

Research Highlights

The past year saw over 100 papers published or accepted for publication in peer-reviewed journals, across all of the Centre's eleven experimental and seven theoretical research programs. The next few pages bring together the key research highlights for 2007 across the Centre.

Silicon-based Si:P Qubits: Top-Down Fabrication and Measurement

The development of spin-based qubits in the Si:P material system is the key focus of the Centre's solid-state programs. At UNSW the Integrated Quantum Computer Devices program and the Quantum Measurement program interact closely to design, construct and measure silicon-based qubits and qubit test devices. During 2006, in close collaboration with the Ion Beam program at the University of Melbourne, they configured (using the Centre's ion-counting technology – see Jamieson et al., *Appl. Phys. Lett.* **86**, 202101 (2005)) and measured (using rf-SET technology – see Buehler et al., *J. Appl. Phys.* **96**, 4508 (2004)) the time-resolved control and detection of single-electron transfers in a silicon device implanted with *exactly two phosphorus donors* – see Andresen et al., *Nano Letters* **7**, 2000 (2007).

The next milestone for our silicon qubit program is the coherent control and readout of the *electron spin* of a single P donor in silicon. To this end, we plan to implement two methods that have been successfully demonstrated in GaAs quantum dots: spin-to-charge conversion for single-shot readout of the electron spin and *local* electron spin resonance for coherent spin manipulation. Both these methods rely upon technology that, to a large extent, has already been demonstrated in Si:P devices.

Underpinning our future qubit structures is a novel MOS-based silicon quantum dot structure with tunable barrier gates – see Angus et al., *Nano Letters* **7**, 2051 (2007). An SEM image of a typical device is shown in Figure 1(a). Two Al barrier gates (~ 20 nm wide) are fabricated on an ultra-high-quality SiO₂ gate oxide, above a near-intrinsic silicon (i-Si) substrate. The barrier gates are oxidized and covered with a top gate that is positively biased to induce an electron accumulation layer underneath the SiO₂. By further lowering the potential of the barrier gates, we isolate the small region of electrons between the gates and obtain a quantum dot, whose coupling to the leads can be easily tuned. By appropriate biasing of the gates it is possible to tune the dot into the few electron limit. Figure 1(e) shows bias spectroscopy data on a dot with only ~ 10 electrons, exhibiting clear evidence of discrete energy levels within the quantum dot.

Besides the quantum dot physics in itself, such a dot has a large charging energy and can be used as a silicon radio-frequency single electron transistor (rf-SET) when connected to a resonant LC tank circuit – see Figures 1(c,d). In this mode, with a not yet optimized Si-SET, we achieved a charge sensitivity of order $10 \mu\text{e}/\text{Hz}$, which equals or surpasses that of typical aluminum SETs – see Angus et al., *Appl. Phys. Lett.* **92**, 112103 (2008). This excellent sensitivity enables measurement of a charge transfer equivalent to 1% of an electron, with a measurement time of $\sim 1 \mu\text{s}$. The important advantage of the Si-SET, as compared to the more common Al-SETs, is that its fabrication is MOS-compatible and does not require complex double-angle evaporation. The absence of the thin (and non-tunable) Al₂O₃ tunnel barriers is an additional bonus, reducing the vulnerability of the SET to electrostatic discharge.

Coherent manipulation of a spin qubit requires the application of microwave fields, with a frequency matching the Zeeman splitting of the electron spin states. During 2007 our team at UNSW demonstrated the ability to perform *local* ESR on a small number of donors by fabricating an MOS structure where the top gate has the double function of: (i) inducing a 2DEG in the MOSFET channel; and (ii) supplying the microwave field – see Figure 2(a). For the latter purpose, the gate is shaped as a coplanar transmission line, terminated by a short circuit. This arrangement yields maximum magnetic field in the channel of the MOSFET, and *zero (rf) electric field*. Because there is no resonant structure, this ESR line can be used over a very wide frequency range (10 MHz–50 GHz). To verify the effectiveness of the local ESR line we performed an electrically-detected magnetic resonance (EDMR) experiment, observing all expected features of the ESR of P donors, and (for the first time) extended this type of investigation to the milli-Kelvin regime – see Figure 2(b).

