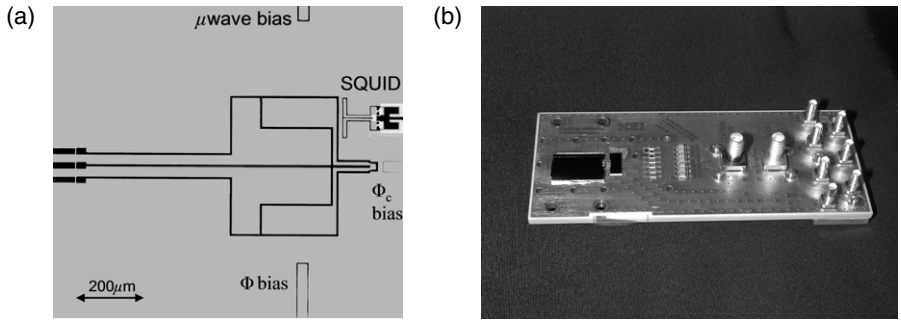

Figure 3. (a) Micrograph of the qubit. The qubit and harmonic oscillator are easily recognizable from the sketch in figure 1. The Josephson junctions are not visible on this scale. (b) Photograph of the circuit board. The qubit chip is mounted on the left side. A separate smaller chip to the right of the qubit chip allows for additional filtering. The control flux and flux lines are 50 Ω transmission lines and terminated by SSMA connectors. Three of the seven SSMC connectors are used for SQUID biasing while the rest are currently unused. Two rows of surface mounted components can be populated to provide discrete low pass filtering for the SQUID bias lines.

energy levels. Because the unparking procedure is adiabatic we can still label the first four energy levels as $|00\rangle$, $|10\rangle$, $|01\rangle$ and $|11\rangle$. Therefore, the microwave pulse generates transitions between the logical states $|01\rangle$ and $|10\rangle$, which correspond to a two-qubit gate.

The matrix element for this excitation can be computed and we found it to be roughly equal to that of single-qubit gates. Thus, there should not be significant additional engineering requirements to apply microwaves on the control flux. It is also possible to drive the transitions between $|00\rangle$ and $|11\rangle$ but only for small values of the control flux and strong microwave amplitudes. Therefore, we feel that short pulses ($<$20 ns) are not easily achievable without significant leakage. It is not clear whether there are convincing advantages to using this transition but we should be aware of this possibility.

The ability to generate two-qubit gates with microwaves has several advantages. The coupling strength is determined by the amplitude of the drive field and therefore can have a large ON/OFF ratio. All tools and techniques available to improve gate fidelities for single qubits are directly applicable for the two-qubit gate as well. Leakage out of the two-qubit manifold or other unwanted transitions within the two-qubit system can be minimized by using a shaped pulse; Bloch–Siegert oscillations can be largely avoided as well. As with single-qubit gates, two-qubit gates can also be implemented using quasi-dc pulses as outlined in [14] to generate controlled phase gates.

5. Experimental results

We next describe experimental results of the IBM qubit by first outlined the fabrication process flow. The qubit was fabricated on a 200 mm high resistivity ($>$1000 Ω cm) silicon (Si) wafer with a 50 nm layer of dry thermal silicon oxide (SiO_2). The niobium layer for the bias lines is 200 nm thick and deposited by physical vapor deposition (PVD). The subsequent patterning is done using deep UV lithography followed by a reactive ion etch (RIE). The dielectric layer used for the SQUID shunting capacitor is made of silicon nitride (SiN). It is deposited using plasma-enhanced chemical vapor deposition (PEVCD) and etched using an RIE. In order to fabricate the Josephson junctions a separate mask must be created by first depositing 650 nm of PMMA and baking it for 1 h at 175 °C. This is followed by depositing 50 nm of germanium (Ge) and another layer of PMMA that is subsequently patterned by e-beam lithography. After developing the top PMMA layer the germanium is etched using a CF_4/Ar plasma. The bottom PMMA layer is then etched using an oxygen plasma etch. Finally, the substrate surface is pre-cleaned using an ion mill, followed by depositing a 35 nm aluminum layer at a 25° angle from vertical. The junctions are then defined by oxidation, typically around 20 Torr seconds, and the final deposition of a 85 nm aluminum layer at −25° from vertical. This process allows a total shift between the two aluminum layers of 0.5 μm.

5.1. Parking

An optical micrograph of the fabricated qubit is shown in figure 3(a). Compared to other flux qubits [15, 16, 28] the IBM qubit is quite large, measuring as much as 0.5 mm across the two large flux loops. The device is mounted on the inside of an RF tight box on top of a superconducting ground plane to shield the qubit from normal metal. The box integrates a carefully designed circuit board featuring 50 Ω impedance transmission lines to SSMA connectors for the flux and control flux biases. The SQUID connections do not need to be impedance-matched to 50 Ω because neither the SQUID current bias nor voltage bias lines need to support high speed pulses. A picture of this circuit board is shown in figure 3(b).

