

CMOS-based cryogenic control of silicon quantum circuits

Xue, Xiao; Patra, Bishnu; van Dijk, Jeroen P.G.; Samkharadze, Nodar; Corna, Andrea; Paquelet Wuetz, Brian; Sammak, Amir; Scappucci, Giordano; Veldhorst, Menno; Sebastiano, Fabio

DOI

[10.1038/s41586-021-03469-4](https://doi.org/10.1038/s41586-021-03469-4)

Publication date

2021

Document Version

Final published version

Published in

Nature

Citation (APA)

Xue, X., Patra, B., van Dijk, J. P. G., Samkharadze, N., Corna, A., Paquelet Wuetz, B., Sammak, A., Scappucci, G., Veldhorst, M., Sebastiano, F., Babaie, M., Charbon, E., & Vandersypen, L. M. K. (2021). CMOS-based cryogenic control of silicon quantum circuits. *Nature*, *593*(7858), 205-210. <https://doi.org/10.1038/s41586-021-03469-4>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' – Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

CMOS-based cryogenic control of silicon quantum circuits

<https://doi.org/10.1038/s41586-021-03469-4>

Received: 1 October 2020

Accepted: 18 March 2021

Published online: 12 May 2021

 Check for updates

Xiao Xue^{1,2,8}, Bishnu Patra^{1,2,3,8}, Jeroen P. G. van Dijk^{1,2,3,8}, Nodar Samkharadze^{1,4}, Sushil Subramanian⁵, Andrea Corna^{1,2}, Brian Paquelet Wuetz^{1,2}, Charles Jeon⁵, Farhana Sheikh⁵, Esdras Juarez-Hernandez⁶, Brando Perez Esparza⁶, Huzaifa Rampurawala⁵, Brent Carlton⁵, Surej Ravikumar⁵, Carlos Nieva⁵, Sungwon Kim⁵, Hyung-Jin Lee⁵, Amir Sammak^{1,4}, Giordano Scappucci^{1,2}, Menno Veldhorst^{1,2}, Fabio Sebastiano^{1,3,9}, Masoud Babaie^{1,3,9}, Stefano Pellerano^{5,9}, Edoardo Charbon^{1,2,5,7,9}✉ & Lieven M. K. Vandersypen^{1,2,5,9}✉

The most promising quantum algorithms require quantum processors that host millions of quantum bits when targeting practical applications¹. A key challenge towards large-scale quantum computation is the interconnect complexity. In current solid-state qubit implementations, an important interconnect bottleneck appears between the quantum chip in a dilution refrigerator and the room-temperature electronics. Advanced lithography supports the fabrication of both control electronics and qubits in silicon using technology compatible with complementary metal oxide semiconductors (CMOS)². When the electronics are designed to operate at cryogenic temperatures, they can ultimately be integrated with the qubits on the same die or package, overcoming the ‘wiring bottleneck’^{3–6}. Here we report a cryogenic CMOS control chip operating at 3 kelvin, which outputs tailored microwave bursts to drive silicon quantum bits cooled to 20 millikelvin. We first benchmark the control chip and find an electrical performance consistent with qubit operations of 99.99 per cent fidelity, assuming ideal qubits. Next, we use it to coherently control actual qubits encoded in the spin of single electrons confined in silicon quantum dots^{7–9} and find that the cryogenic control chip achieves the same fidelity as commercial instruments at room temperature. Furthermore, we demonstrate the capabilities of the control chip by programming a number of benchmarking protocols, as well as the Deutsch–Josza algorithm¹⁰, on a two-qubit quantum processor. These results open up the way towards a fully integrated, scalable silicon-based quantum computer.

A practical quantum computer comprises two main building blocks—a quantum processor with millions of qubits and classical instrumentation to generate control signals (input) and to process readout signals (output)^{1,3}. A standard setup for semiconducting or superconducting qubits has the qubits operating in a dilution refrigerator at about 20 mK, whereas bulky microwave vector sources and arbitrary waveform generators are at room temperature (RT) and connected to the qubits via long cables and attenuators (Fig. 1a, left). This approach has recently enabled an experimental demonstration of the advantage of quantum computing over classical computing in a random circuit sampling experiment that utilizes a superconducting quantum processor consisting of 53 qubits¹¹. This system requires more than 200 coaxial control lines from RT to the quantum chip, which is operated below 20 mK. This brute-force approach to reach higher qubit numbers will soon reach its limits.

A promising path forward is to bring the control electronics close to the quantum chip, at cryogenic temperatures. Although important steps in this direction have been taken^{4,5,12–19}, high-fidelity multi-qubit control and a universal gate set remain to be demonstrated using cryogenic controllers. A central challenge is that the power dissipation of the control electronics easily surpasses the typical cooling power of 10 μ W available at 20 mK (refs. ^{20,21}). Because silicon spin qubits can be operated and measured above 1 K (refs. ^{22–24}), they are well positioned for overcoming the wiring bottleneck by on-die or on-package co-integration with classical electronics (Fig. 1a, right) at a temperature of 1–3 K, where the cooling power is orders of magnitude higher than at millikelvin temperatures.

A cryogenic quantum controller for practical quantum information processing must meet multiple criteria: a form factor compatible with

¹QuTech, Delft University of Technology, Delft, The Netherlands. ²Kavli Institute of Nanoscience, Delft University of Technology, Delft, The Netherlands. ³Department of Quantum and Computer Engineering, Delft University of Technology, Delft, The Netherlands. ⁴Netherlands Organization for Applied Scientific Research (TNO), Delft, The Netherlands. ⁵Intel Corporation, Hillsboro, OR, USA. ⁶Intel Guadalajara, Zapopan, Mexico. ⁷École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland. ⁸These authors contributed equally: Xiao Xue, Bishnu Patra, Jeroen P. G. van Dijk. ⁹These authors jointly supervised this work: Fabio Sebastiano, Masoud Babaie, Stefano Pellerano, Edoardo Charbon, Lieven M. K. Vandersypen. ✉e-mail: edoardo.charbon@epfl.ch; L.m.vandersypen@tudelft.nl

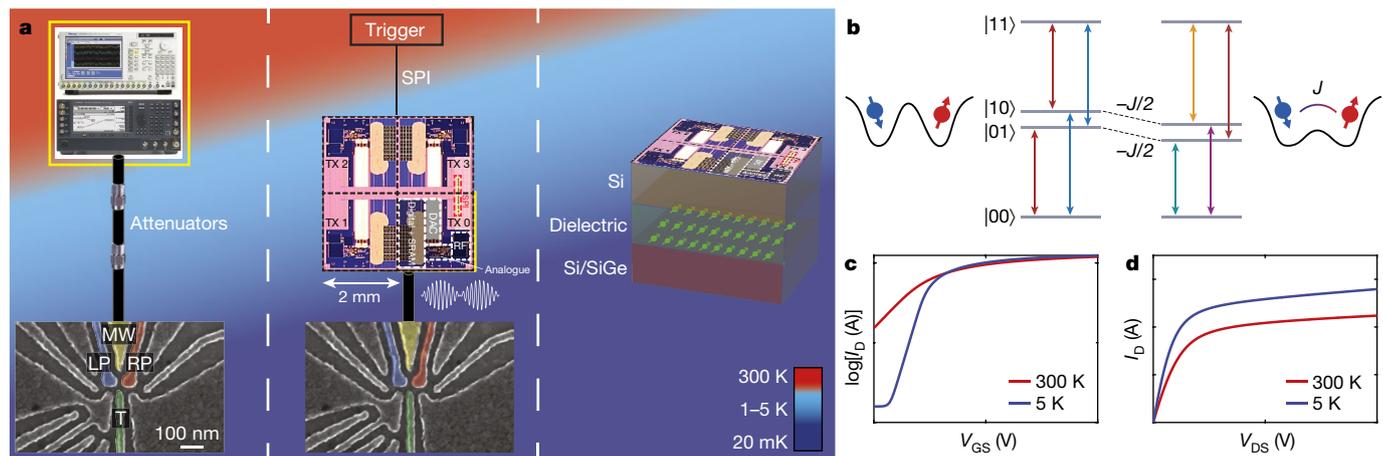


Fig. 1 | The cryogenic quantum control system. **a**, Three stages of development of the control system towards full integration. Left, RT instruments connected to qubits via coaxial lines and attenuators. Middle, cryo-controller at 1–5 K (see Extended Data Fig. 2) directly connected to the qubits and triggered from RT using an SPI, which leaves the wiring from 1–5 K to the qubit sample unaltered, but brings a considerable reduction in the wire count from RT to 1–5 K when targeting many qubits. Right, a future perspective of fully integrated control electronics and qubits on the same package/die, eliminating dense wiring all the way down to the package/die (note that we do not envision one transmitter above every qubit, so the qubit and transmitter

form factor can be different^{3,42}). Two single electron spins used as qubits are located underneath gates LP (blue) and RP (red), as shown in the SEM image. Multiplexed microwave signals are sent to gate MW (yellow) to control both qubits. Gate T (green) is used to tune the coupling between the qubits. **b**, Energy level diagram without (left) and with (right) exchange coupling (J). The resonance frequency of each qubit depends on the other qubit only when the coupling is on (low tunnel barrier between the dots). **c, d**, FinFET NMOS device characteristics at RT versus 5 K (temperature of the sample holder). Drain current (I_D) versus gate–source voltage (V_{GS}) at drain–source voltage of $V_{DS} = 1$ V (**c**) and I_D versus V_{DS} at $V_{GS} = 0.4$ V (**d**).

integration in a cryogenic refrigerator; frequency multiplexing to facilitate scalability; low power consumption within the limit of the refrigerator cooling power; sufficiently high output power to enable fast operations compared to the qubit coherence times; high signal-to-noise ratio (SNR) and spurious-free dynamic range (SFDR) for high-fidelity control; the ability to generate tailored pulse shapes and perform a universal set of quantum operations; an integrated instruction set memory for the efficient execution of complex algorithms. All these requirements can be met by commercial CMOS circuits designed to operate at a few kelvin.

In this work, we utilize a quantum control chip operating at 3 K (cryo-controller, named Horse Ridge) and fabricated using Intel 22-nm-FinFET low-power CMOS technology¹³ to coherently control two electron spin qubits in a silicon double quantum dot. To benchmark the limits of the controller, we keep the qubits at about 20 mK, where they are most coherent and the non-idealities of the control chip can be assessed best (Fig. 1a, middle). Extensive electrical characterization and benchmarking using the quantum processor show that the cryo-controller meets all the above criteria.

The specifications for the cryo-controller derive from the demands on the qubit control. Here we target qubits that can be resonantly controlled with drive frequencies in the 2–20 GHz band, covering the typical resonance frequencies of both superconducting and spin qubits^{7–9,12,22}. The cryo-controller has four output ports, each with up to 32 frequency-multiplexed tones. Because the controller must dissipate minimal power and have a small form factor, we analyse in detail the signal specifications that are sufficient to achieve a 99.99% gate fidelity²⁵. Among other performance metrics, the most stringent ones that dominate the architecture and power consumption of the controller are the SNR (>48 dB) and SFDR (>44 dB) for frequency-multiplexed control²⁵.

