

Realization of three-qubit quantum error correction with superconducting circuits

M. D. Reed¹, L. DiCarlo², S. E. Nigg¹, L. Sun¹, L. Frunzio¹, S. M. Girvin¹ & R. J. Schoelkopf¹

Quantum computers could be used to solve certain problems exponentially faster than classical computers, but are challenging to build because of their increased susceptibility to errors. However, it is possible to detect and correct errors without destroying coherence, by using quantum error correcting codes¹. The simplest of these are three-quantum-bit (three-qubit) codes, which map a one-qubit state to an entangled three-qubit state; they can correct any single phase-flip or bit-flip error on one of the three qubits, depending on the code used². Here we demonstrate such phase- and bit-flip error correcting codes in a superconducting circuit. We encode a quantum state^{3,4}, induce errors on the qubits and decode the error syndrome—a quantum state indicating which error has occurred—by reversing the encoding process. This syndrome is then used as the input to a three-qubit gate that corrects the primary qubit if it was flipped. As the code can recover from a single error on any qubit, the fidelity of this process should decrease only quadratically with error probability. We implement the correcting three-qubit gate (known as a conditional-conditional NOT, or Toffoli, gate) in 63 nanoseconds, using an interaction with the third excited state of a single qubit. We find 85 ± 1 per cent fidelity to the expected classical action of this gate, and 78 ± 1 per cent fidelity to the ideal quantum process matrix. Using this gate, we perform a single pass of both quantum bit- and phase-flip error correction and demonstrate the predicted first-order insensitivity to errors. Concatenation of these two codes in a nine-qubit device would correct arbitrary single-qubit errors. In combination with recent advances in superconducting qubit coherence times^{5,6}, this could lead to scalable quantum technology.

Quantum error correction relies on detecting the presence of errors without gaining knowledge of the encoded quantum state. In the three-qubit error-correcting code, the subspace of the two additional ‘ancilla’ qubits uniquely encodes which of the four possible single-qubit errors has occurred, including the possibility of no flip. Crucially, errors consisting of finite rotations can also be corrected using these schemes because the error syndromes are allowed to be in superpositions of the possible outcomes, flipped and not flipped². Previous works implementing error correcting codes in liquid-^{7–9} and solid-state¹⁰ NMR and with trapped ions^{11,12} have demonstrated two possible strategies for using the error syndromes. The first is to measure the ancillas (thereby projecting the syndrome) and use a classical logic operation to correct the detected error. This ‘feed-forward’ capability is challenging in superconducting circuits as it requires a fast and high-fidelity quantum non-demolition measurement, but is probably a necessary component to achieve scalable fault tolerance^{2,13}. The second strategy, as recently demonstrated with trapped ions¹² and used here, is to replace the classical logic with a quantum controlled-controlled NOT (CCNOT) gate that performs the correction coherently, leaving the entropy associated with the error in the ancilla qubits, which can then be reset¹⁴ if the code is to be repeated. The CCNOT gate performs exactly the action that would follow the measurement in the first scheme: flipping the primary qubit if and only if the ancillas encode the associated error syndrome.

The CCNOT gate is also vital for a wide variety of applications such as Shor’s factoring algorithm¹⁵ and has attracted substantial experimental interest, with recent implementations in linear optics¹⁶, trapped ions¹⁷ and superconducting circuits^{18,19}. Here we use the circuit quantum electrodynamics architecture²⁰ to couple four superconducting transmon qubits²¹ to a single microwave cavity bus²², where each qubit transition frequency can be controlled on nanosecond timescales with individual flux bias lines²³ and collectively measured by interrogating transmission through the cavity²⁴. (The details of the device can be found in Methods Summary and in ref. 3.) The frequencies of the qubits, labelled Q₁–Q₄, are tuned respectively to 6, 7, 7.85 and ~ 13 GHz, with Q₄ unused. In this Letter, we first demonstrate the three-qubit interaction used in the gate, which is an extension of interactions used in previous two-qubit gates^{3,23,25}, and show how this interaction can be used to create the desired CCNOT gate. We then verify its action and use it to demonstrate error correction for an error on a single qubit with the bit-flip code and then for simultaneous errors on all three qubits with the phase-flip code. We find a quadratic dependence of process fidelity on error probability, indicating that the algorithm is correcting errors as predicted.