The coaxial lines connecting to the circuit board were filtered using low temperature bronze powder filters. Low temperature attenuators are used to reduce low and high frequency thermal noise reaching the qubit, at the cost of requiring large room temperature bias currents for the qubit flux and control flux. Attenuators are therefore not optimal, but because they are impedance-matched to 50 Ω they are a good choice when high speed signals are required to reach the qubit. The full circuit diagram is shown in figure 4.

The qubit was first calibrated by mapping out the flux–control flux space as in [11]. For certain ranges of these fluxes the qubit has a hysteretic response, indicating a double-well

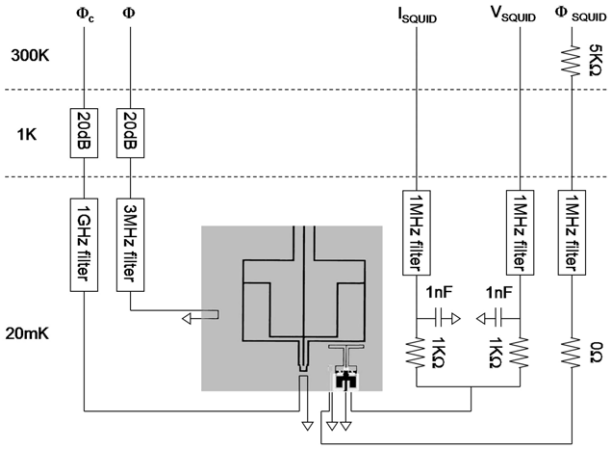

Figure 4. Circuit diagram. The low temperature filters for the flux and control flux lines are bronze powder filters to provide excellent filtering of room temperature noise and ensuring low electrical temperatures of the qubit. The control flux filter was designed to be 50 Ω impedance-matched because high speed pulses (∼1 GHz) are passed through it. The RC filter for the SQUID current and voltage lines is located on the PC board and implemented with discrete surface mounted components.

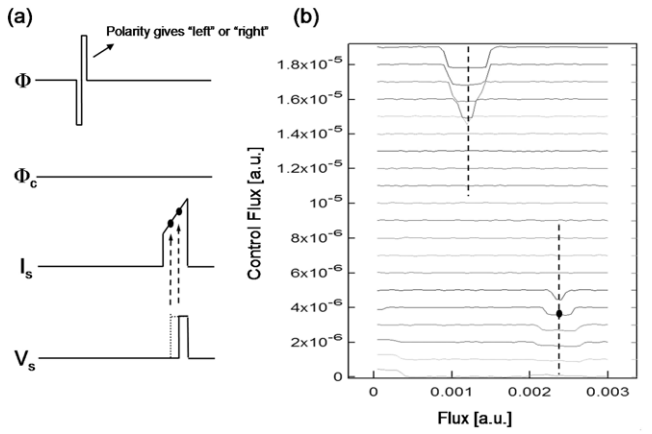

Figure 5. Flux mapping of the qubit. (a) Basic pulse sequence. Because a linear ramp for the SQUID read-out is used, the critical current of the SQUID correlates with time, which is what is measured in the actual experiment. (b) Experimental results. For each control flux value the difference of the SQUID response for the two oppositely polarized reset pulses is measured versus flux. The data is vertically offset for various control flux values. The fluxes over which the qubit is hysteretic are clearly visible. The qubit potential is symmetric at the mid-point of these areas, indicated by the vertical dashed line, and corresponds to a flux bias of an integer multiple of Φ_0.

potential energy such that the qubit is stable in either the 'left' or the 'right' well. It is possible to probe the hysteresis by applying a large flux pulse to the qubit which resets the qubit in one or other of these wells, depending on the polarity of the applied flux pulse, similar to the resetting of phase qubits [7]. This is followed by measuring the critical current of the SQUID which changes depending on which well the qubit is located in. This measurement is repeated approximately 1000 times to gather sufficient statistics. Finally, the difference in SQUID critical currents for the two reset pulses is recorded and plotted. The full pulse sequence and the result of the flux map is shown in figure 5.

The hysteretic behavior of the qubit is clearly visible. This map allows straightforward identification of the flux bias corresponding to an integer multiple Φ_0 by locating the mid-point of the flux over which the qubit is hysteretic. This plot also shows that increasing the control flux decreases the width of the hysteretic region in flux, indicating lower barrier heights.