Further challenges arise in designing complex CMOS circuits at deep cryogenic temperatures. Key device characteristics such as the threshold voltage (V_{th}) and the mobility (μ) increase compared to RT, as seen in Fig. 1c, d²⁶. Moreover, the degradation of active device matching²⁷ and the improvement of the quality factor of on-chip passive components²⁸ necessitate careful characterization and modelling for circuits operated at cryogenic temperatures.

As a benchmark of performance, we use the cryo-controller to coherently control a two-qubit quantum processor. The quantum processor is made of a double quantum dot (DQD) electrostatically confined in an undoped ²⁸Si/SiGe heterostructure. By tuning the voltage on the plunger gates LP and RP, two single electrons are locally accumulated underneath each gate, shown in blue and red in the scanning electron microscope (SEM) image in Fig. 1a. By applying an external magnetic field of 380 mT, combined with the longitudinal magnetic field induced by a micro-magnet on top of the DQD (see Extended Data Fig. 7), we can encode the qubit states into the Zeeman split states of the two electrons, where spin up is used as |1> and spin down as |0>. The resonance frequencies of qubit 1 (Q_1 ; underneath gate LP) and qubit 2 (Q_2 ; underneath gate RP) are 13.62 GHz and 13.51 GHz, respectively. Rotations around the \hat{x} and \hat{y} axes are implemented by sending microwave bursts, with the microwave phase controlling the rotation axis; for example, an in-phase (quadrature) microwave burst implements a rotation about \hat{x} (\hat{y}). The microwave bursts are applied to gate MW, which drives an electric-dipole spin resonance enabled by the transverse magnetic field gradient from the micro-magnet²⁹, whereas rotation around the \hat{z} axis (phase control) is achieved by changing the reference phase in the cryo-controller, which adds a phase shift to all the subsequent bursts³⁰. The two-qubit interaction is mediated by the exchange coupling (J) between the two spins³¹, controlled by gate T. Its effect here is to shift the anti-parallel spin states down in energy³². As a result, the resonance frequency of each qubit now depends on the state of the other qubit, allowing conditional operations on each qubit via narrow-band microwave bursts^{8,9} (Fig. 1b). The corresponding four different frequencies can be individually addressed using frequency multiplexing. Both qubits are read out in single-shot mode³³ (see Methods).

Figure 2 shows the system-level architecture of one transmitter module (TX) in the cryo-controller, which consists of a digital signal-generation unit with an analogue/radiofrequency (RF) front end. At the core of the digital signal generation, a numerically controlled oscillator (NCO) outputs a sequence of bit strings every clock period³⁴. This bit string encodes a phase that is intended to track the reference phase of one particular qubit. The output of 16 NCOs is multiplexed

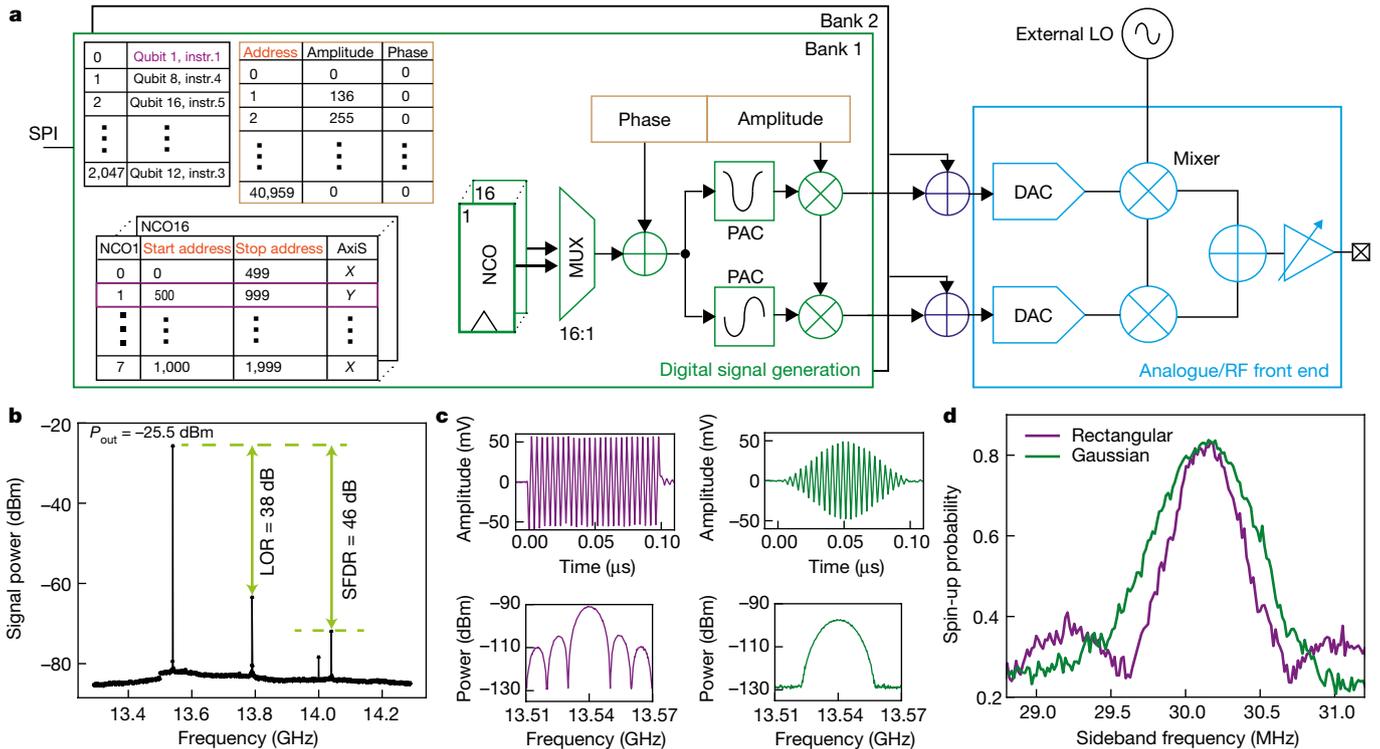


Fig. 2 | The Horse Ridge cryogenic controller characterized at 3 K.

a, System-level representation showing the digital signal generation and analogue/RF front end of the cryo-controller, programmable via the SPI. **b**, Continuous-wave output spectrum from the cryo-controller at 13.54 GHz using only bank 1 (for both banks, see Extended Data Fig. 5c), showing the main output tone, LO rejection ratio (LOR) and SFDR limited by the image tone.

(by ‘MUX’ in Fig. 2a) and fed to a phase-to-amplitude converter (PAC) implemented as a hard-wired look-up table, to generate a sinusoidal (in-phase) and cosinusoidal (quadrature) signal. The NCO phases are constructed via a phase accumulator, which increments the phase in steps determined by a digital frequency tuning word (FTW). The 22-bit ($N = 22$) FTWs in combination with the 1-GHz clock frequency (f_s) of the phase accumulator gives a frequency resolution of about 238 Hz, that is, $f_s/2^N$.

The sine and cosine signals are amplitude- and phase-modulated using an envelope memory (orange box) containing up to 40,960 points, each specifying an amplitude and phase value. An instruction table memory can store up to eight different instructions per qubit/NCO by referring to start and stop addresses in the envelope memory. Finally, these instructions are listed in the instruction list to execute up to 2,048 instructions from multiple instruction tables, initiated by a single external trigger, and acting sequentially on up to 16 qubits. The output of two such banks, each generating a digital signal, are summed to simultaneously control two qubits, consequently increasing the number of supported (uncoupled) qubits from 16 to 32 per TX module.

The generated digital signals are translated to the analogue domain using 1 GS s^{-1} ($\text{GS}, 10^9$ samples) 10-bit current-steering digital-to-analogue converters (DACs), low-pass-filtered with an adjustable bandwidth to remove the sampling replicas, and upconverted to the required qubit frequency using an I/Q mixer and an external local oscillator (LO). Finally, an output driver is incorporated to produce the required voltage amplitude (through a tunable gain of 40 dB) in the frequency range of 2–20 GHz while driving the 50- Ω coaxial cable connecting to the qubits. Such a wide frequency and output power range allows the control of various solid-state qubits, such as spin qubits

and superconducting qubits. The controller dissipates 384 mW with all the NCOs simultaneously operating at a clock frequency of 1 GHz (digital signal generation, 330 mW; analogue/RF front end, 54 mW) (see Methods). This architecture is replicated four times in a die area of 16 mm² (TX0–TX3 in Fig. 1), with an ability to control up to 4×32 frequency-multiplexed qubits.

c, Rectangular (purple) and Gaussian (green) shaped bursts before up-conversion (baseband signal) and the corresponding spectra after up-conversion. **d**, Qubit response for different burst envelopes, obtained when sweeping the NCO frequency around the qubit resonance across a span of about 3 MHz with a resolution of 15 kHz.

The purity of the generated signal can be quantified using the output signal spectrum shown in Fig. 2b. The generated signal has an SFDR of 46 dB at 13.54 GHz in a 1-GHz bandwidth, excluding the residual LO leakage (see Extended Data Fig. 5 for a two-tone test). The noise floor is flat across the 1-GHz bandwidth, and the cryo-controller leaves the electron temperature of the quantum device unaffected (see Extended Data Fig. 8). The SNR is 48 dB when integrating over 25 MHz, corresponding to the targeted maximum qubit Rabi frequency. Along with the low quantization noise and frequency noise, the output signal quality is predicted to achieve a single-qubit gate fidelity of 99.99%, assuming ideal qubits²⁵. The amplitude and phase modulation capabilities of the controller allow the chip to generate arbitrary waveforms to precisely shape the spectral content of the pulse used to manipulate the qubits, as shown in Fig. 2c. In illustration, Fig. 2d shows the response of Q_2 to a microwave burst with a rectangular envelope versus a Gaussian one, both calibrated to invert the qubit state when the drive is on-resonance with the qubit.