In late-2007 we conceived a proposal for the coherent control and readout of a Si:P single spin qubit device that combines our new Si-SET technology and local ESR capability – see Figure 3. The essential features of this structure are: (i) a silicon rf-SET for single-shot detection of electron transfers between an individually implanted P donor and a reservoir; (ii) a local ESR line running above the dot, to apply both an electrostatic potential and a microwave magnetic field; and (iii) use of the SET island as the gate-induced electron reservoir, by positioning the implanted P donor close enough to allow (spin-dependent) electron tunneling. Besides its compactness, the proposed structure is innovative in the idea of using the SET island as an electron reservoir. The electrostatic potential on the ESR line is used to shift the Zeeman-split electron spin levels with respect to the Fermi level of the reservoir.

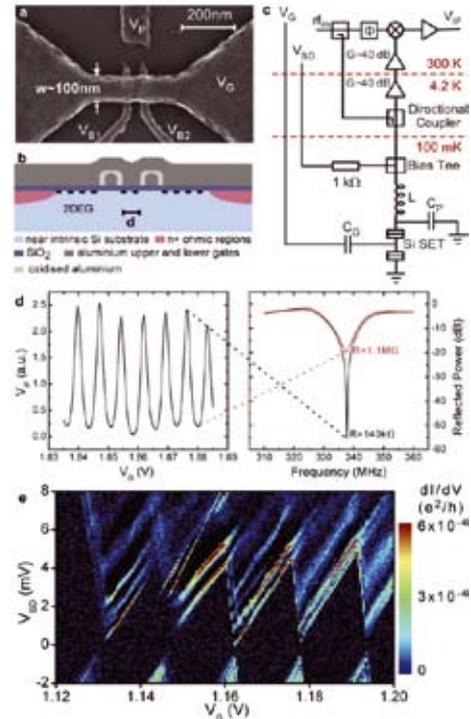


FIGURE 1 All-silicon radio-frequency single electron transistor (Si rf-SET). (a) Scanning electron micrograph of a typical device. (b) Schematic cross section of a device, illustrating the 2DEG induced in the near-intrinsic silicon by the upper MOSFET gate, and locally depleted by the lower barrier gates. (c) Schematic of rf-measurement setup. (d) When inserting the device in a tank circuit, the Coulomb blockade peaks (left) are obtained from the modulation of the reflected microwave power (right) when changing the potential of the top gate. The charge sensitivity achieved by biasing the device on the steepest slope of the Coulomb peaks is better than $10 \mu\text{e}/\text{Hz}$. (e) DC bias spectroscopy on a similar device, with gate width $w=50$ nm, biased to a regime with only ~ 10 electrons in the dot. Evidence of excited states in the dot are clearly visible.

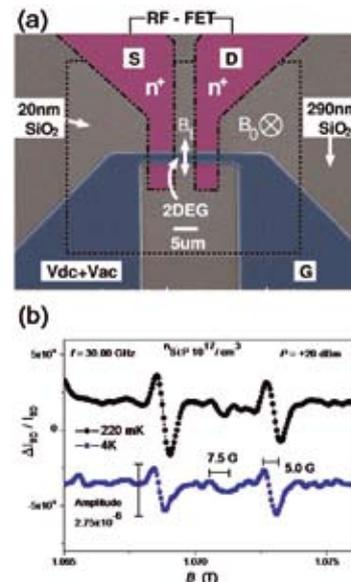


FIGURE 2 (a) SEM image (in false colors) of an EDMR device designed to probe P donors in the Si:P system. (b) Resonant microwave absorption of P donors as detected by EDMR, at $T = 4$ K (blue) and 220 mK (black), with excitation frequency $f = 30$ GHz. Notice the ability to resolve the hyperfine-split resonance peaks due to the interaction between the electron spin and the ^{31}P nucleus.