In the next series of experiments the qubit parking experiment is performed as follows, similar to [11]. The flux is biased near Φ_0 and the control flux is biased in a regime where the qubit is barely hysteretic (indicated by the black dot in figure 5). We then apply a control flux pulse with a fast rise time (1 ns), followed by a delay, followed by another fast control flux ramp down back to the original bias value. The SQUID is then read-out to determine the qubit location (left or right well). This experiment is repeated about 250 times to gather sufficient statistics. Finally, the SQUID signal is plotted versus delay time, control flux pulse amplitude and flux bias near Φ_0. Whenever the flux is very near Φ_0 and the proper control flux pulse amplitude is applied, sinusoidal oscillations in the signal are observed as shown in figure 6. In addition we observe a shift in phase by 180°, depending on which well the qubit is initialized in. These observations are consistent with predictions, as we now discuss.

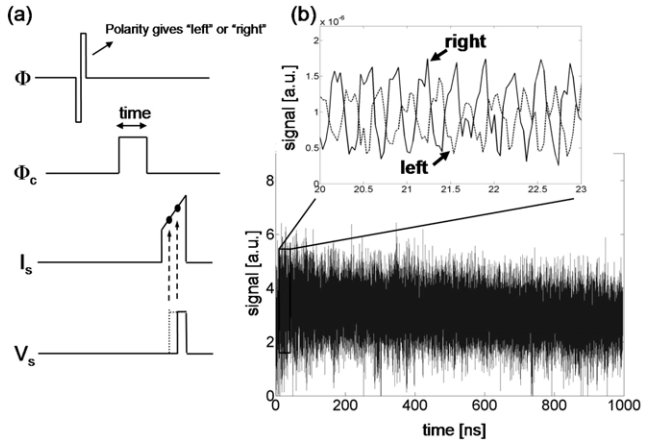

Figure 6. Parking experiment. (a) Pulse sequence. The qubit is biased at the calibrated flux and control flux values. A fast control flux pulse of varying time duration is applied. (b) When the flux bias and control flux pulse amplitude are optimal, oscillations in the SQUID signal versus pulse duration are observed. The phase of the oscillations changes 180°, depending on which reset pulse is first applied, consistent with predictions. The oscillations correspond to a frequency of approximately 3 GHz, which is the design value for the harmonic oscillator.

When the qubit potential is symmetric ($\Phi = \Phi_0$) and hysteresis is observed, the qubit can be initialized in the left (right) well. The first half of the fast control flux pulse then non-adiabatically lowers the potential barrier so that the system is no longer in its eigenstate. In fact, if at the initialization point the state $|L\rangle$ ($|R\rangle$) nearly represents the ground (first excited) state, the application of the first half of the pulse will implement the non-adiabatic transformation

$U_1|L\rangle = \frac{1}{\sqrt{2}}(|g\rangle + |e\rangle)$ ($U_1|R\rangle = \frac{1}{\sqrt{2}}(|g\rangle - |e\rangle)$), where the unitary operator U_1 represents the action of the first half of the pulse. This result means that the non-adiabatic pulse essentially created a superposition of the first two instantaneous eigenstates $|g\rangle$ and $|e\rangle$. This is quite similar to the first charge qubit experiments [29]. In the spin language one could envision the state $|L\rangle$ ($|R\rangle$) as being along the z axis ($-z$ axis) of the Bloch sphere. The non-adiabatic pulse quickly turns on a magnetic field about the x axis and the spin begins precessing about it at a frequency equal to the difference between the first eigenenergy and the ground state. This first step corresponds to moving the control flux from about $0.2\Phi_0$ to $0.295\Phi_0$ in figure 1.

If the control flux pulse amplitude is sufficiently large, the second half of the pulse lowers the potential barrier further to a point in control flux where the instantaneous eigenstates are those of the harmonic oscillator. Under this operation, the system has to pass through the avoided level crossing, which constitutes a crucial point in its evolution: if, compared to the avoided crossing gap, the system undergoes a fast evolution at this point, the probability of transitions to the second and higher excited states become appreciable. This *leakage* to higher excited states reduces the final probability of finding the system at its two lowest lying states, thus decreasing the visibility observed from the measurements. Since we envision a qubit–HO coupling of the order of 1 GHz, it is sufficient to have a rise time of the control flux of a few ns, in order to implement this operation as an adiabatic evolution. Thus, our final states after the pulse second half are given ideally by $U_2U_1|L\rangle = \frac{1}{\sqrt{2}}(|g\rangle + \mathrm{e}^{-\mathrm{i}\theta}|e\rangle)$ ($U_2U_1|R\rangle = \frac{1}{\sqrt{2}}(|g\rangle - \mathrm{e}^{-\mathrm{i}\theta}|e\rangle)$), where the unitary operator U_2 represents the physical implementation of the second half of the pulse and θ is the relative phase accumulation acquired during this process. In spin language the spin continues to precess with the frequency ω_{10}. This step corresponds to the control flux continuing to change from $0.295\Phi_0$ to $0.33\Phi_0$.