Our three-qubit gate uses an interaction with the third excited state of one transmon. Specifically, it relies on the unique capability among computational states (eigenstates of the Pauli operator Z) of $|111\rangle$ to interact with the non-computational state $|003\rangle$ (the notation $|abc\rangle$ refers to the excitation levels of Q₁–Q₃, respectively). As the direct interaction of these states is prohibited to first order, we first transfer the quantum amplitude of $|111\rangle$ to the intermediate state $|102\rangle$, which itself couples strongly to $|003\rangle$. Calculated energy levels and time-domain data showing interaction between $|011\rangle$ and $|002\rangle$ (which is identical to that between $|111\rangle$ and $|102\rangle$ except for a 6-GHz offset) as a function of the flux bias on Q₂ are shown in Fig. 1a. Once the amplitude of $|111\rangle$ has been transferred to $|102\rangle$ with a sudden swap interaction, a three-qubit phase is acquired by moving Q₁ up in frequency adiabatically, near the avoided crossing with $|003\rangle$. Figure 1b shows the avoided crossing between these states as a function of the flux bias on Q₁. This crossing shifts the frequency of $|102\rangle$ relative to the sum of the frequencies of $|100\rangle$ and $|002\rangle$ to yield the three-qubit phase. The detailed procedure of the gate is shown in Fig. 2a, and is implemented in 63 ns. Further details can be found in Supplementary Information.

We demonstrate the gate by first measuring its classical action. The controlled-controlled phase (CCPhase) gate, which maps $|111\rangle$ to $-|111\rangle$, has no effect on pure computational states so we implement a CCNOT gate by concatenating pre- and post-gate rotations on Q₂, as indicated in the unshaded regions of Fig. 2a. Such a gate ideally swaps $|101\rangle$ and $|111\rangle$ and does nothing to the remaining states. To verify this, we prepare the eight computational states, implement the gate and measure its output using three-qubit state tomography³ to generate the classical truth table. The intended state is reached with $85 \pm 1\%$ fidelity on average. This measurement is sensitive only to classical action, however, so we complete our verification by performing full quantum process tomography on the CCPhase gate, which can detect

¹Departments of Physics and Applied Physics, Yale University, New Haven, Connecticut 06520, USA. ²Kavli Institute of Nanoscience, Delft University of Technology, 2628 CJ Delft, The Netherlands.