Spin-to-charge conversion is possible when the reservoir Fermi level is halfway between the ground (spin-up) and excited (spin-down) electron spin states – see Figure 3(b). Single-shot detection of charge transfer is therefore equivalent to projective measurement of the qubit state. Since the electron tunnels directly onto the SET island, the magnitude of the charge transfer signal is expected to be very large, $\sim 0.1e$, easily detectable in single shot by our Si-SETs. The coherent manipulation of the qubit will be achieved by applying controlled microwave pulses to the ESR line, which induce Rabi oscillations of the spin. A range of devices based upon this new architecture are currently in production at UNSW and UM.

Of central importance to our ion-implanted Si:P qubit devices is the quality of the Si-SiO₂ interface, formed during thermal oxidation. Our Materials program at UM works closely with the fabrication team at UNSW to assess the quality of our Si-SiO₂ interfaces, using capacitance-voltage (CV) measurements and deep-level transient spectroscopy (DLTS). Figure 4 illustrates the changes in interface-trap density (D_{it}) for oxides grown with and without an (undesired) source of residual water vapour in the oxide growth furnace. The red curve shows the trap density versus defect energy-level for an oxide grown with a low level of water vapour. The trap density near the conduction-band edge ranged from $1-3 \times 10^{11} \text{ eV}^{-1}\text{cm}^{-2}$, a figure unacceptable for Si:P qubits. The blue curve shows the interface trap density obtained after the source of water vapour contamination had been removed. Now, the trap density near the band edge is an order of magnitude better at $1.0 \times 10^{10} \text{ eV}^{-1}\text{cm}^{-2}$. This exciting result is the lowest trap density we have achieved thus far and is expected to be suitable for successful fabrication of qubit devices.

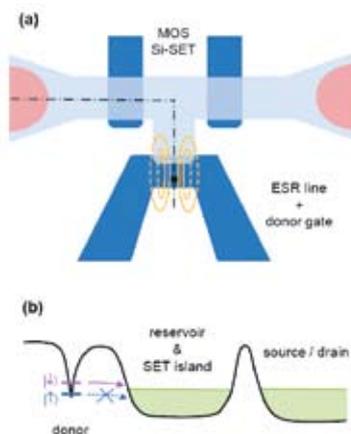


FIGURE 3

(a) Conceptual sketch of complete device for control and readout of the electron spin of a single P donor. In this design, the Si-SET has an extended island that protrudes towards the donor site to enhance spin-dependent tunnelling. (b) The energy landscape along the dash-dotted line shown in (a).

Beyond the few-qubit level, solid-state QCs will require classical control-measure electronics for qubit initialization, control and readout. In 2006 our Control Chip team at UNSW designed, fabricated (with Peregrine Semiconductor Australia) and demonstrated (at 300 K, 4 K and sub-1 K) a CMOS silicon-on-sapphire (SOS) voltage-pulse generator for integration with our Si:P qubit chips in a cryogenic environment. During 2007 we developed and demonstrated a second-generation pulse-generator chip (see Figure 5) using digital pulse timing, offering much lower signal jitter and improved reliability.

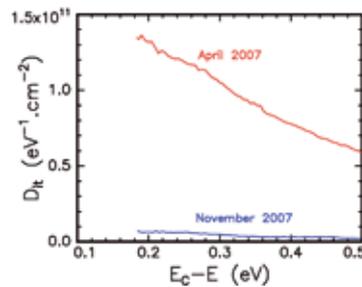


FIGURE 4

DLTS measurements of our thermally-grown gate oxide now show Si-SiO₂ interface-trap densities of order $10^{10} \text{ eV}^{-1}\text{cm}^{-2}$ near the band edge and approaching $10^9 \text{ eV}^{-1}\text{cm}^{-2}$ in the mid-gap region. These trap densities are comparable with the lowest reported. The data shows a typical result (blue curve) under optimal conditions, together with higher trap densities (red curve) obtained when our growth furnace was exhibiting water-vapour contamination that has now been removed.