After applying the control flux pulse described above, the system eigenstates $|g\rangle$ and $|e\rangle$ will evolve in time (i.e. the spin keeps precessing) with a relative frequency given by that of the HO. After spending a delay time Δt at the parking point, their relative phase will be given by $\mathrm{e}^{-\mathrm{i}\omega_{\mathrm{HO}}\Delta t}$. Consequently, at this point in the evolution, our initial states have evolved to: $U(\Delta t)U_2U_1|L\rangle = \frac{1}{\sqrt{2}}(|g\rangle + \mathrm{e}^{-\mathrm{i}\theta}\mathrm{e}^{-\mathrm{i}\omega_{\mathrm{HO}}\Delta t}|e\rangle)$ and $U(\Delta t)U_2U_1|R\rangle = \frac{1}{\sqrt{2}}(|g\rangle - \mathrm{e}^{-\mathrm{i}\theta}\mathrm{e}^{-\mathrm{i}\omega_{\mathrm{HO}}\Delta t}|e\rangle)$, where the unitary transformation $U(\Delta t)$ determines the system evolution during this stage. Because the harmonic oscillator frequency ω_{HO} is fixed and does not change much with control flux we refer to this location in control flux as parking the qubit (similar to [11]). Finally, the qubit is then unparked by applying the previous operations in reverse order. The falling edge of the control flux pulse first adiabatically unparks the qubit and then non-adiabatically back to the original control flux value. Hence, our final states after applying the reverse control flux pulse are given by $U_1^\dagger U_2^\dagger U(\Delta t)U_2U_1|i\rangle$ ($i = L, R$). It is simple to show that those states do not depend on the relative phase θ, and that the probability of the final state being found in the state $|L\rangle$ is $P_{i\to L}(\Delta t) = |\langle L|U_1^\dagger U_2^\dagger U(\Delta t)U_2U_1|i\rangle|^2 = \frac{1}{2}(1 \pm \cos(\omega_{\mathrm{HO}}\Delta t))$. Thus we observe that the probability of the final state corresponding to $|L\rangle$ varies sinusoidally with frequency ω_{HO}, and that, depending on the initial state, the oscillations should shift in phase by 180° as verified experimentally. In the spin language the ramp-down of the control flux pulse corresponds to turning off the magnetic field non-adiabatically. The probability of point up or down then corresponds to the $|L\rangle$ and $|R\rangle$ states.

The data shows a long T_2 lifetime when the qubit is parked. Although clearly the signal decays for long delays near 1 μs it is not possible to discern a clear exponential decay. By fitting the data to an exponentially decaying sinusoid, nonetheless we obtain a dephasing time of approximately 2.5 μs, corresponding to a quality factor of the harmonic oscillator of $Q = 50\,000$. Such quality factors at low temperature and power are expected [13]. This experiment is an encouraging first step in building a quantum memory for storing quantum information.

The visibility of the oscillations is approximately 80%. We attribute the reduced visibility to energy loss during the ramp-up and ramp-down periods of the control flux pulse as well as non-adiabatic effects near $\Phi_C \sim 0.295\Phi_0$. Optimizing the visibility to 100% is not our goal in these first tests and therefore we did not spend time to track down methods to increase the visibility.

The described experiment can be repeated for decreased values for the control flux amplitude. According to figure 1 the frequency of the measured oscillations should decrease slightly at first and then significantly once the control flux amplitude is near $0.3\Phi_0$. We have observed the predicted change in frequency; however, only over a range of about 50 MHz when the oscillations completely disappeared (data not shown). This result could mean that coherence times for the tunable flux qubit itself could be very short because for lower values of control flux the wavefunctions resemble those of the flux qubit instead of the harmonic oscillator. We next measure a new device without the resonator attached in order to quantify performance of the bare qubit.

5.2. Bare qubits

The bare qubit is fabricated in the same fashion as the previously described IBM qubit except that the harmonic oscillator is left out. The sample is then mounted in a similar fashion and cooled down using the same coaxial configuration. Calibrations are performed as described earlier, as we expect to observe similar hysteresis patterns as in figure 5. Finally, we use these calibrations and repeat the pulsed control flux experiments. However, in this case, because the qubit is not coupled to the resonator we expect the qubit frequency to strongly depend on the control flux pulse amplitude instead of being fixed by the harmonic oscillator frequency.

The results of this experiment are shown in figure 7. As before, the phase of the oscillations changes 180°, depending on the polarity of the reset pulse. The frequency of the oscillations varies with control flux amplitude by an amount consistent with simulations. The results are similar to [30]. The data also shows a clear decay of these oscillations that