Next, we test the functionality of the cryo-controller for controlling uncoupled qubits. The LO frequency is set to 13.54 GHz. Q_1 is then offset from the LO by 24 MHz and Q_2 by -90 MHz. The qubit resonances are found by sweeping one single-sideband tone generated by one NCO (Fig. 3a), using the 22-bit FTW. Then, we use one NCO from each bank to generate two tones on resonance with the two qubits and drive

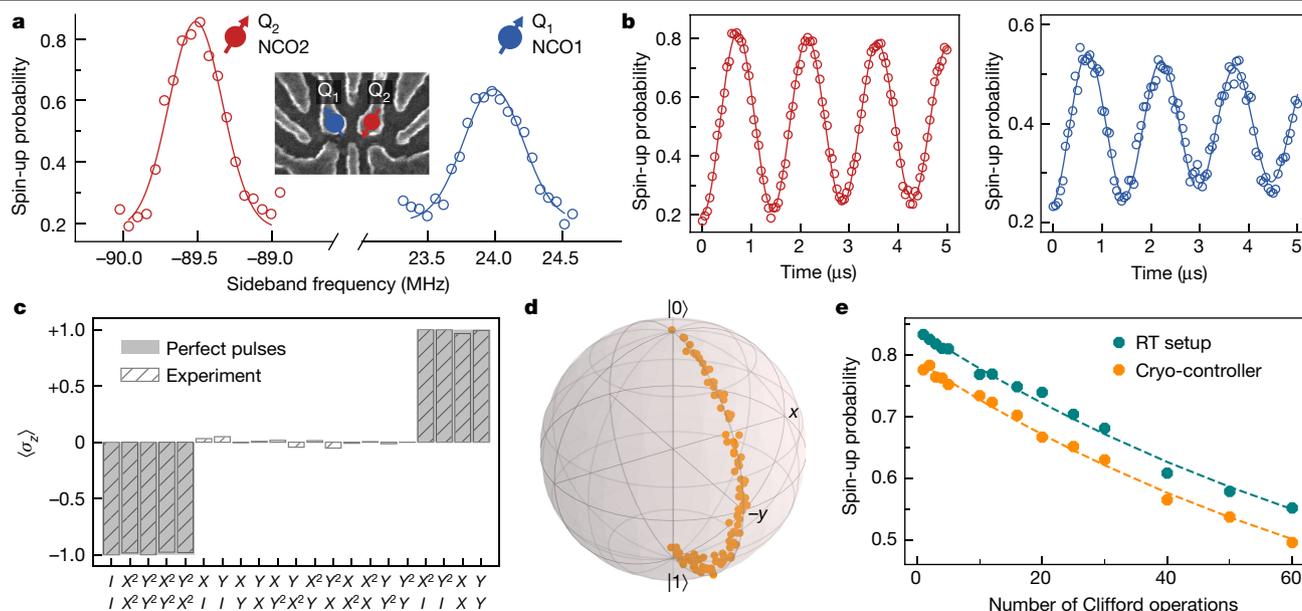


Fig. 3 | Frequency-multiplexed qubit control and fidelity benchmarks with the cryo-controller. **a**, Spectra showing the qubit resonances. Inset, SEM image indicating the positions of the qubits. **b**, Frequency-multiplexed control producing simultaneous Rabi oscillations of Q_2 (left) and Q_1 (right). The decay arises mainly from the residual coupling between the two qubits (see Extended Data Fig. 9 for Rabi oscillations in individual driving mode). **c**, $\langle \sigma_z \rangle$ of Q_2 measured after an AllXY sequence consisting of 21 different pairs of gates, each listed vertically on the X axis. The output power is calibrated to achieve a Rabi

frequency of about 1 MHz (the same applies to the QST and RB experiments). The visibility is normalized by removing the readout error (see Methods). **d**, Trajectory of the state of Q_2 under an X^2 gate reconstructed by QST. Orange data points indicate the qubit state after incrementing microwave burst times. **e**, Randomized benchmarking of Q_2 performed by the cryo-controller and the RT setup. We offset the orange data points by -0.05 along the Y axis to facilitate comparison of the two traces.

simultaneous Rabi oscillations on both qubits (Fig. 3b). Here a 5- μ s rectangular envelope is uploaded to the envelope memory, and saved as an instruction. The duration of the microwave burst is swept by updating the start or stop address of this instruction.

The pulses for single-qubit rotations are precisely calibrated using the AllXY sequence³⁵. In the AllXY experiment, 21 different pairs of single-qubit gates from the set $\{I, X, Y, X^2, Y^2\}$ are applied to a qubit initialized to $|0\rangle$. Here I is the identity operation, X and Y are $\pi/2$ rotations around the \hat{x} and \hat{y} axis, respectively, and X^2 and Y^2 are π rotations. The final-state \hat{z} -projection $\langle \sigma_z \rangle$ takes values from $\{-1, 0, +1\}$ under perfect operations (shown as the grey shaded areas in Fig. 3c). Any miscalibration in the amplitude, frequency or phase of the pulse results in deviations from the ideal outcome (hatched bars in Fig. 3c). In addition, we reconstruct the trajectory of an X^2 gate by performing quantum state tomography (QST)³⁶ at incremental burst times of a rectangular microwave signal (Fig. 3d), which shows an average state fidelity of 97.92%. Here the infidelity comes from the error during the operation, as well as the errors in the initialization and readout of the qubit (see Methods). The AllXY and QST results indicate that the single-qubit gate set is well calibrated, offering a good starting point for benchmarking the gate fidelity.

The gate fidelity is a crucial metric used to express the performance of a quantum processor and its classical controller. We use single-qubit randomized benchmarking (RB)^{37,38} to compare the performance of the cryo-controller with the conventional RT setup, which consists of an arbitrary waveform generator (Tektronix 5014C) and a vector signal generator (Keysight E8267D). A programmable mechanical microwave switch placed at the 3-K plate allows us to conveniently alternate between the cryo-controller and the RT setup. In the RB experiment, sequences of increasing numbers of randomly selected Clifford operations are applied to the qubit (Q_2), followed by a final Clifford operation that returns the qubit to its initial state in the ideal case. For each data point in Fig. 3e, 32 different sequences are randomly sampled, and each is repeated 200 times. Envelopes of all gates to be used are uploaded to

the envelope memory and saved as instructions. The random sequences are constructed by updating the instruction list. The instructions in the list are executed sequentially after an external trigger is received via the serial peripheral interface (SPI) in Fig. 2a. Exactly the same random sequences are used in an RB experiment using the RT setup. We find an average single-qubit gate fidelity of $99.71 \pm 0.03\%$ with the RT setup and $99.69 \pm 0.02\%$ with the cryo-controller (see Methods; all uncertainties are one standard deviation). The fidelities are consistently identical within the error bars and well above the threshold for fault tolerance³⁹, with the infidelity limited by the qubit. These experiments demonstrate the high quality of the signal from the cryo-controller, as well as its capability of generating complex sequences.

To further test the programmability of the cryo-controller, we use it to implement two-qubit logic in the quantum processor (Fig. 4a). Taking advantage of the frequency shift of each qubit being conditional on the state of the other qubit (Fig. 1b), we use controlled-rotation (CROT) gates as the native two-qubit gates. These are achieved by frequency-selective addressing^{8,9}, thus demanding two NCOs per qubit (see Methods). A π rotation at the higher or lower frequency implements the canonical controlled NOT (CNOT) gate or the zero-controlled NOT (Z-CNOT) gate, respectively, up to a single-qubit $\pi/2$ \hat{z} rotation on the control qubit. Because of cross-talk, an additional phase correction in the form of a \hat{z} rotation is needed. All \hat{z} rotations are implemented by updating the reference phase of the NCO (see Extended Data Fig. 3). Single-qubit gates are implemented by addressing both frequencies of the same qubit sequentially. Making use of four NCOs, we program the cryo-controller to run the two-qubit Deutsch–Josza algorithm, which determines whether a function gives constant or balanced outcomes¹⁰. The two constant (balanced) functions that map one input bit on one output bit are implemented by the CNOT and Z-CNOT (I and X^2) operations. Here, we choose Q_1 to be the output qubit and Q_2 to be the input qubit. Fig. 4b shows the pulse sequence and the measurement results, where the constant (balanced) functions lead to a high probability of 78%–80% (79%–82%) for

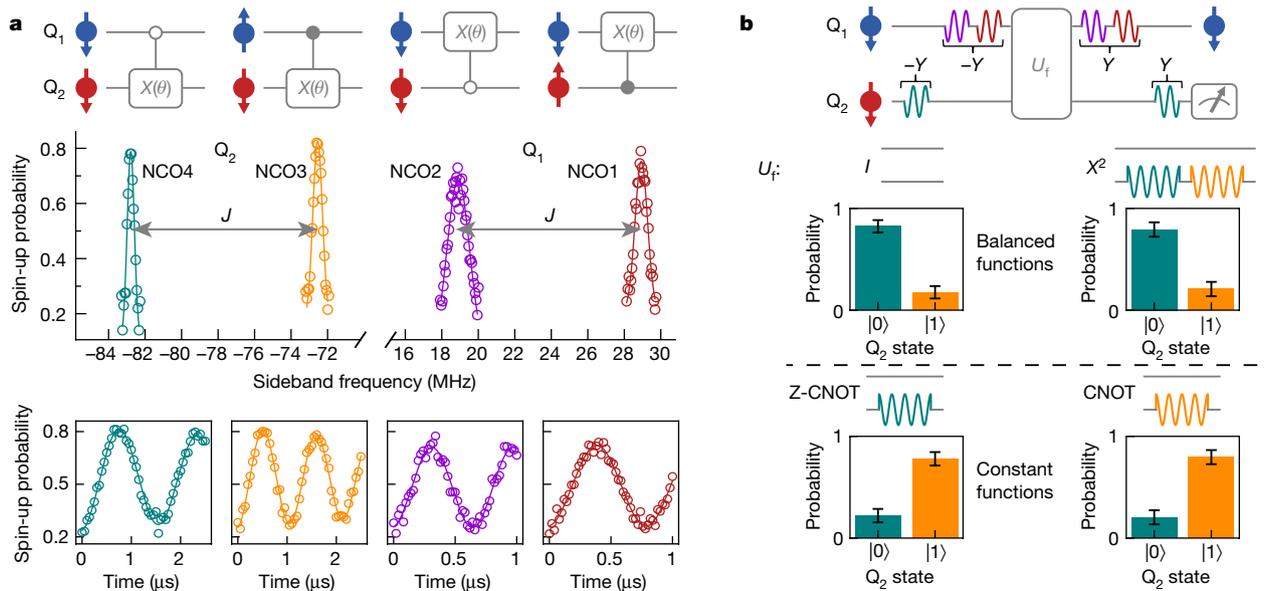


Fig. 4 | Programming a quantum processor with the cryo-controller.

a, Two-qubit logic with the cryo-controller. The middle panel shows the spectra of two qubits obtained using the cryo-controller with the exchange coupling (J) between the qubits turned on. Selective excitation of each of the four resonances can be used for implementing various two-qubit CROT gates, shown in the upper panel. The lower panels (shared Y -axis labels) show the Rabi oscillations at each frequency. **b**, Pulse sequences of the Deutsch–Jozsa algorithm programmed into the cryo-controller (top) and measured probabilities of the output qubit state (Q_2) after running the algorithm

(bottom) for constant or balanced functions, implemented through the unitary operation U_f . A constant function is composed of either a CNOT or a Z-CNOT gate, which consists of a CROT gate on Q_2 and a phase correction on Q_1 (not plotted). Only the lower frequency (green branch, Z-CROT) is used for the $-Y$ and Y gates on Q_2 , because Q_1 (ideally) starts from and ends up in $|0\rangle$. Error bars are standard deviations of the measured probability data. The visibility of Q_2 is normalized by removing the readout error. Empirically, we attribute the remaining errors mostly to charge noise in the presence of a finite J (see Methods).

measuring the data qubit as $|1\rangle$ ($|0\rangle$), as expected. This experiment highlights the ability to program the cryo-controller with arbitrary sequences of operations.

The cryo-controller allows for much more complex sequences, containing up to 2,048 instructions for each of the four transmitters. Each instruction defines a microwave burst at one of 32 independent frequencies, with an amplitude and phase profile that can be arbitrarily shaped. The cryo-controller can be conveniently embedded in existing micro-architectures and programmed via standard QASM variants⁴⁰. This quantum–classical architecture can thus be directly applied to multi-qubit algorithms and noisy intermediate-scale quantum devices⁴¹.

The versatile programmability, combined with a signal quality allowing up to 99.99% gate fidelities, a footprint of just 4 mm², a power consumption of 384 mW, the ability to integrate multiple transmitters on one die, and operation at 3 K demonstrate the potential of the cryo-controller to address key challenges in building a large-scale quantum computer. Optimized design of cryogenic CMOS circuits—for example, the use of a narrower frequency band—can substantially reduce the power consumption (see Methods) and make it possible to work at 1 K or even lower temperatures. Furthermore, FinFET quantum dots that are fully compatible with CMOS processing² and increased operating temperatures (about 1 K) of spin qubits show only a modest reduction in coherence times^{22,23}. These advances imply that it may be possible to fully integrate the quantum processor with the classical controller on chip or by flip-chip technology, lifting a major roadblock in scaling.