Silicon-based Si:P Qubits: Scanning Probe Fabrication and Modelling

The Atomic-Scale Fabrication and Crystal Growth program at UNSW has the ultimate goal of fabricating the Si:P qubit architecture with atomic precision using a combination of scanning tunneling microscopy (STM) and molecular beam epitaxy (MBE). Over the past 5 years they have developed a unique fabrication strategy to fabricate buried P-doped nano- to atomic-scale devices in silicon using the atomic placement accuracy of the STM. In 2007 the main focus has been

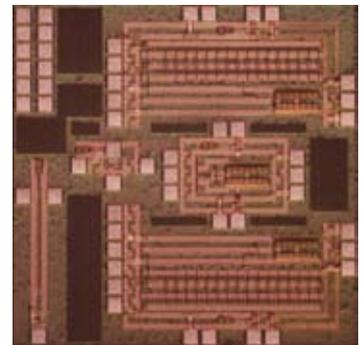


FIGURE 5

Optical micrograph of the Centre's second-generation voltage pulse generator for (classical) solid-state qubit control. The chip was designed and tested at UNSW and fabricated using rf-CMOS silicon-on-sapphire technology at Peregrine Semiconductor Australia in Sydney.

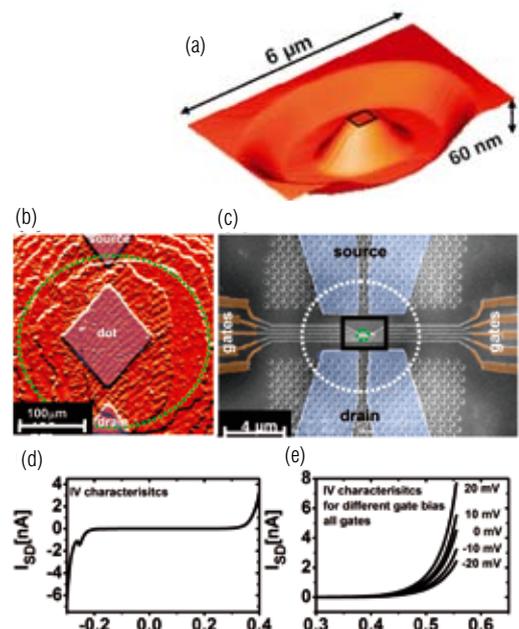


FIGURE 6

STM-fabricated Si:P dot with EBL-aligned metal surface gates. (a) Large, atomically-flat terrace, produced using Si step engineering. (b) image of a 100 nm^2 dot patterned between source-drain leads on an atomically flat terrace. (c) Alignment of 9 gates across the top of the dot. (d) I-V characteristics across the dot, showing a large conductance gap despite the small $\sim 20 \text{ nm}$ separations between source/drain and dot. (e) I-V characteristics across the dot for different gates voltages (all gates at the same potential).

to demonstrate controlled charge transfer in atomically controlled Si:P nanodots patterned between lateral source-drain leads. This is a major step towards mapping the energy levels and spin structure of P dopants in silicon.

To date, our STM-based devices have been electrically contacted using Al contacts aligned above the buried dopants using etched registration markers defined in the Si by optical lithography. To improve alignment of contacts and gates we have transferred all *ex-situ* fabrication steps from optical to electron-beam lithography, thereby increasing alignment accuracy from 500 nm to below 100 nm. A key component of this scheme is the use of Si step engineering to create atomically-flat plateaus for accurate device positioning on a single atomic Si plane – see Figure 6(a). Our largest terraces have measured up to $\sim 1 \mu\text{m}$ across and we can reliably achieve step-free regions several-hundred nm in diameter. Figure 6(b) shows a Si:P dot patterned by STM between source-drain leads in a four-terminal configuration before ten metal gates (width $\sim 50 \text{ nm}$) were aligned above the dot – see Figure 6(c). I-V characteristics across the dot (Figure 6d) demonstrate a large non-conductive gap with the appearance of resonances at the conductance onset. By applying voltages to all gates simultaneously we can alter the conductance across the gap, demonstrating clear gating action – see Figure 3(e).