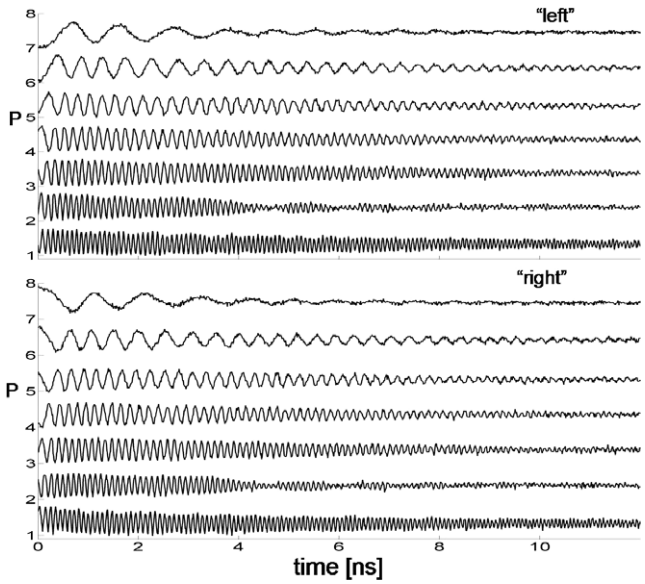

Figure 7. Bare qubit 'parking' experiment. The two panels display the probability of being in one well (e.g. the 'right' well) versus time delay for various pulsed control flux amplitudes (vertically offset for clarity). Each panel corresponds to either the 'left' or 'right' well initialization.

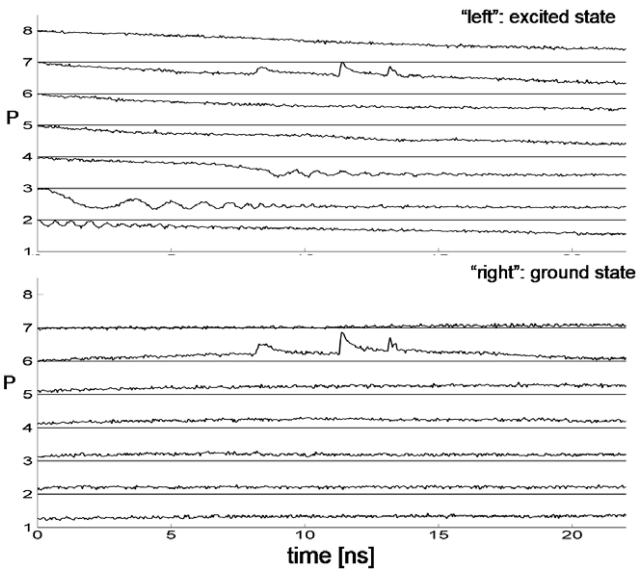

Figure 8. Bare qubit energy relaxation experiment. The two panels display the probability of being in one well (e.g. the 'right' well) versus time delay for various pulsed control flux amplitudes (vertically offset for clarity). Each panel corresponds to either the 'left' or 'right' well initialization corresponding to being in the excited and ground, respectively, when offsetting the flux bias by about $1m\Phi_0$. The horizontal lines are only a guide to the eye and do not have any physical meaning or relevance.

could be due to either energy relaxation (T_1) or dephasing (T_2) processes. Which process is dominant can be determined by performing the following additional experiment.

Offsetting the flux from Φ_0 by a small amount (of the order of 1 $m\Phi_0$) leads to a sufficiently asymmetric potential so that the frequency difference between the ground and excited state is greater than 1 GHz for all values of control flux (see [14]). Now the control flux pulse is no longer non-adiabatic but rather adiabatic! Therefore, by initializing the qubit into the excited state (which is simply the higher lying energy well for small values of control flux), followed by the adiabatic control flux pulse, the system remains in the exited state. By waiting for a certain duration and adiabatically moving the system back to its original starting point we measure the qubit. If the qubit did not decay to the ground state the measurement should reveal that the qubit is in the same well as it was prior to the control flux pulse, otherwise the qubit is in the other well location. This allows measurement of the relaxation time of the qubit from the excited to the ground state. The results of this experiment are shown in figure 8.

The top panel shows the data when initializing the qubit into the exited state. Similar to figure 7 each trace corresponds to a different control flux amplitude and therefore corresponds to a different frequency of the qubit. The data clearly shows that the qubit state decays rather quickly (∼10 ns) for all control flux amplitudes (or qubit frequencies), confirming that the observed oscillations in figure 7 are limited by energy relaxation.

However, the data also shows additional oscillations which are consistent with the presence of two-level systems (TLS) that couple to the qubit [31]. As we sweep the control flux the qubit energy may sweep through the energy state of a TLS. If the change in qubit energy is neither fully adiabatic nor non-adiabatic a superposition of the qubit and TLS states is created, similar to the parking experiment earlier and therefore oscillations are observed. Some of these oscillations do not necessarily begin at zero time, which implies that the actual control flux value is chirping. Upon careful inspection of the data in figure 7 we indeed observe the frequency of oscillations to change with time, confirming the presence of some chirping. The presence of the qubit–TLS oscillations complicates the analysis of the data but even for control flux amplitude for which no oscillations are visible the excited state nonetheless decays rapidly.