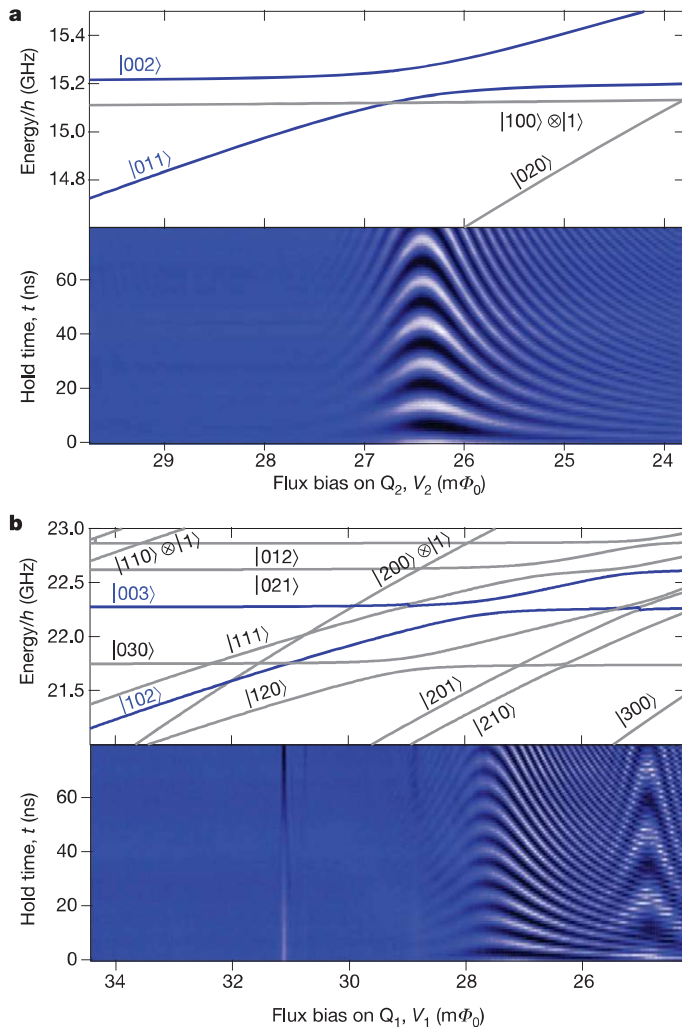


Figure 1 | Calculated energy spectra and time-domain measurements of the interactions used in the three-qubit gate. **a**, The energy spectrum of doubly excited states demonstrating the avoided crossing between $|011\rangle$ and $|002\rangle$ (identical to that between $|111\rangle$ and $|102\rangle$ except for a 6-GHz offset) is shown with both a numerical diagonalization of the system Hamiltonian (top) and a time-domain measurement as a function of the applied magnetic flux on Q_2 (bottom). Top: the frequencies for the involved eigenstates are blue and the non-interacting eigenstates of similar energy are grey. The notation $|abc\rangle \otimes |d\rangle$ indicates the excitation level of each qubit and the cavity photon number, respectively. When the second ket is omitted, $d = 0$. Bottom: the state $|011\rangle$ is prepared and a square flux pulse of duration t and amplitude V_2 is applied. Coherent oscillations produce a ‘chevron’ pattern, with darker colours corresponding to population left in $|002\rangle$. h , Planck’s constant. **b**, The spectrum of triply excited states showing the avoided crossing between $|102\rangle$ and $|003\rangle$ as a function of the flux bias on Q_1 is characterized in the same way as above. The state $|102\rangle$ is prepared by first making $|111\rangle$ and then performing the swap as described in Fig. 2. Many additional eigenstates are close in energy but are irrelevant because they do not interact with the populated states. The large avoided crossing between the relevant eigenstates that is used to produce an adiabatic three-qubit interaction happens near $28 m\Phi_0$ (where Φ_0 is the magnetic flux quantum). Extra lines near $31 m\Phi_0$ and $29 m\Phi_0$ are due to higher-order interactions predicted by the Hamiltonian ($|102\rangle$ with $|030\rangle$ and $|003\rangle$ with $|111\rangle$), as is the larger first-order interaction at $25 m\Phi_0$ ($|102\rangle$ with a hybridization of $|021\rangle$ and $|111\rangle$), but their effect on the protocol in Fig. 2 is negligible.

the evolution of quantum superpositions of computational states. This is done by preparing 64 input states that span the computational Hilbert space and by performing state tomography on the result of the gate’s action on each state. As detailed in Supplementary Information, we find a fidelity of $78 \pm 1\%$ to a process in which the spurious two-qubit phase between Q_1 and Q_3 is set to the independently