An advantage of STM lithography is that it allows patterning of gates in the plane of the P dopants, providing higher alignment accuracy than surface gates. To achieve this we must fabricate narrow P-doped leads in close proximity to the dot. Over the past few years we have demonstrated the ability to fabricate STM-defined planar, highly P-doped wires down to 2.5 nm in width which still show ohmic behaviour – see Ruess et al., Phys. Rev. B **76**, 085403 (2007). The observation that ohmic behaviour is maintained at such narrow widths with such high conductances can be attributed to a combination of factors, including the use of highly-doped planar layers with atomically-sharp dopant profiles, high structural integrity combined with a separation of the wire from surface and interface states. These represent the lowest resistivity Si:P nanowires realised to date. We have recently integrated such wires as in-plane gates controlling a Si:P nanodot – see Figure 7(a). This method of gating provides very precise alignment – an important factor as devices reach the single atom level. Operating this as a three-terminal device shows non-linear electron transport across the dot (see Figure 7b), with the I-V characteristics showing a larger energy gap for the largest electron path, as expected. By measuring the I-V characteristics between terminals 1 and 3, we can use terminal 2 to gate the device, as shown in Figure 7(c).

Optical Qubits: Experiment and Theory

The Centre's optical qubit programs aim to construct the basic building blocks of an optical QC and develop foundations of a scalable architecture based on the protocols for linear optical quantum computing (LOQC) introduced by Knill, Laflamme and Milburn.

Our experimental team at UQ is focused on development of optical qubits in the polarization basis. In 2003 they constructed and observed unambiguous quantum operation of a CNOT gate, which was then fully characterized using quantum process tomography in 2004. During 2007, the highlight from the UQ team was their implementation of a compiled version of Shor's algorithm in a photonic system, demonstrating for the first time the core processes, coherent control, and resultant entangled states required in a full-scale implementation – see Lanyon et al., Phys. Rev. Lett. **99**, 250505 (2007).

Figure 8 shows the measured density matrices of the argument-register output for Shor's algorithm. Ideally these are maximally-mixed states: in all cases we measure near-unity fidelities. The

output of the routines are the logical state probabilities, i.e. the diagonal elements of the matrices. Combining these with the known state of the redundant qubit, and reversing the argument qubits as required, gives the binary outputs of the algorithm which after classical processing yields the prime factors of $N=15$. In the order-2 circuits the outputs of the algorithm are 00 or 10: the former represents the expected failure mode, the latter a successful determination of the order $r=2$; failure and success should have equal probabilities, we measure them to be 50% to within error. Thus half the time the algorithm yields $r=2$, which gives the factors, 3 and 5. In the order-4 circuit the algorithm finds $r=4$, again yielding the factors 3 and 5.

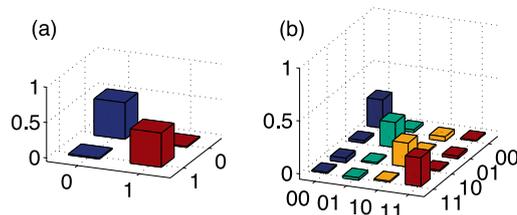


FIGURE 8

Shor's algorithm outputs given by measured argument-register density matrices. The diagonal elements are the logical output probabilities. (a) Order-2 algorithm. The fidelity with the ideal state is $F=99.9 \pm 0.3\%$, the linear entropy is $S_L=100 \pm 1\%$. Combined with the redundant qubit the logical probabilities are $\{P_{00}, P_{10}\} = \{52, 48\} \pm 3\%$. (b) Order-4 algorithm, $F=98.5 \pm 0.6\%$, and $S_L=98.1 \pm 0.8\%$. The logical probabilities are $\{P_{000}, P_{010}, P_{100}, P_{110}\} = \{27, 23, 24, 27\} \pm 2\%$. Real parts shown, imaginary parts are less than 0.6%.