5.3. Coherence time discussion

The observed energy relaxation time of the bare qubit is significantly shorter than predicted by the current theoretical models. The cause of the reduced coherence times could be related to issues in the measurement set-up, materials or qubit design. We have implemented numerous modifications aimed at identifying the root cause.

5.3.1. Measurement. We do not believe that there are significant imperfections in the measurement set-up that could severely degrade measured relaxation times. Malfunction of the low temperature filters or substantial high frequency noise from room temperature electronics could, in principle, account for shortened coherence times. We have eliminated measurement-related issues by measuring a phase qubit from Dr Martinis' group using the same filtering and electronics. We obtained $T_1 \sim 400$ ns and narrow spectroscopic linewidths (1 MHz), confirming the group's measurement

results. Although a phase instead of a flux qubit was measured we believe they are similar enough that we can be confident that measurement-related issues are not a limiting factor.

5.3.2. Materials. Compared with phase qubits, materials have generally appeared to be much less of a limiting factor for the IBM qubit mostly because the junction area is so small that material defects have only a small chance of occurring [31]. For this reason we believe that the junction dielectric is not a limiting factor in our experiments. However, the substrate right underneath the qubit could well drastically impact coherence times. We have fabricated and tested bare qubits on various substrates:

- Low resistivity (few Ω cm) silicon wafers.
- Intrinsic silicon-on-sapphire wafers.
- Sapphire wafers.

Coherence times remained unaffected. Therefore we do not believe that the substrate plays a crucial role in our experiments.

Recently we have discovered through TEM imaging that the ion mill cleaning step prior to the aluminum deposition led to Iron and chrome incorporation into the silicon substrate. As a result these contaminants are present right underneath most of the qubit wiring and could well account for the reduced relaxation times. We are currently in the process of eliminating these contaminants. Although such contamination could well be the source of reduced relaxation times it is perplexing that sapphire substrates showed the same coherence times. One might naively expect the contamination to be less, or at least different, because sapphire is a much stiffer material. In either case, a clean experiment must be performed to confirm or eliminate contamination as a source of reduced relaxation times.

5.3.3. Design. Several design changes aimed at impacting coherence times have been implemented.

- Reduced the qubit wire width from close to 10 to 1 μm to ensure no flux trapping.
- Reduced mutual inductances to the SQUID, flux and control flux bias lines by as much as a factor of two to determine if magnetic coupling to bias lines or the measurement SQUID reduces coherence times.
- Substantially increased the thickness of the superconducting ground plane underneath the qubit to improve shielding of the qubit.
- Included a perforated ground plane around the qubit to improve the grounding of the bias lines and help further shield the qubit.
- Redesigned the qubit geometry to more closely resemble the shape of an hourglass to reduce the effects of capacitive coupling [25][2].

[2] Only one sample was tested and resonators fabricated in parallel with the measured qubit had much reduced quality factors for unknown reasons, possibly a dirty substrate or other one-time processing issues. It is possible that the qubit had actually been better but the improvement was masked by the same issues that the resonators were facing. More samples need to be tested to verify the null improvement.

However, none of the changes affected qubit coherence times. Despite these apparently negative results, we now believe that electromagnetic considerations for flux qubits, such as a capacitive coupling to bias leads, can be a significant source of decoherence. Much more attention should be devoted to this topic than there has been.

As described earlier and in [25] the T_1 coherence times can be conveniently expressed as an RC time constant. As explained in [25] capacitive coupling to bias leads gives an effective resistance $R = 1/Z_0(\omega C_c)^2$, where $Z_0 = 50\ \Omega$ is the characteristic impedance of the bias leads, ω is the qubit frequency and C_c is the capacitance to the bias lead. In another recent experiment described in [19] it is believed that capacitive coupling could indeed be the source for significantly reduced coherence times.

In order to obtain $T_1 = 1\ \mu$s we require $C_c < 0.5$ fF for a qubit frequency of 5 GHz and a junction capacitance of 10 fF. Although it has been argued in [25] that symmetry helps alleviate the effects of capacitive coupling it becomes clear that 0.5 fF worth of 'asymmetric' capacitance to bias leads is a dauntingly small number! Of particular concern is the feasibility of a scalable quantum computer based on 'floating' qubits.

A survey of flux qubit results shows that small and symmetric designs perform better than larger asymmetric ones, consistent with the capacitive coupling model. While it is generally straightforward to design a single flux qubit on its own with many symmetry features, designing a scalable unit cell flux qubit without breaking symmetry appears very challenging. In order to achieve appreciable qubit–qubit coupling the qubits should be rather large (and hence the reason for the large feature size of the IBM qubit). This implies a large capacitance to ground and, owing to the close proximity of neighboring qubits, there exists non-negligible capacitive coupling to these qubits, possibly leading to unwanted coupling or, worse, additional decoherence mechanisms. To what extent capacitive coupling plays a role in a scalable system is not clear. However, it is clear that, because of the logarithmic scaling of capacitance, a large scale quantum computer based on flux qubits certainly has to take these effects into account! To the best of our knowledge this has not yet been done.