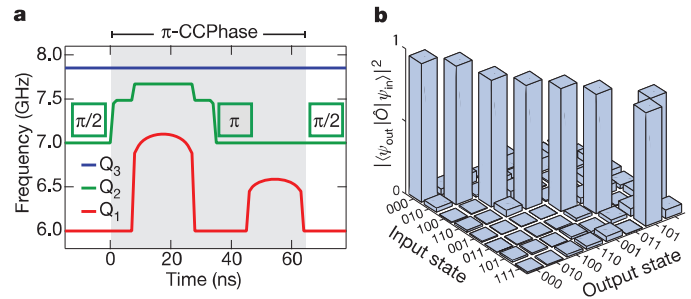


Figure 2 | Pulse sequence and classical action of the three-qubit gate. **a**, The frequencies of the three qubits and the locations of applied rotations during the three-qubit gate as functions of time. Shaded region: to produce the CCPhase interaction, Q_2 is first moved suddenly into resonance with the avoided crossing shown in Fig. 1a, which coherently transfers the population of $|111\rangle$ to $|102\rangle$ (and also that of $|011\rangle$ to $|002\rangle$) in 7 ns. Fine adjustments in the first point of the pulse compensate for finite pulse rise time and temporal precision. The frequency of Q_2 is then abruptly increased to where its two-qubit phase with Q_3 is cancelled during the gate by accumulating a multiple of 2π . The frequency of Q_1 is then increased adiabatically to initiate the interaction between $|102\rangle$ and $|003\rangle$. The duration and amplitude of this excursion is tuned to acquire a three-qubit phase of π . The population in $|102\rangle$ is then transferred back to $|111\rangle$ by reversing the swap procedure. Finally, the two-qubit phase between Q_1 and Q_2 is cancelled with an additional adiabatic interaction, which is sped up with a π -pulse on Q_2 at 37 ns (all rotations here are done about the x axis). The two-qubit phase between Q_1 and Q_3 is uncontrolled and there is an overall π -rotation of Q_2 , making this a π -CC- $e^{i\phi}Z$ gate, taking a total of 63 ns. Unshaded region: pre- and post-gate rotations on Q_2 appended to the CCPhase gate turn its action into that of a CCNOT gate, as described in Supplementary Information. **b**, The classical action of the CCNOT gate is measured by preparing the eight computational basis states, $|\psi_{in}\rangle$, and performing state tomography on the resulting state, $|\psi_{out}\rangle$, after applying the gate, \hat{O} , to them. The projection of these measurements to the computational basis states is taken to generate the displayed truth table. The fidelity to the expected action, where only the states $|101\rangle$ and $|111\rangle$ are swapped, is $85 \pm 1\%$. Full quantum process tomography of the gate is shown in Supplementary Information.

measured value of 57° (see Supplementary Information for an explanation of why this phase is irrelevant here). Owing to this extraneous phase, ϕ , the gate is most accurately described as a CC- $e^{i\phi}Z$ gate. The loss of fidelity is consistent with the expected energy relaxation of the three qubits during the 85-ns tomography procedure, which includes preparation and analysis pulses in addition to the gate, with some remaining error due to qubit transition frequency drift during the 90 min it takes to collect the full data set.

With our CCPhase gate in hand, we now demonstrate three-qubit error correction. We first examine the bit-flip code, which, as shown in Fig. 3a, starts by encoding the quantum state to be protected in a three-qubit entangled state² through the use of conditional phase (CPhase) gates. The state $\alpha|0\rangle + \beta|1\rangle$ is encoded as $\alpha|000\rangle + \beta|111\rangle$, which has the property that the value of any two-qubit ZZ product is +1 regardless of the values of α and β . (For quantum states on the equator of the Bloch sphere, $|\alpha| = |\beta| = 1/\sqrt{2}$ and the encoding is a maximally entangled three-qubit Greenberger–Horne–Zeilinger state^{3,4,26} that we independently measure to have a state fidelity of $89 \pm 1\%$.) If any single qubit is flipped, one or more of the ZZ products will flip sign as well. For example, if Q_1 were flipped, the Z_1Z_2 product would become -1 whereas the Z_2Z_3 product would remain $+1$, uniquely indicating that Q_1 needs to be corrected. Indeed, the four possible combinations of Z_1Z_2 and Z_2Z_3 exactly encode the possible single bit flips, including the possibility of no flip. In a fault-tolerant code, these products would be stored in separate qubits for later measurement², but here we instead reverse the encoding so that the ancillas Q_1 and Q_3 can no longer witness bit-flip errors and instead store the values of the two ZZ products. These ancillas are then used as the control bits for the CCNOT gate described above, so that Q_2 will be flipped back if and only if both ancillas are excited, which indicates that Q_2 was flipped. The detailed

evolution of the qubits during the error correction procedure can be found in Supplementary Information.

Whereas errors on classical bits are discrete, quantum error correction must be able to correct arbitrary rotations as well as complete flips because superpositions of states are allowed. Remarkably, the code described above already satisfies this criterion. If an error causes a rotation θ on Q_2 , the quantum state after decoding will be $\sqrt{1-p}(\alpha|0\rangle + \beta|1\rangle) \otimes |00\rangle + \sqrt{p}(\beta|0\rangle + \alpha|1\rangle) \otimes |11\rangle$, where $p = \sin^2(\theta/2)$ is the effective probability of a full flip and where we have listed first the state of Q_2 followed by those of Q_1 and Q_3 for notational simplicity. That is, the state will be a superposition of Q_2 in the correct state with the ancillas indicating no error plus Q_2 flipped with the ancillas indicating as such. The application of the CCNOT gate to this state will successfully correct it because it acts only on the subspace where both ancillas are excited, making the state $(\alpha|0\rangle + \beta|1\rangle) \otimes (\sqrt{1-p}|00\rangle + \sqrt{p}|11\rangle)$. (Alternatively, if the ancilla