Online content

Any methods, additional references, Nature Research reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions

and competing interests; and statements of data and code availability are available at <https://doi.org/10.1038/s41586-021-03469-4>.

1. Van Meter, R. & Horsman, C. A blueprint for building a quantum computer. *Commun. ACM* **56**, 84–93 (2013).
2. Pillarisetty, R. et al. Qubit device integration using advanced semiconductor manufacturing process technology. In *2018 IEEE International Electron Devices Meeting 6.3.1–6.3.4* (IEEE, 2018).
3. Vandersypen, L. M. K. et al. Interfacing spin qubits in quantum dots and donors—hot, dense, and coherent. *npj Quantum Inf.* **3**, 34 (2017).
4. Patra, B. et al. Cryo-CMOS circuits and systems for quantum computing applications. *IEEE J. Solid-State Circuits* **53**, 309–321 (2018).
5. Pauka, S. J. et al. A cryogenic CMOS chip for generating control signals for multiple qubits. *Nat. Electronics* **4**, 64–70 (2021).
6. Geck, L., Kruth, A., Bluhm, H., van Waasen, S. & Heinen, S. Control electronics for semiconductor spin qubits. *Quantum Sci. Technol.* **5**, 015004 (2019).
7. Watson, T. F. et al. A programmable two-qubit quantum processor in silicon. *Nature* **555**, 633–637 (2018).
8. Zajac, D. M. et al. Resonantly driven CNOT gate for electron spins. *Science* **359**, 439–442 (2018).
9. Huang, W. et al. Fidelity benchmarks for two-qubit gates in silicon. *Nature* **569**, 532–536 (2019).
10. Deutsch, D. & Jozsa, R. Rapid solution of problems by quantum computation. *Proc. R. Soc. Lond. A* **439**, 553–558 (1992).
11. Arute, F. et al. Quantum supremacy using a programmable superconducting processor. *Nature* **574**, 505–510 (2019).
12. Bardin, J. C. et al. Design and characterization of a 28-nm bulk-CMOS cryogenic quantum controller dissipating less than 2 mW at 3 K. *IEEE J. Solid-State Circuits* **54**, 3043–3060 (2019).
13. Patra, B. et al. A scalable cryo-CMOS 2-to-20GHz digitally intensive controller for a 4×32 frequency multiplexed spin qubits/transmons in 22nm FinFET technology for quantum computers. In *2020 IEEE International Solid-State Circuits Conference* 304–306 (IEEE, 2020).
14. Le Guevel, L. et al. A 110mK 295μW 28nm FDSOI CMOS quantum integrated circuit with a 2.8 GHz excitation and nA current sensing of an on-chip double quantum dot. In *2020 IEEE International Solid-State Circuits Conference* 306–308 (IEEE, 2020).
15. Bonen, S. et al. Cryogenic characterization of 22-nm FDSOI CMOS technology for quantum computing ICs. *IEEE Electron. Device Lett.* **40**, 127–130 (2018).
16. Esmailyan, A. et al. A fully integrated DAC for CMOS position-based charge qubits with single-electron detector loopback testing. *IEEE Solid-State Circuits Lett.* **3**, 354–357 (2020).
17. Ekanayake, S. R. et al. Characterization of SOS-CMOS FETs at low temperatures for the design of integrated circuits for quantum bit control and readout. *IEEE Trans. Electron Dev.* **57**, 539–547 (2010).

18. Mukhanov, O. et al. Scalable quantum computing infrastructure based on superconducting electronics. In *2019 IEEE International Electron Devices Meeting* 31.2.1–31.2.4 (IEEE, 2019).
19. Xu, Y. et al. On-chip integration of Si/SiGe-based quantum dots and switched-capacitor circuits. *Appl. Phys. Lett.* **117**, 144002 (2020).
20. Batey, G., Matthews, A. J. & Patton, M. A new ultralow-temperature cryogen-free experimental platform. *J. Phys. Conf. Ser.* **568**, 032014 (2014).
21. Green, M. A. The cost of coolers for cooling superconducting devices at temperatures at 4.2 K, 20 K, 40 K and 77 K. In *IOP Conference Series: Materials Science and Engineering* Vol. 101, 012001 (IOP, 2015).
22. Petit, L. et al. Universal quantum logic in hot silicon qubits. *Nature* **580**, 355–359 (2020).
23. Yang, C. H. et al. Operation of a silicon quantum processor unit cell above one kelvin. *Nature* **580**, 350–354 (2020).
24. Urdampilleta, M. et al. Gate-based high fidelity spin readout in a CMOS device. *Nat. Nanotechnol.* **14**, 737–741 (2019).
25. van Dijk, J. P. G. et al. Designing a DDS-based SoC for high-fidelity multi-qubit control. *IEEE Trans. Circuits Syst. I* **67**, 5380–5393 (2020).
26. Beckers, A., Jazaeri, F. & Enz, C. Characterization and modeling of 28-nm bulk CMOS technology down to 4.2 K. *IEEE J. Electr. Dev. Soc.* **6**, 1007–1018 (2018).
27. Hart, P. A. T., Babiak, M., Charbon, E., Vladimirescu, A. & Sebastiano, F. Subthreshold mismatch in nanometer CMOS at cryogenic temperatures. *IEEE J. Electr. Dev. Soc.* **8**, 797–806 (2020).
28. Patra, B. et al. Characterization and analysis of on-chip microwave passive components at cryogenic temperatures. *IEEE J. Electr. Dev. Soc.* **8**, 448–456 (2020).
29. Pioro-Ladrière, M. et al. Electrically driven single-electron spin resonance in a slanting Zeeman field. *Nat. Phys.* **4**, 776–779 (2008).
30. Vandersypen, L. M. K. & Chuang, I. L. NMR techniques for quantum control and computation. *Rev. Mod. Phys.* **76**, 1037–1069 (2005).
31. Petta, J. R. et al. Coherent manipulation of coupled electron spins in semiconductor quantum dots. *Science* **309**, 2180–2184 (2005).
32. Meunier, T., Calado, V. E. & Vandersypen, L. M. K. Efficient controlled-phase gate for single-spin qubits in quantum dots. *Phys. Rev. B* **83**, 121403 (2011).
33. Xue, X. et al. Repetitive quantum nondemolition measurement and soft decoding of a silicon spin qubit. *Phys. Rev. X* **10**, 021006 (2020).
34. Saul, P. H. & Mudd, M. S. J. A direct digital synthesizer with 100-MHz output capability. *IEEE J. Solid-State Circuits* **23**, 819–821 (1988).
35. Reed, M. Entanglement and Quantum Error Correction with Superconducting Qubits. PhD Thesis, Yale Univ. (2013).
36. Altepetter, J. B., Jeffrey, E. R. & Kwiat, P. G. Photonic state tomography. *Adv. At. Mol. Opt. Phys.* **52**, 105–159 (2005).
37. Knill, E. et al. Randomized benchmarking of quantum gates. *Phys. Rev. A* **77**, 012307 (2008).
38. Magesan, E., Gambetta, J. M. & Emerson, J. Characterizing quantum gates via randomized benchmarking. *Phys. Rev. A* **85**, 042311 (2012).
39. Fowler, A. G., Mariantoni, M., Martinis, J. M. & Cleland, A. N. Surface codes: towards practical large-scale quantum computation. *Phys. Rev. A* **86**, 032324 (2012).
40. Svore, K. M., Aho, A. V., Cross, A. W., Chuang, I. & Markov, I. L. A layered software architecture for quantum computing design tools. *Computer* **39**, 74–83 (2006).
41. Preskill, J. Quantum computing in the NISQ era and beyond. *Quantum* **2**, 79 (2018).
42. Boter, J. M. et al. A sparse spin qubit array with integrated control electronics. In *2019 IEEE International Electron Devices Meeting* 31.4.1–31.4.4 (IEEE, 2019).

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

© The Author(s), under exclusive licence to Springer Nature Limited 2021

Methods

Programming the cryo-controller

The setup (Extended Data Fig. 1) contains a field-programmable gate array (FPGA) that configures the cryo-controller (for example, FTW), programs the various memories inside the cryo-controller (for example, envelope memories, instruction tables and instruction lists) and controls the start of the execution of the instruction list. The FPGA is connected to the host PC, which sends the data that need to be uploaded to the cryo-controller over the SPI. The instruction list integrated in the cryo-controller does not support classical instructions that allow for, for example, branching or wait statements, as required for performing certain qubit experiments and for synchronization with other equipment. Therefore, switching between different instruction lists and synchronization with the rest of the equipment are controlled by two trigger lines from the AWG to the FPGA. The application of the *execute* trigger starts the execution of the instruction list that is programmed in the cryo-controller for performing repeated measurements. The application of the *sweep* trigger loads the next instruction list from the static random-access memory (SRAM) of the FPGA into the cryo-controller's instruction list.

Power budget

All memory blocks except NCO phase-update registers (to perform phase corrections; 'Z-Corr.' in Extended Data Fig. 3a) are implemented using SRAM. The high power consumption of the digital circuitry of the cryo-controller (Extended Data Fig. 4) is due to the lack of clock gating in registers, thus causing them to operate continuously instead of only during the read/write cycle. This could easily be reduced by further optimizations (for example, by replacing more registers with SRAM memory and by adding clock gating) that were not included in the first-generation cryo-controller. On the basis of the Cadence simulation with clock-gating, the power consumption of the digital circuitry should be lower than 40 mW instead of 330 mW in the current design. Migrating to a finer technology node would result in further substantial power savings²⁵.

Moreover, this chip was designed to address both transmons and spin qubits, and hence an ultrawide output frequency range was supported, that is, 2–20 GHz with an LO frequency of 2.5–14.5 GHz, using both fundamental and third-harmonic outputs (see Extended Data Fig. 3b). Once the qubit frequency is fixed within a range of a few gigahertz, the power consumption of the analogue circuitry can be substantially reduced to limit the power consumption to ~20 mW instead of 54 mW. In the current architecture (Extended Data Fig. 3 and ref.¹³), such power savings can be achieved by: (1) eliminating the output drivers and compensating the gain loss by increasing the matching network impedance transformation, as allowed by the reduced frequency range, which also results in a higher mixer load and consequently lower VGA bias current; (2) replacing the folded-current topology of the anti-aliasing filter with a stacked topology; (3) replacing the active balun before the LO drivers with a passive matching network. Further power savings could be achieved by architecture improvements, for example, by replacing the gm-C filter with a passive filter, thus eliminating the power consumption of the filter²⁵. Moreover, the integration of a phase-locked loop can eliminate I/Q active baluns completely.