To this date it is not clear what exactly is causing the reduced coherence times of our qubits. The experiment that made the qubit geometrically more symmetrical to reduce the effects of capacitive coupling showed no effect on relaxation times, possibly indicating our qubit may not be limited by capacitive coupling. However, at the same time we have not modeled the entire qubit using an electromagnetic simulator either. A detailed model that incorporates junction dynamics does not yet exist and, even when linearizing the response of the junctions, many simulation tools are not optimized to reliably predict accurate quality factors of resonators. It may well be possible that we are currently limited by some electromagnetic effect that we have not yet incorporated into our modeling. In order to discover the root source of reduced relaxation times we decided to fabricate flux qubits based on a successfully demonstrated design that can be adapted step by step into the IBM qubit design.

5.4. 3JJ flux qubits

A flux qubit design based on three Josephson junctions inserted into a superconducting loop, and surrounded by the measurement SQUID [15], has been used to demonstrate long T_1 and T_2 times at the symmetry point. We fabricated qubits based on this design using at first *only* the lift-off process described earlier. Unlike the experiments described in [15] no on-chip resistors for the SQUID were fabricated. This reduces the fabrication resources to a minimum. Because the ion mill clean step has been identified as a possible source of contamination this step too has been eliminated[3].

We have only preliminary results and thus do not show any data here nor do we describe the experimental procedure due to space considerations. However, the experiments repeat the procedure from [15] closely. Similar to [15] we have obtained qubit spectroscopy with a gap frequency of 2.8 GHz. Although the spectroscopic linewidths were rather large (∼100 MHz) even at the symmetry point for yet unknown reasons (implying a short T_2), the measured energy relaxation time $T_1 = 1\ \mu$s is much longer than those measured in an IBM qubit! At least as far as energy relaxation is concerned we now have a proof-of-principle of long energy relaxation times.

A second sample was fabricated in a more complex process by adding some fabrication, complexity permitting the SQUID resistors to be on-chip. This step helps protect the SQUIDs from electrostatic discharge to minimize device failure. The ion mill step is again included but has been modified so that a tantalum sleeve is now used leading to tantalum contamination which is still not optimal but is magnetically not as serious as Iron. This qubit (and another measured nominally identical qubit) had a smaller gap frequency of 1 GHz (<0.5 GHz for the other qubit). The energy relaxation was measured to be $T_1 \sim$ 100–150 ns (not measurable for the other qubit) and the spectroscopy was much narrower at the gap. Due to the low signal-to-noise ratio we were not able to measure the low power linewidth. Rabi oscillations with a decay envelope of about 200 ns were measured (but not measurable for the other qubit).

Finally, a third sample was fabricated with a slightly decreased junction area ratio to increase the qubit frequency at the symmetric point. The processing was identical to the second sample but the ion mill step was left out (the on-chip SQUID resistors are still present, unlike the first sample). The gap frequency was 9.65 GHz, the T_1 time was measured to be 400 ns after optimizing the dc SQUID current [32, 33] and the Rabi decay times were approximately 250 ns although it was not optimized.

6. Conclusions

In summary we have shown a tunable flux qubit design exhibiting features which we believe are critical in a scalable architecture. High fidelities and long coherence times are predicted by current theoretical models. First experimental results successfully demonstrate parking of quantum information with long coherence times. However, qubit coherence times are much shorter than predicted. A new theoretical model that includes capacitive effects reveals loss mechanisms that should be taken seriously. Because of the small self-capacitance of the flux qubit a large impedance of the environment is even more crucial. Finally, recent tests using 3JJ flux qubits stripped of various scalability features were measured with reasonably long coherence times. By adding these features in one-by-one we hope to identify the step at which coherence times are drastically affected in the IBM qubit. Some or even most of the degraded coherence times could be related to the contamination discussed earlier, and experiments will soon tell the whole story.

Our experiments clearly show the richness and complexity of the field of superconducting qubits. Loss mechanisms related to issues ranging anywhere from materials over hardware to electromagnetic effects can unexpectedly degrade qubit performance. It is therefore important to consider all possible loss mechanisms and objectively revisit them periodically. We believe that decoherence currently poses the greatest challenge towards a scalable design and should be focused on.

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressly or implied, of the US Government.

[3] The ion mill clean step is required for the IBM qubit because a good contact between the niobium and shadow evaporated aluminum is required. Leaving out this step leads to a resistive contact that renders the measurement SQUID useless.

Acknowledgments

This work is dedicated to Roger H Koch who initiated and passionately led the experimental quantum computing effort at IBM until his tragic passing.