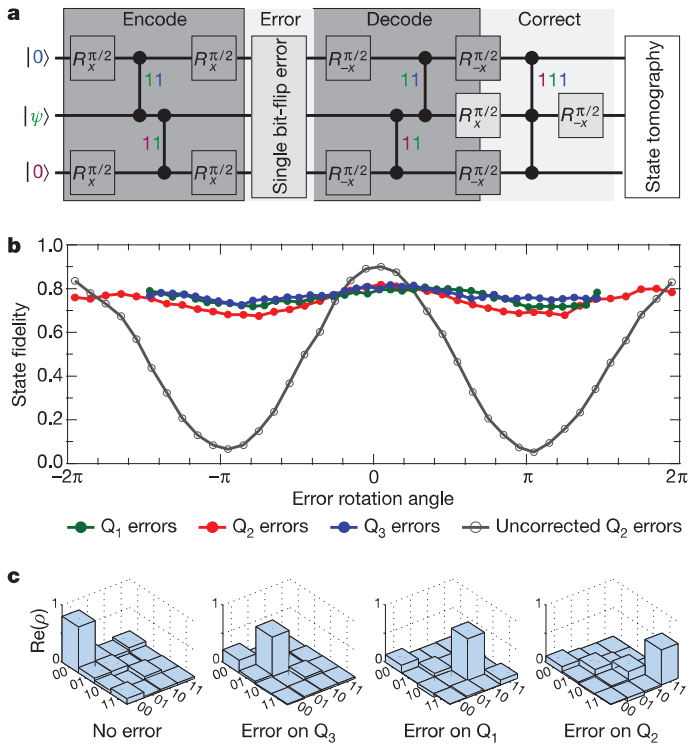


Figure 3 | Bit-flip error correction demonstrating recovery from a single arbitrary rotation. **a**, The error correction protocol starts by encoding the quantum state to be protected in a three-qubit state by entangling the two ancilla qubits, Q_1 and Q_3 , with Q_2 through the use of two CPhase gates (vertical lines terminating in solid circles) and $\pi/2$ -rotations ($R_{\hat{n}}^{\alpha}$ is a single-qubit rotation, where \hat{n} indicates the rotation axis and α is the rotation angle). The number adjacent to each CPhase gate indicates which state receives a phase shift of π . A single y -rotation error of a known angle is then performed on a single qubit (as explained in Supplementary Information, this is compiled together with other rotations when acting on the ancillas). The state is then decoded, leaving the ancillas in a product state indicating which single-qubit error occurred. For finite rotations, the ancillas will be in a superposition of states in which the error did and, respectively, did not occur, with each tensor multiplied by the associated single-qubit state of Q_2 . If an error occurred on Q_2 , the CCNOT gate implemented with our CCPhase gate (represented by three solid circles linked by a vertical line) at the end of the code will correct it. We then perform three-qubit state tomography to verify the result. **b**, State fidelity to the created state $|\psi\rangle = |+\rangle_X$ after causing an error on one of the qubits, with and without error correction. Ideally, the error-corrected curves would be horizontal lines at unit fidelity. Finite excited-state lifetimes cause oscillations and lower fidelity because errors change the excitation level of the system. **c**, Two-qubit density matrices (ρ) of the ancillas after each of the four possible full bit-flip errors has occurred. The fidelities of these states to the ideal error syndromes, $|00\rangle$, $|01\rangle$, $|10\rangle$ and $|11\rangle$, are respectively 81.3%, 69.7%, 73.1% and 61.2%.

qubits were measured, they would project the state onto one of those possibilities, essentially forcing the computer to ‘decide’ whether a full flip had occurred.) We demonstrate this with the procedure shown in Fig. 3a for the state $|\psi\rangle = |+\rangle_X$ (the positive eigenstate of the Pauli operator X), performing single deterministic rotations of a known angle on each of the three qubits to simulate errors. As shown in Fig. 3b, we compare this with the case of uncorrected errors on Q_2 . Ideally, the error-corrected curves would have unit fidelity and be independent of θ , but they are slightly lower in fidelity and oscillate in θ owing to finite coherence. They are, however, substantially improved relative to the uncorrected case, demonstrating that the errors are in fact being ameliorated. As shown in Fig. 3c, we also measure the two-qubit density matrix of the ancilla qubits after each of the four possible full bit-flip errors, showing that they end up in a computational product state correctly indicating the induced error.