Si/SiGe heterostructure

The ²⁸Si/SiGe heterostructure is grown on a 100-mm n-type Si(001) substrate using an Epsilon 2000 (ASMI) reduced-pressure chemical vapour deposition reactor equipped with a ²⁸SiH₄ gas cylinder (1% dilution in H₂) for the growth of isotopically enriched ²⁸Si. The ²⁸SiH₄ gas is obtained by reducing ²⁸SiF₄ with a residual ²⁹Si concentration of 0.08%, as described in ref.⁴³. Starting from the Si substrate, the layer sequence comprises a 900-nm linearly graded Si_(1-x)Ge_x layer (Ge concentration *x* varied from 0 to 0.3), followed by a 300-nm strain-relaxed Si_{0.7}Ge_{0.3} buffer layer, an

8-nm tensile strained ²⁸Si quantum well, a 30-nm Si_{0.7}Ge_{0.3} barrier and a 1-nm sacrificial Si cap. These undoped ²⁸Si/SiGe heterostructures are insulating at low temperature and support via gating the accumulation of a two-dimensional electron gas with transport mobility of up to 10⁵ cm²V⁻¹s⁻¹ at 55 mK (ref.⁴⁴). Extended Data Fig. 7 shows concentration profiles obtained by secondary ion mass spectroscopy (SIMS) of a control ²⁸Si/SiGe heterostructure. This control ²⁸Si/SiGe heterostructure has an increased quantum well thickness of 20 nm, which facilitates the investigation of the chemical composition therein by reducing the impact of the knock-on SIMS artefact⁴⁵. The concentration profiles of isotopes of Ge, ²⁸Si, ²⁹Si and ³⁰Si show a high-purity and homogeneous ²⁸Si quantum well. The residual concentration of non-zero-spin nuclei ²⁹Si is reduced from 3.29% in the Si_{0.7}Ge_{0.3} buffer and barrier to 0.08% in the quantum well, demonstrating that the ²⁸SiH₄ precursor purity is maintained during the heterostructure deposition process. As a result, we achieve a suitable quantum grade environment for the qubits in which decoherence due to hyperfine interaction is minimized. Furthermore, the concentration of common background contaminants C and O is below the detection limit of around 3 × 10¹⁶ cm⁻³ and 1 × 10¹⁷ cm⁻³, respectively, reducing scattering sources in the qubit-surrounding environment that can be sources of charge noise.

Quantum dot device fabrication

On top of the heterostructure, a 7-nm-thick AlO_x layer is deposited using atomic-layer deposition, followed by a 20-nm Al metal film, which is patterned using electron beam lithography in order to define the first gate layer, which shapes the potential landscape. Next, another 7-nm AlO_x layer is deposited, followed by a 70-nm Al layer that uniformly covers the quantum dot area. Finally, a 200-nm Co film is deposited and patterned into a micro-magnet (see Extended Data Fig. 7).

Qubit readout

The readout scheme is described in Extended Data Fig. 6. After each operation sequence, Q₂ is measured by spin-selective tunnelling to the electron reservoir, where a spin-up (|1>) electron can tunnel out and a spin-down (|0>) electron is blocked from tunnelling out. Such a spin-to-charge conversion changes the charge occupancy in the quantum dot, conditional on the spin state. This in turn changes the current signal in an adjacent capacitively coupled single-electron transistor (SET). Single-shot readout of the qubit state can be done by thresholding the current signal through the SET³³. The post-measurement state in this readout protocol is the |0> state, serving as reinitialization. Q₁ is tuned to be only weakly coupled to the SET, which serves as the electron reservoir for Q₁. This is to minimize the back-action from the SET, but also makes it less efficient to readout Q₁ by spin-selective tunnelling to the SET. Therefore, with Q₂ reinitialized, a CROT gate is applied to map the state of Q₁ onto Q₂. Then, Q₁ is read out by measuring Q₂ again³³. The readout fidelity of Q₂ is mainly limited by the thermal broadening of the electron reservoir, and the readout fidelity of Q₁ is limited by both the error in the CROT gate and in the readout of Q₂. Thus, the readout visibility of Q₁ is lower than that of Q₂. Alternatively, Q₁ could be read out by shuttling the electron to the location of Q₂ after emptying the dot hosting Q₂. We here chose to map the state of Q₁ onto that of Q₂ using a CROT gate, because this reduces charging effects on the bias tees on the printed circuit board (PCB). The CROT gate used here is also executed by the cryo-controller.

Readout error removal

In the AIIXY experiments and in the implementation of the Deutsch-Jozsa algorithm, the readout probabilities of Q₂ are normalized with the calibrated readout fidelities ($F_{|0\rangle}, F_{|1\rangle}$). After preparing Q₂ in |0>, $F_{|0\rangle}$ can be calibrated directly through the measured spin-down probability, and $F_{|1\rangle}$ is calibrated through the measured spin-up probability after a spin-flip operation (the spin-flip fidelity is above 99%). On the basis of the measured state probabilities in the AIIXY and Deutsch-Jozsa

Article

experiments, $P^M = (P_{|0\rangle}^M, P_{|1\rangle}^M)^T$, the actual state probabilities ($P_{|0\rangle}, P_{|1\rangle}$) can be reconstructed by $P = F^{-1}P^M$, where

$$F = \begin{pmatrix} F_{|0\rangle} & 1 - F_{|1\rangle} \\ 1 - F_{|0\rangle} & F_{|1\rangle} \end{pmatrix}. \quad (1)$$

Error sources

In the simultaneous Rabi oscillation experiment (Fig. 3b), we attribute the visible decays in both curves to the residual exchange coupling between the two qubits. Simultaneous Rabi oscillations recorded (in this case, using the RT setup) over larger numbers of oscillations show beating patterns. These patterns are well reproduced by numerical models of the spin evolution in the presence of a finite residual exchange coupling. Such a beating effect looks like a decay in the beginning. It is absent in the individually driven Rabi oscillation (Extended Data Fig. 9). In the two-qubit experiments shown in Fig. 4, the decay in the controlled-rotation Rabi oscillations and the finite visibilities in the Deutsch–Jozsa algorithm are largely attributed to charge noise. With the exchange coupling turned on, as needed for two-qubit gates, the energy levels are much more sensitive to charge noise.

Quantum state tomography

In the QST experiment, the qubit state is measured by projecting it onto the $(-\hat{z}, +\hat{x}, -\hat{y}, +\hat{z})$ axes. The projection on the $-\hat{z}$ axis is measured by direct readout of the spin state, whereas the projections on the other axes are measured by applying an X , Y or X^2 gate, which are calibrated by the ALLXY experiment, before the readout. The trajectory of the qubit state in the course of an X^2 gate can be reconstructed by performing QST at incremental burst times of a rectangular microwave signal (Fig. 3c), with each measurement repeated 1,000 times. To visualize the qubit state in the Bloch sphere, we remove the readout error from the data. Given that error removal can lead to unphysical states, such as data points outside the Bloch sphere, a maximum-likelihood estimation is implemented to find the closest physical state of the qubit³⁷. The ideal quantum state at each burst time can be predicted using the Rabi frequency fitted from experimental data. We then calculate the quantum state fidelities at each burst time and obtain an average state fidelity of 97.92%. The difference between this number and the fidelity obtained from randomized benchmarking is mostly

from residual state preparation and measurement error after imperfect readout error removal.

Data availability

Data supporting this work are available at <https://doi.org/10.5281/zenodo.4061970>.

Code availability

The codes used for data acquisition and processing are from the open-source Python packages QCoDeS (available at <https://github.com/QCoDeS/Qcodes>), QTT (available at <https://github.com/QuTech-Delft/qtt>) and PycQED (available at https://github.com/DiCarloLab-Delft/PycQED_py3).

43. Sabbagh, D. et al. Quantum transport properties of industrial $^{28}\text{Si}/^{28}\text{SiO}_2$. *Phys. Rev. Appl.* **12**, 014013 (2019).
44. Lawrie, W. I. L. et al. Quantum dot arrays in silicon and germanium. *Appl. Phys. Lett.* **116**, 080501 (2020).
45. Zinner, E. Depth profiling by secondary ion mass spectrometry. *Scanning* **3**, 57–78 (1980).
46. Srinivasa, V. et al. Simultaneous spin-charge relaxation in double quantum dots. *Phys. Rev. Lett.* **110**, 196803 (2013).

Acknowledgements This research was funded by Intel Corporation. We acknowledge discussions with the members in the Spin Qubit team, the Cryo-CMOS team and Intel Corporation, and technical assistance by O. Benningshof, M. Sarsby, R. Schouten and R. Vermeulen.

Author contributions X.X., B.P. and J.P.G.v.D. performed the experiment. N.S. fabricated the quantum device. A.S., B.P.W. and G.S. designed, grew and characterized the Si/SiGe heterostructure. A.C. contributed to the preparation of the experiment. X.X. and B.P. analysed the data presented in the main manuscript. F. Sebastiano, M.B., S.P., E.C. and L.M.K.V. conceived and supervised the project. X.X., B.P. and L.M.K.V. wrote the manuscript with input from all authors.

Competing interests The authors declare no competing interests.