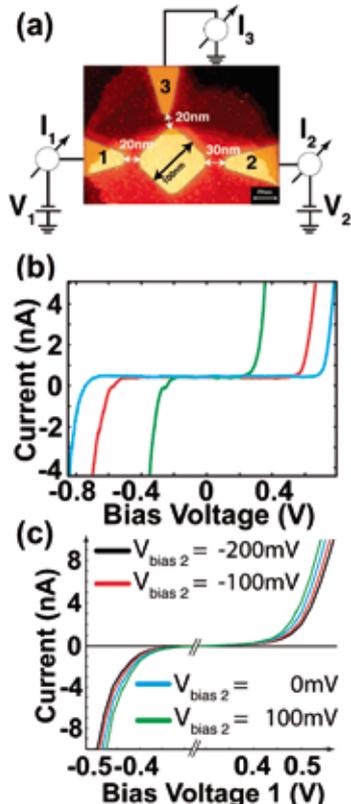


FIGURE 7

STM-fabricated 3-terminal dot device, where one terminal is used to gate the dot between source-drain leads. (a) STM image of the device. (b) Current-voltage characteristics across the device for each different configuration. (c) Current-voltage characteristics across the dot from terminals 1–3 whilst applying $\pm 200 \text{ mV}$ to gate 2.

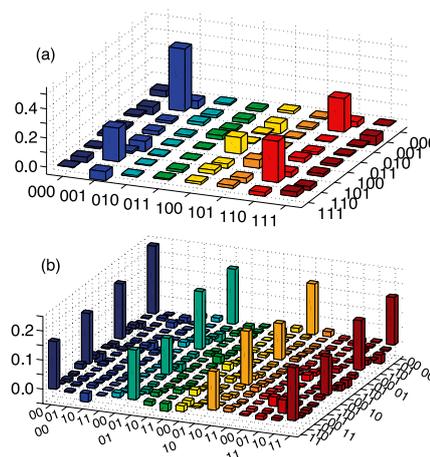


FIGURE 9

Measured density matrices of the state of both registers after modular exponentiation in our implementation of Shor's algorithm. (a) Order-2 circuit. The ideal state is locally equivalent to a GHZ state: we find $F_{\text{GHZ}}=59 \pm 4\%$. The state is partially-mixed, $S_L=62 \pm 4\%$, and entangled, violating the optimal GHZ entanglement witness $W_{\text{GHZ}}=1/2-F_{\text{GHZ}}=-9 \pm 4\%$. (b) Order-4 circuit. Measured fidelity with the ideal state, a tensor product of two Bell-states, is $F=68 \pm 3\%$. The state is partially-mixed, $S_L=52 \pm 4\%$, and entangled, with tangles of the component Bell-States of $41 \pm 5\%$ and $33 \pm 5\%$. Real parts shown, imaginary parts are respectively less than 7% and 4%. The fidelity of the four-qubit state, (b), is higher than the three-qubit state, (a), chiefly because the latter requires nonclassical interference of photons from independent sources, which suffer higher distinguishability, lowering gate performance.

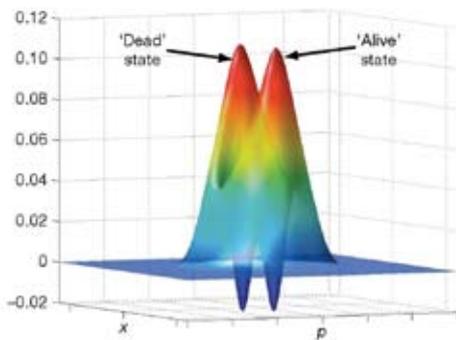


FIGURE 10

Quasi-probability distribution (Wigner Function) of the cat state produced in collaboration with Institut d'Optique. The negative regions of the distribution indicate the strong non-classical nature of the state [A.Ourjoutsev, H.Jeong, R.Tuaille-Brouri & P.Grangier, Nature **448**, 784 (2007)].