References

[1] Bertet P, Chiorescu I, Burkard G, Semba K, Harmans C J P M, DiVincenzo D P and Mooij J E 2005 *Phys. Rev. Lett.* **95** 257002
[2] Siddiqi I, Vijay R, Metcalfe M, Boaknin E, Frunzio L, Schoelkopf R J and Devoret M H 2006 *Phys. Rev.* B **73** 054510
[3] Wallraff A, Schuster D I, Blais A, Frunzio L, Majer J, Devoret M H, Girvin S M and Schoelkopf R J 2005 *Phys. Rev. Lett.* **95** 060501
[4] Steffen M, Ansmann M, McDermott R, Katz N, Bialczak R C, Lucero E, Neeley M, Weig E M, Cleland A N and Martinis J M 2006 *Phys. Rev. Lett.* **97** 050502
[5] Yamamoto T, Pashkin Y A, Astafiev O, Nakamura Y and Tsai J 2003 *Nature* **425** 941
[6] Berkley A, Xu H, Ramos R, Gubrud M, Strauch F, Johnson P, Anderson J, Dragt A, Lobb C and Wellstood F 2003 *Science* **300** 1548
[7] McDermott R, Simmonds R, Steffen M, Cooper K, Cicak K, Osborn K, Oh S, Pappas D and Martinis J 2005 *Science* **307** 1299
[8] Hime T, Reichardt P, Plourde B, Robertson T, Wu C-E, Ustonov A and Clarke J 2006 *Science* **314** 1427
[9] Steffen M, Ansmann M, Bialczak R, Katz N, Lucero E, McDermott R, Neeley M, Weig E, Cleland A and Martinis J 2006 *Science* **313** 1423

[10] DiCarlo L, Chow J, Bambetta J, Bishop L S, Schuster D, Majer J, Blais A, Frunzio L, Girvin S and Schoelkopf R 2009 *Nature* **460** 240
[11] Koch R H, Keefe G A, Milliken F P, Rozen J R, Tsuei C C, Kirtley J R and DiVincenzo D P 2006 *Phys. Rev. Lett.* **96** 127001
[12] Gao J, Zmuidzinas J, Mazin B, LeDuc H and Day P 2007 *Appl. Phys. Lett.* **90** 102507
[13] Wang H *et al* 2009 arXiv:/pdf/0909.0547
[14] Brito F, DiVincenzo D and Steffen M 2008 *New J. Phys.* **10** 033027
[15] Plourde B, Robertson T, Reichardt P, Hime T, Linzen S, Wu C and Clarke J 2005 *Phys. Rev.* B **73** 060506
[16] Yoshihara F, Harrabi K, Niskanen A, Nakamura Y and Tsai J 2006 *Phys. Rev. Lett.* **97** 167001
[17] Aliferis P, Brito F, DiVincenzo D P, Preskill J, Steffen M and Terhal B 2009 *New J. Phys.* **11** 013061
[18] DiVincenzo D P, Brito F and Koch R H 2006 *Phys. Rev.* B **74** 014514
[19] Paauw F, Fedorov A, Harmans C and Mooij J 2008 arXiv:0812.1912 [cond-mat]
[20] Koch R H, Rozen J R, Keefe G A, Milliken F M, Tsuei C C, Kirtley J R and DiVincenzo D P 2005 *Phys. Rev.* B **72** 092512
[21] Shahriar M S, Pradhan P and Morzinski J 2004 *Phys. Rev.* A **69** 032308
[22] Khaneja N, Reiss T, Kehlet C, Schulte-Herbrüggen T and Glaser S 2005 *J. Magn. Reson.* **172** 296
[23] Steffen M and Koch R H 2007 *Phys. Rev.* B **75** 062326
[24] Burkard G, Koch R and DiVincenzo D 2004 *Phys. Rev.* B **69** 064503
[25] Steffen M, Brito F, DiVincenzo D, Kumar S and Ketchen M 2009 *New J. Phys.* **11** 033030
[26] van der Wal C, Wilhelm F, Harmans C and Mooij J 2003 *Eur. Phys. J.* B **31** 111
[27] Neeley M, Ansmann M, Bialczak R, Hofheinz M, Katz N, Lucero E, O'Connell A, Wang H, Clelang A and Martinis J 2008 *Phys. Rev.* B **77** 180508
[28] Chiorescu I, Nakamura Y, Harmans C and Mooij J 2003 *Science* **299** 1869
[29] Nakamura Y, Pashkin Y and Tsai J 1999 *Nature* **3398** 786
[30] Poletto S, Chiarello F, Castellano M, Lisenfeld J, Lukashenko A, Cosmelli C, Torrioli G, Carelli P and Ustinov A 2009 *New J. Phys.* **11** 013009
[31] Martinis J M *et al* 2005 *Phys. Rev. Lett.* **95** 210503
[32] Burkard G, DiVincenzo D, Bertet P, Chiorescu I and Mooij J 2005 *Phys. Rev.* B **71** 134504
[33] Bertet P, Harmans C J P M and Mooij J E 2006 *Phys. Rev.* B **73** 064512