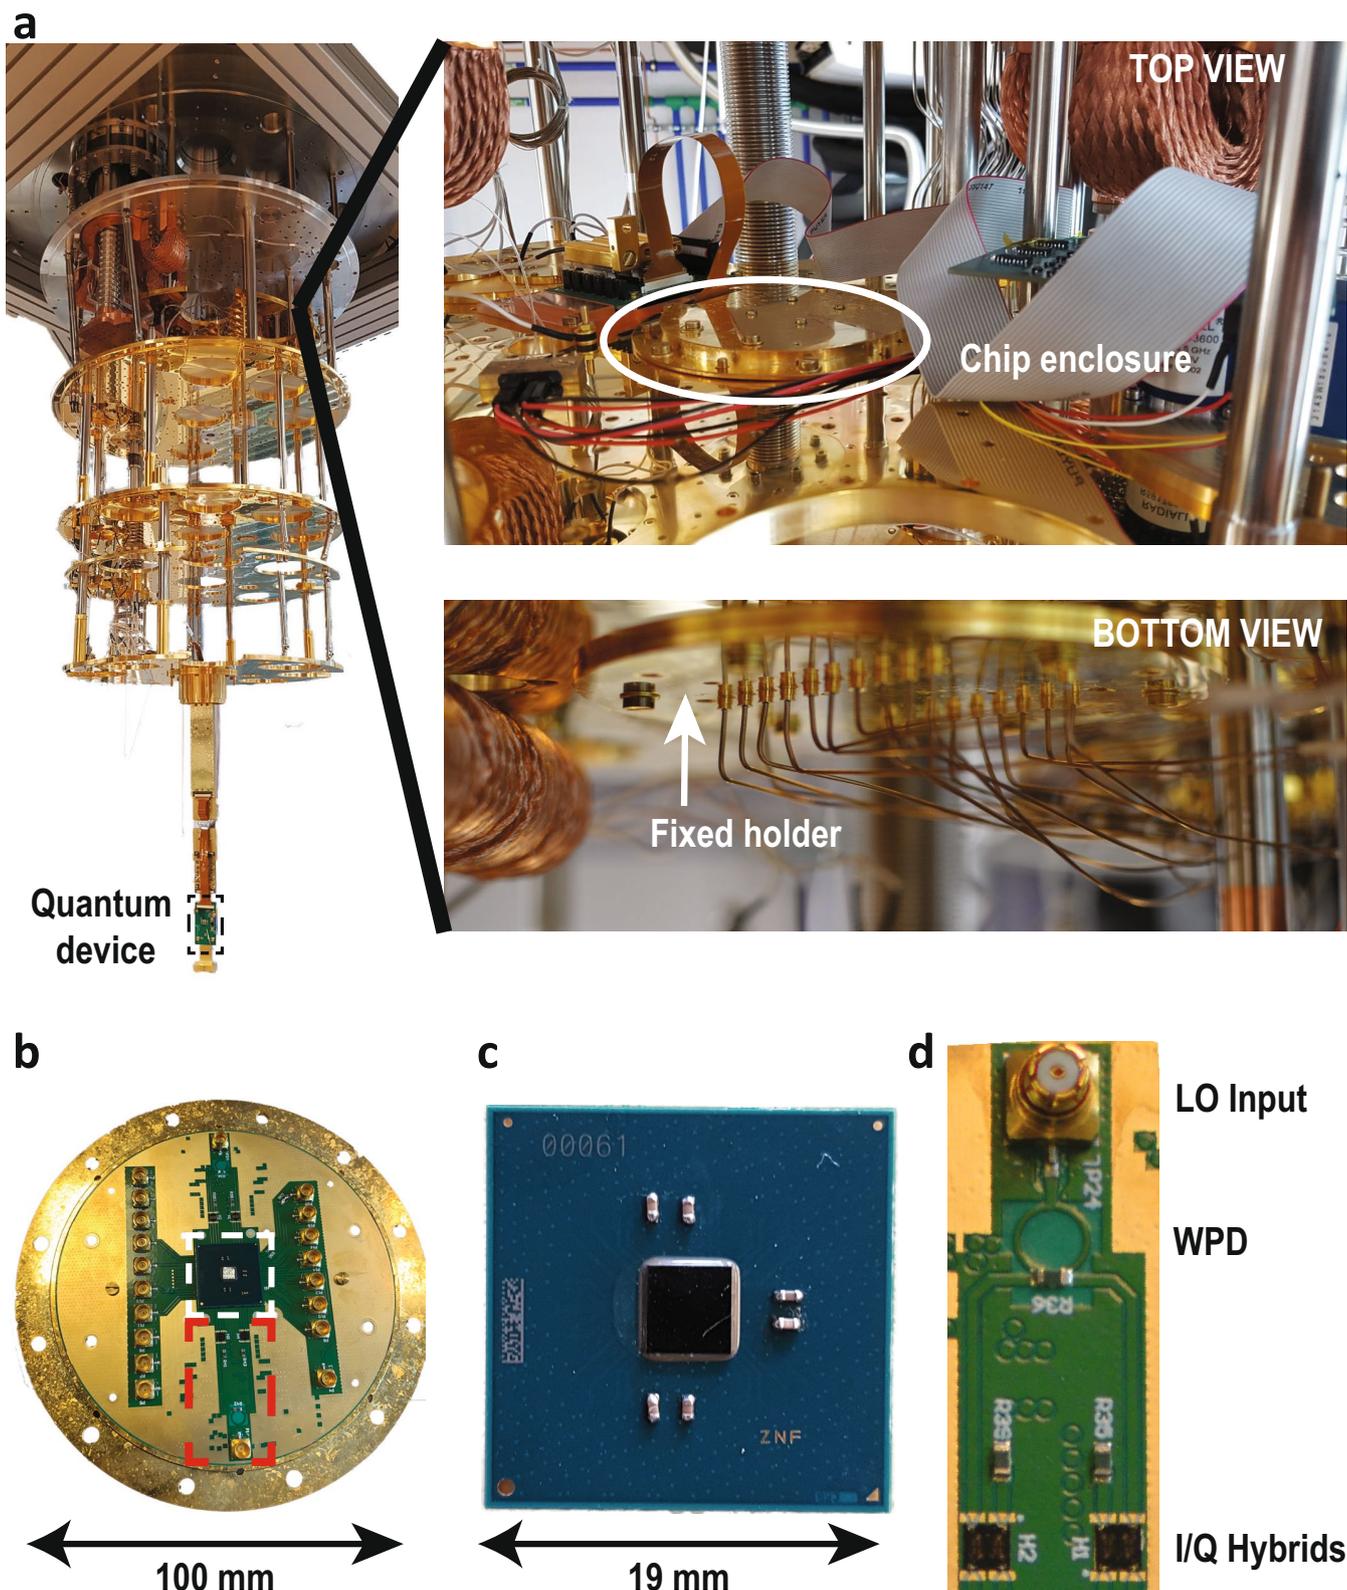
Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41586-021-03469-4>.

Correspondence and requests for materials should be addressed to E.C. or L.M.K.V.

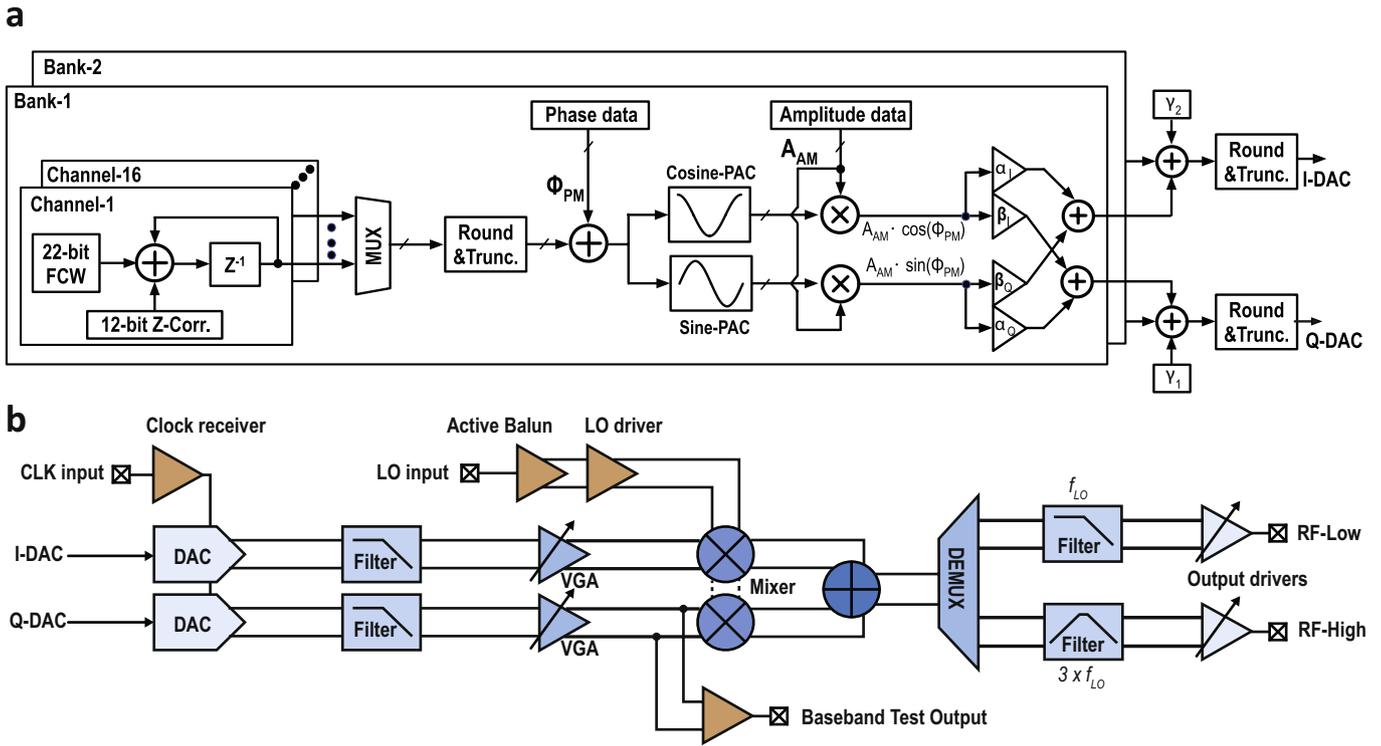
Peer review information Nature thanks Fernando Gonzalez-Zalba, Stefan van Waasen and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Peer reviewer reports are available.

Reprints and permissions information is available at <http://www.nature.com/reprints>.

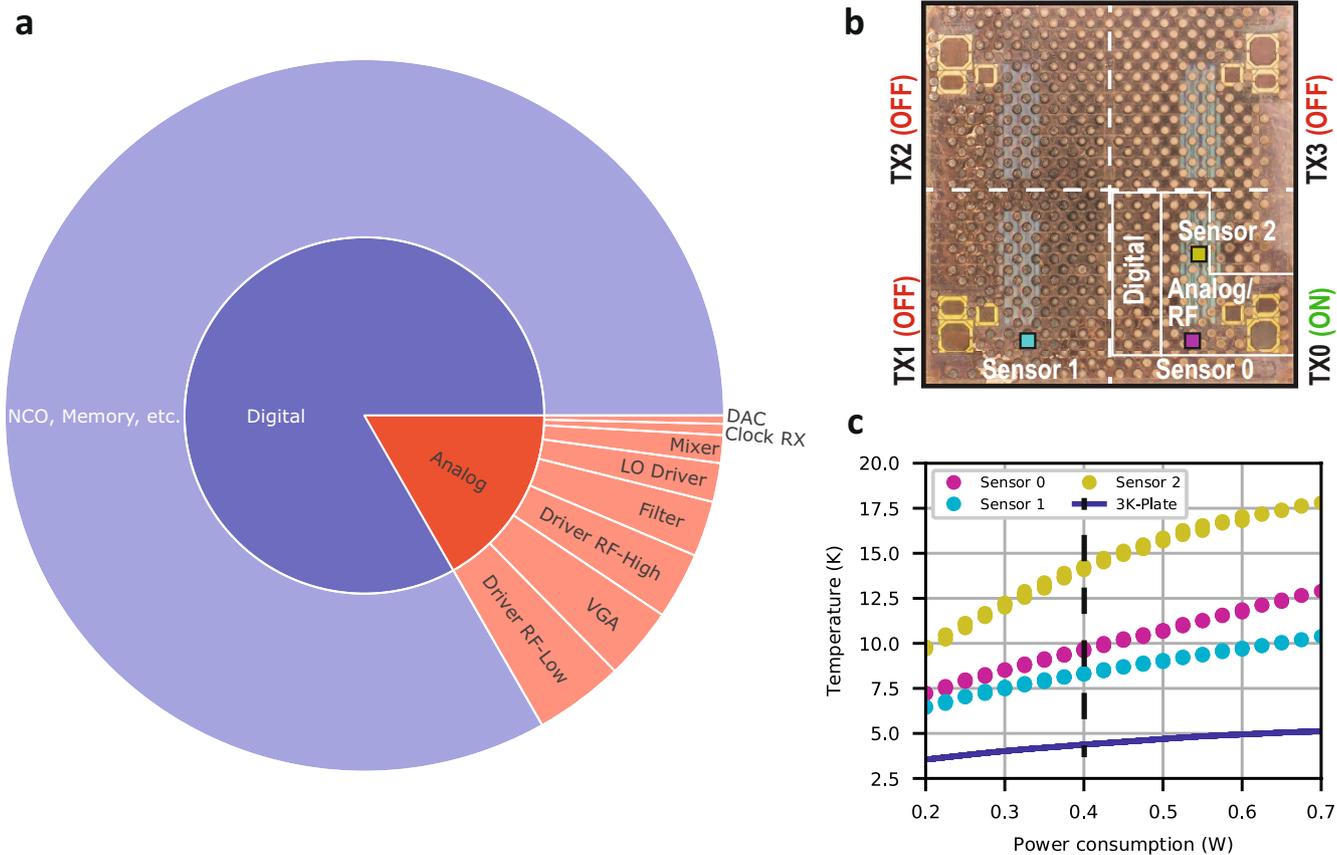


Extended Data Fig. 2 | Dilution refrigerator setup. **a**, Location of the cryo-controller and the quantum device inside the dilution refrigerator (left). Top and bottom views of the 3-K plate, showing the mounted chip enclosure and the fixed holder for the enclosure, respectively (right). **b**, Top view of the gold-plated annealed copper enclosure (without the lid), which is used to mount and thermalize the cryo-controller. **c**, Ball-grid array (BGA 324) package hosting the cryo-controller chip with on-package decoupling capacitors

(shown as a white box in **b**). **d**, The Wilkinson power divider (WPD) splits the input LO power into two equal paths with half power in each, implemented on a PCB. Discrete I/Q hybrids that create the in-phase and quadrature-phase components of the input LO are wire-bonded on the PCB for LO distribution between the different transmitters inside the cryo-controller (shown as a red box in **b**).