In real physical systems, errors occur at approximately the same rate on all constituent qubits rather than on one at a time. The correction

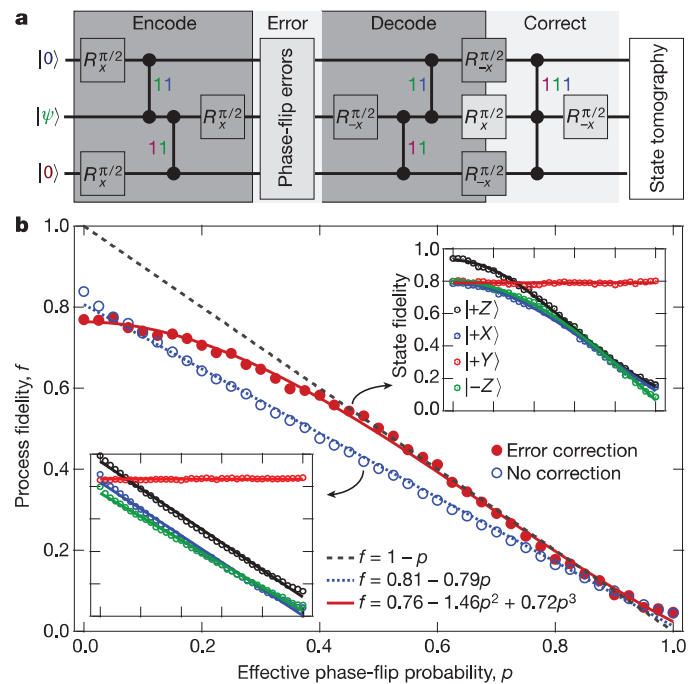


Figure 4 | Demonstration of first-order insensitivity to simultaneous phase-flip errors. **a**, The phase-flip error correction protocol differs from the bit-flip protocol described in Fig. 3a only by single-qubit gates. Those gates effectively rotate the coordinate system, mapping phase flips to bit flips, and vice versa, so the remainder of the procedure is exactly the same as in the bit-flip case². We perform errors on all three qubits simultaneously with z gates of known rotation angle, which is equivalent to phase-flip errors with probability $p = \sin^2(\theta/2)$. As with the bit-flip code, if a single error has occurred on the primary qubit, the CCNOT gate at the end of the code will undo it. **b**, Fidelity of the process matrix of the protected qubit to the identity operation plotted as a function of p . As the code corrects only single errors, it will fail on the three-qubit subspace where more than one has occurred, which happens with a probability $3p^2 - 2p^3$. The coefficients here are reduced for processes with finite fidelity. The process fidelity is fitted with $f = (0.760 \pm 0.005) - (1.46 \pm 0.03)p^2 + (0.72 \pm 0.03)p^3$. If a linear term is allowed, its best-fit coefficient is 0.03 ± 0.06 . We compare this with the case of no error correction to simulate the improvement made when the decoherence of Q_2 is normalized away (blue symbols). We also plot the simulated fidelity of a perfect but non-corrected system (black dashed line), which indicates that for our gate fidelities we do not show a net improvement for artificial errors. Insets: the constituent state fidelities of the four basis states used to produce the process fidelity data in the case with error correction (right) and in the case with no correction (left). The x axes of the plots are the same as the main panel, and they share the same y axis. The state $|+\rangle_Y$ (the positive eigenstate of the Pauli operator Y) is immune to errors because its encoded state is an eigenstate of single, double and triple qubit phase flips.