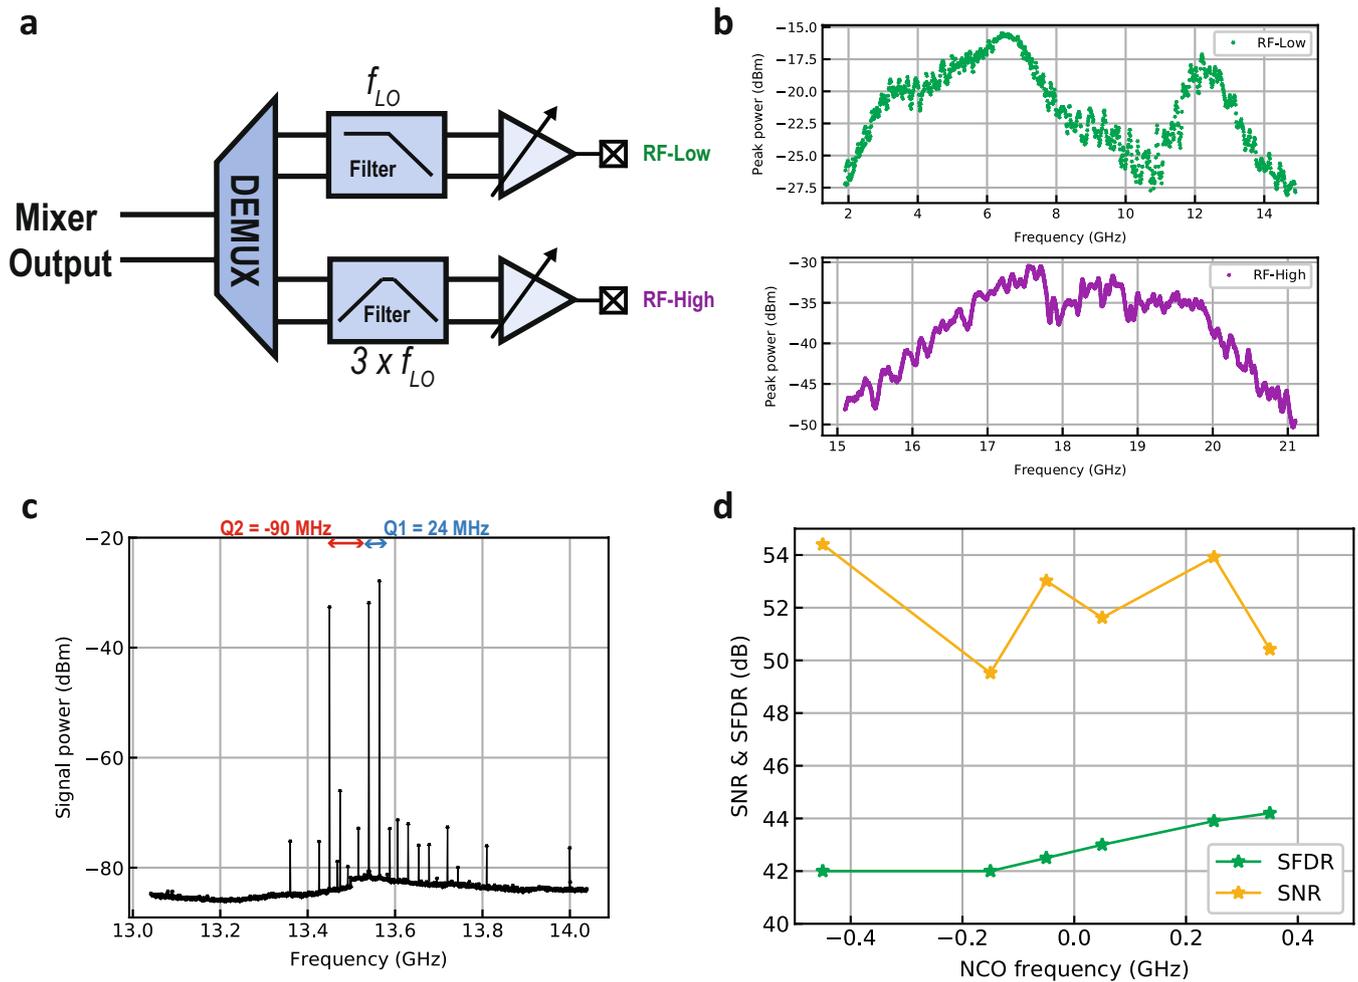


Extended Data Fig. 3 | Detailed cryo-controller schematic. a. Detailed representation of the digital circuitry. **b.** Detailed system-level schematic of the analogue circuitry inside the controller.



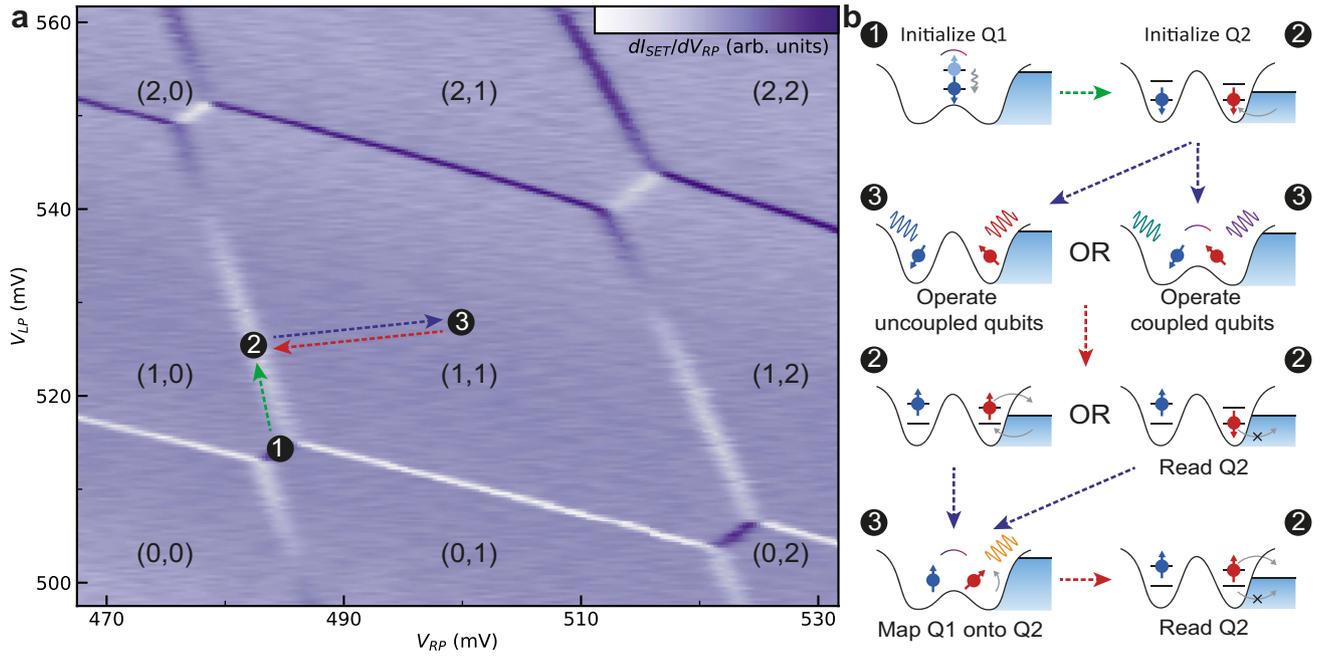
Extended Data Fig. 4 | Power consumption and self-heating of the cryo-controller. **a**, Power-consumption pie chart showing the contribution of the digital and analogue circuits in the cryo-controller. A breakdown of the power consumption of individual circuit blocks is shown for the analogue circuits. The digital circuits use a 0.7-V supply and the analogue circuits use a 1-V supply. **b**, Chip micrograph showing the on-chip bumps used as inputs/outputs. The locations of on-chip temperature-sensing diodes and the

analogue and digital circuitry (in TX0) are highlighted. **c**, On-chip and 3-K plate temperature, measured using different sensors, versus the power consumption of TX0, as reported in ref.¹³. The power consumption is varied by changing the clock frequency of the chip. The nominal operating point for the work presented here and the corresponding temperatures are highlighted with a dashed vertical line. All the other transmitters (TX1, TX2, TX3) are switched off in this measurement.



Extended Data Fig. 5 | Detailed electrical characterization of the cryo-controller. **a**, Schematic of the output driver (complete version in Extended Data Fig. 3) showing the two different RF outputs, which use the same external LO to generate two different frequencies—that is, a 1-GHz band around the LO frequency (f_{LO}) or a 1-GHz band around $3 \times f_{LO}$ —by selecting the ‘RF-Low’ or ‘RF-High’ path, respectively. RF-High uses the third harmonic output of the mixer to generate the tone around $3 \times f_{LO}$. **b**, Peak output power versus