scheme will only succeed, therefore, on the three-qubit subspace with zero or one errors. The probability of more than one error occurring is $3p^2 - 2p^3$, where p is the single-qubit error rate², so the fidelity of error correction should be $1 - 3p^2 + 2p^3$. For a scheme with gate fidelity limited by decoherence, the coefficients of the quadratic and cubic terms will be smaller but, crucially, any linear dependence on p will be strongly suppressed. If the error rates for each qubit were different, these coefficients would again be modified but any linear dependence would still be abated. For the sake of completeness, here we use the phase-flip code, which differs from the previously discussed bit-flip code by only single-qubit rotations, as shown in Fig. 4a. This difference can be viewed as a rotation of the coordinate system, converting phase flips to bit flips and vice versa, so the remainder of the code is exactly the same as the previous case^{2,12,27}. Phase errors of known rotation angle are applied by rotating the frame of reference of subsequent x and y rotations. As shown in Fig. 4b, we measure the process fidelity of the error correction scheme as a function of p and compare this with the case of no error correction in which identical single-qubit rotations are applied to Q_2 but the ancillas are not involved (this comparison is without gates, but with appropriate delays to have the same total procedure duration, to indicate the lack of fidelity due to the decoherence of Q_2). Whereas without error correction we find a purely linear dependence on p , with the correction applied the data are extremely well modelled by only quadratic and cubic terms, demonstrating the desired first-order insensitivity to errors. We have therefore realized a successful implementation of quantum error correction, although improvement of the fidelity of a real physical process will require considerable advances in both gate fidelity and device complexity.

We have realized both bit- and phase-flip error correction in a superconducting circuit. In doing so, we have tested both main conceptual components of the nine-qubit Shor code¹, which can defend against arbitrary single-qubit errors by concatenating the bit- and phase-flip codes. The implementation relies on our efficient three-qubit gate, which uses non-computational states in the third excitation manifold of our system, demonstrating that the simple Hamiltonian of the system accurately predicts the dynamics even at these high excitation levels. The gate takes approximately half the time of an equivalent construction with one- and two-qubit gates. We expect it to work between any three nearest-neighbour qubits in frequency regardless of the number of qubits sharing the bus, as interactions involving other qubits will be first-order prohibited.

METHODS

Arbitrary qubit rotations around the x and y axes of the Bloch sphere are performed with pulse-shaped resonant microwave tones. Rotations around the z axis are made by rotating the reference phase of subsequent x and y pulses. One-qubit dynamical phases resulting from flux excursions are measured with modified Ramsey experiments comparing the phase acquired by an unmodified prepared state with the phase acquired by that same state after a flux pulse, and are cancelled with z rotations. Two- and three-qubit phases are measured with a similar Ramsey experiment comparing the phase acquired when a control qubit is in its ground state with the phase acquired when it is in an excited state. For example, the two-qubit phase between Q_2 and Q_3 is measured by preparing Q_3 along the y axis and Q_2 either in its ground or excited state and then performing the flux pulse in both cases. The single-qubit phase of Q_3 is the same for both states, so the two-qubit phase is directly measurable as the difference in phase between them. All phases are initially tuned to within 1° , limited by the resolution of control equipment and drifts of system parameters such as the qubit transition frequencies.

Received 21 September; accepted 7 December 2011.

Published online 1 February 2012.

1. Shor, P. W. Scheme for reducing decoherence in quantum computer memory. *Phys. Rev. A* **52**, R2493–R2496 (1995).

2. Nielsen, M. A. & Chuang, I. L. *Quantum Computation and Quantum Information* (Cambridge Ser. Inf. Nat. Sci, Cambridge Univ. Press, 2000).
3. DiCarlo, L. *et al.* Preparation and measurement of three-qubit entanglement in a superconducting circuit. *Nature* **467**, 574–578 (2010).
4. Neeley, M. *et al.* Generation of three-qubit entangled states using superconducting phase qubits. *Nature* **467**, 570–573 (2010).
5. Paik, H. *et al.* Observation of high coherence in Josephson junction qubits measured in a three-dimensional circuit QED architecture. *Phys. Rev. Lett.* **107**, 240501 (2011).
6. Kim, Z. *et al.* Decoupling a Cooper-pair box to enhance the lifetime to 0.2 ms. *Phys. Rev. Lett.* **106**, 120501 (2011).
7. Cory, D. *et al.* Experimental quantum error correction. *Phys. Rev. Lett.* **81**, 2152–2155 (1998).
8. Knill, E., Laflamme, R., Martinez, R. & Negrevergne, C. Benchmarking quantum computers: the five-qubit error correcting code. *Phys. Rev. Lett.* **86**, 5811–5814 (2001).
9. Boulant, N., Viola, L., Fortunato, E. & Cory, D. Experimental implementation of a concatenated quantum error-correcting code. *Phys. Rev. Lett.* **94**, 130501 (2005).
10. Moussa, O., Baugh, J., Ryan, C. A. & Laflamme, R. Demonstration of sufficient control for two rounds of quantum error correction in a solid state ensemble quantum information processor. *Phys. Rev. Lett.* **107**, 160501 (2011).
11. Chiaverini, J. *et al.* Realization of quantum error correction. *Nature* **432**, 602–605 (2004).
12. Schindler, P. *et al.* Experimental repetitive quantum error correction. *Science* **332**, 1059–1061 (2011).
13. Shor, P. W. in *Proc. 37th Symp. Foundations Comput.* 56–65 (IEEE, 1996); preprint at (<http://arxiv.org/abs/quant-ph/9605011>) (1996).
14. Reed, M. D. *et al.* Fast reset and suppressing spontaneous emission of a superconducting qubit. *Appl. Phys. Lett.* **96**, 203110 (2010).
15. Shor, P. W. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Sci. Statist. Comput.* **26**, 1484–1509 (1997).
16. Lanyon, B. P. *et al.* Simplifying quantum logic using higher-dimensional Hilbert spaces. *Nature Phys.* **5**, 134–140 (2009).
17. Monz, T. *et al.* Realization of the quantum Toffoli gate with trapped ions. *Phys. Rev. Lett.* **102**, 040501 (2009).
18. Mariantoni, M. *et al.* Implementing the quantum von Neumann architecture with superconducting circuits. *Science* **334**, 61–65 (2011).
19. Fedorov, A., Steffen, L., Baur, M. & Wallraff, A. Implementation of a Toffoli gate with superconducting circuits. *Nature* **481**, 170–172 (2012).
20. Wallraff, A. *et al.* Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics. *Nature* **431**, 162–167 (2004).
21. Schreier, J. A. *et al.* Suppressing charge noise decoherence in superconducting charge qubits. *Phys. Rev. B* **77**, 180502 (2008).
22. Majer, J. *et al.* Coupling superconducting qubits via a cavity bus. *Nature* **449**, 443–447 (2007).
23. DiCarlo, L. *et al.* Demonstration of two-qubit algorithms with a superconducting quantum processor. *Nature* **460**, 240–244 (2009).
24. Reed, M. *et al.* High-fidelity readout in circuit quantum electrodynamics using the Jaynes-Cummings nonlinearity. *Phys. Rev. Lett.* **105**, 173601 (2010).
25. Strauch, F. *et al.* Quantum logic gates for coupled superconducting phase qubits. *Phys. Rev. Lett.* **91**, 167005 (2003).
26. Greenberger, D. M., Horne, M. A. & Zeilinger, A. in *Bell's Theorem, Quantum Theory and Conceptions of the Universe* (ed. Kafatos, M.) 69–72 (Kluwer Academic, 1989).
27. Tornberg, L., Wallquist, M., Johansson, G., Shumeiko, V. & Wendin, G. Implementation of the three-qubit phase-flip error correction code with superconducting qubits. *Phys. Rev. B* **77**, 214528 (2008).

Supplementary Information is linked to the online version of the paper at www.nature.com/nature.

Acknowledgements We thank G. Kirchmair, M. Mirrahimi, I. Chuang and M. Devoret for discussions. We acknowledge support from LPS/NSA under ARO contract no. W911NF-09-1-0514 and from the NSF under grants no. DMR-0653377 and no. DMR-1004406. Additional support was provided by CNR-Istituto di Cibernetica, Pozzuoli, Italy (L.F.), the Swiss NSF (S.E.N.) and the Dutch NWO (L.D.C.).

Author Contributions M.D.R. carried out measurements and performed data analysis. L.D.C. designed the three-qubit gate and conducted initial measurements. L.S. provided further experimental contributions. S.E.N. and S.M.G. provided theoretical support. L.F., L.D.C. and L.S. fabricated the devices. M.D.R. wrote the manuscript, with feedback from all authors. S.M.G. and R.J.S. designed and supervised the project.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of this article at www.nature.com/nature. Correspondence and requests for materials should be addressed to M.D.R. (matthew.reed@yale.edu) or R.J.S. (robert.schoelkopf@yale.edu).