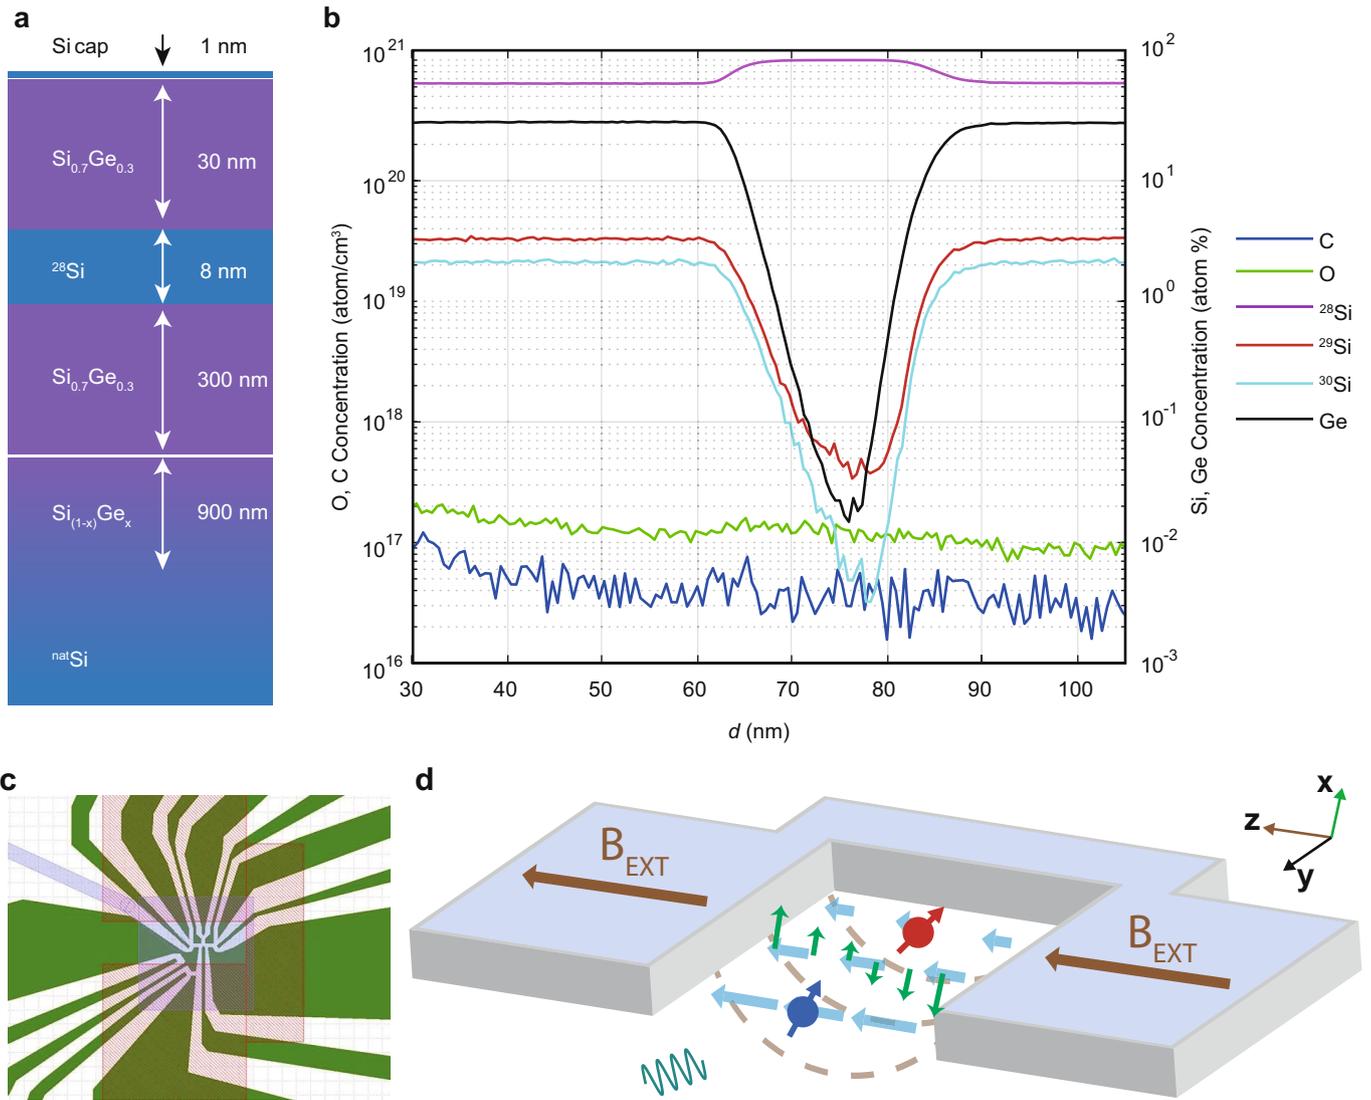
frequency generated using the RF-Low and RF-High path, respectively, as reported in ref. ¹³. The output power can be lowered by up to 40 dB below the peak power in the entire frequency range; for example, at 6 GHz the amplitude range is from -56 dBm to -16 dBm. **c**, Two-tone output spectrum of the cryo-controller used in the simultaneous Rabi oscillation experiment. **d**, SNR and SFDR of the cryo-controller at various NCO frequencies around 13.54 GHz.



Extended Data Fig. 6 | Pulsing scheme used in qubit experiments.

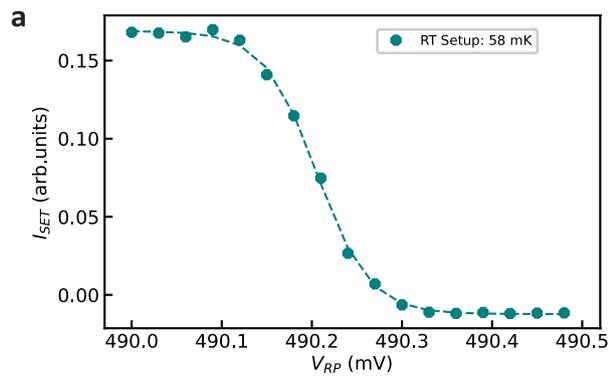
a, Charge-stability diagram of the DQD system, showing the differential current signal (dI_{SET}/dV_{RP}) and charge occupation (M, N) (where M denotes the number of electrons in the dot below LP and N represents the number of electrons in the dot below RP) as a function of the voltages applied to gates LP (V_{LP}) and RP (V_{RP}). The three main stages of a typical pulse sequence are marked by the numbered circles. The gate voltages of stage 3 vary between different experiments: in the experiments with exchange coupling turned on, owing to the cross-capacitance between the barrier (gate T) and the plungers (gates RP and LP), the LP and RP voltages differ from those in the experiments without exchange coupling by -15 mV. **b**, Schematic representations of the DQD system during the experiment cycle. Q_1 is first initialized to its ground state (spin down) via fast relaxation by pulsing to the charge transition line between (1, 0) and (0, 1) (stage 1), which is a spin-relaxation hotspot⁴⁶. Then, Q_2 is initialized by

pulsing it to the transition line between (1, 0) and (1, 1) (stage 2), where the Fermi energy of the electron reservoir is placed in between the two spin states of Q_2 . This allows a spin-down electron to tunnel into the dot but forbids spin-up electrons from tunnelling in—a mechanism called spin-selective tunnelling. During the qubit operations, the system is pulsed to the middle of the (1, 1) region (stage 3), so both electrons are well confined inside the DQD. The barrier (gate T) voltage is used to turn off the exchange coupling between the two spins in the operation of uncoupled qubits (all measurements in Fig. 3) and to turn on the coupling for two-qubit logic operations (all measurements in Fig. 4). After the operations, the Q_2 state is read out via spin-selective tunnelling and reinitialized into the spin-down state (stage 2). The state of Q_1 is read out by mapping its state onto Q_2 via a two-qubit CROT gate (stage 3), followed by readout of Q_2 again (stage 2).

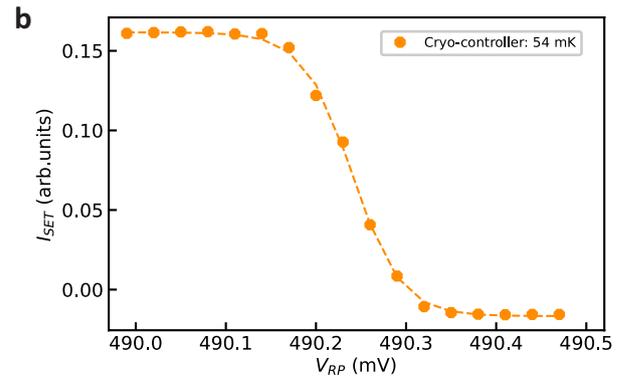


Extended Data Fig. 7 | Magnetic field gradient. **a**, Wafer stack schematic with corresponding layer thicknesses. **b**, Depth concentration SIMS profile of ^{28}Si (red), ^{29}Si (blue), ^{30}Si (purple), Ge (black), oxygen (green) and carbon (blue). The residual ^{29}Si concentration in the quantum well is 0.08%, considerably reducing qubit decoherence due to hyperfine interaction. Both carbon and oxygen concentrations are below their respective detection limits of $3 \times 10^{16} \text{ cm}^{-3}$ and $1 \times 10^{17} \text{ cm}^{-3}$. **c**, Schematic showing the first and second Al gate layers in green and purple, respectively. A cobalt micro-magnet is located on top of the metallic gates (pink-shaded area). **d**, The micromagnet is magnetized by sweeping the external magnetic field (in the \hat{z} direction) from 0 to 3 T and back to 380 mT. The magnetized micro-magnet provides an additional magnetic

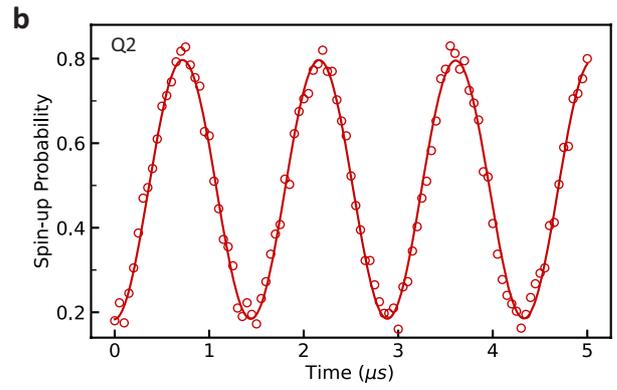
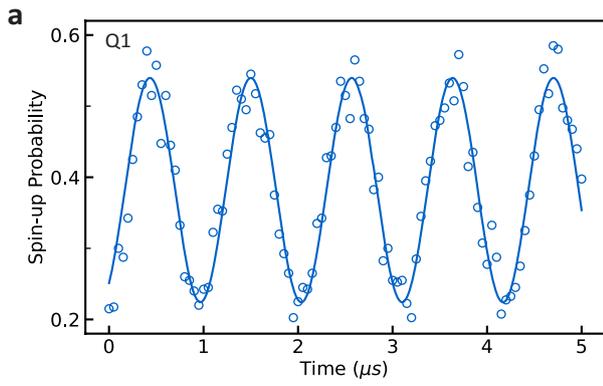
field (brown dashed lines) that has a longitudinal (\hat{z}) component with a field gradient along the double quantum dots. This longitudinal magnetic field gradient (light blue arrows) makes the Zeeman splittings (resonance frequencies) of the two qubits different by -110 MHz. Additionally, the micro-magnet also induces a transverse (\hat{x}) magnetic field gradient (green arrows). When a microwave pulse is sent to the device through gate MW, the wave functions of the electrons are oscillating in the \hat{z} direction. If the microwave frequency is in resonance with the qubit frequency, the electron is subject to an oscillating magnetic field along the \hat{x} direction, which induces electric-dipole spin resonance²⁹.



Extended Data Fig. 8 | Electron temperature measured at different configurations. a, b, SET current signal (I_{SET}) as a function of RP voltage (V_{RP}) measured at the charge transition between (1, 0) and (1, 1) when the quantum device is connected to the VSG (a) and to the cryo-controller (b) (at zero magnetic field). The electron temperatures are extracted by fitting the curves



with the Fermi-Dirac distribution, with a lever arm of 0.172 eV V^{-1} . The measurements indicate that the output noise of the cryo-controller does not affect the electron temperature more than the noise from the RT setup reduced by 6 dB at the 3-K plate.



Extended Data Fig. 9 | Rabi oscillations of qubits individually driven by the cryo-controller. The output frequency of two NCOs are set to the frequencies of Q_1 and Q_2 , respectively, but only one NCO is active each time. Using the same method as described in the main text, the Rabi oscillations of Q_1 (**a**) and Q_2 (**b**)

are measured individually. Compared to the simultaneous Rabi oscillations shown in Fig. 3b, the decay is much slower in the individual driving experiments.