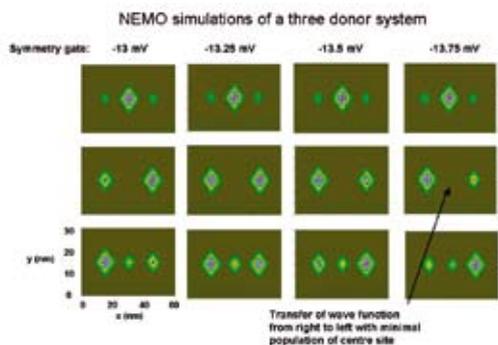


FIGURE 11

Lowest eigenstates (probability density plotted) of a three donor system, calculated using NEMO, as a function of surface gate bias. To facilitate NEMO simulations, donors were placed 15 nm apart in a device volume corresponding to some 3 million atoms. In the large space of gate configurations the bias conditions were found corresponding to the required CTAP adiabatic pathway (state 2 sequence in plot).

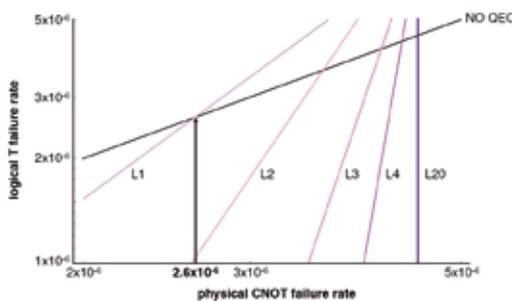


FIGURE 12

Error thresholds (lower bounds) for the encoded T-gate gate at various levels of concatenation, showing pseudo-thresholds (Level 1 = 2.6×10^{-6}) and asymptotic threshold estimate (Level $\infty = 4.5 \times 10^{-6}$) for a conservative set of inputs for physical times and error rates for gates, transport and measurement.

These results indicate near-ideal algorithm performance, *far better* than we have any right to expect given the known errors inherent in the logic gates. This highlights that the *algorithm* performance is not always an accurate indicator of *circuit* performance since the algorithm produces mixed states. In the absence of the gates the argument-register qubits would remain pure; as they are mixed they have become entangled to *something* outside the argument-register. From algorithm performance we cannot distinguish between desired mixture arising from entanglement with the function-register, and undesired mixture due to environmental decoherence. Circuit performance is crucial if it is to be incorporated as a sub-routine in a larger algorithm. The *joint* state of both registers after modular exponentiation indicates circuit performance; we find entangled states that partially overlap with the expected states (see Figure 9) indicating some environmental decoherence.

Our team at the Australian Defence Force Academy is focused on development of optical qubits in the frequency basis. Currently there are two methods used for the detection of optical phase, homodyne and heterodyne or dual homodyne. Homodyne is good, but only if the phase of the signal is known to within a quarter of a wavelength, which is not always practical, in which case one must use heterodyne or dual homodyne detection. Our experimental setup uses an Ornstein-Uhlenbeck process (OUP) to vary the optical phase of 100 MHz side bands on our coherent state. The results indicate that the adaptive phase estimate is well correlated to the input OUP. We find that the dual homodyne phase measurements are not as well correlated.

Our Optical Quantum Information program at Griffith aims to develop experimental quantum optical tools for quantum information tasks and QC. This involves development of novel measurement techniques, preparation of new types of quantum states, and implementation of alternative methods of characterizing quantum systems. A major 2007 highlight was the first demonstration of an experimentally achievable method of performing a Heisenberg-limited phase estimation, based on a quantum computer algorithm – see Higgins, Berry, Bartlett, Wiseman and Pryde, Nature **450**, 393–6 (2007). This work is discussed further in the Quantum Computing Theory section which follows.

Our program in Linear Optical Quantum Computing (LOQC) Theory at UQ addresses a broad range of issues associated with optical QC from close collaborations on experimental demonstrations to alternative architectures and fundamental issues of scaling. A quite different version of the LOQC paradigm involves encoding the quantum information in multi-photon

coherent states, rather than single photon states – see Ralph et al., Phys. Rev. A **68**, 042319 (2003). This is commonly referred to as Coherent State Quantum Computing (CSQC). A major issue for CSQC is how to produce the necessary resource cat states. This year we developed and implemented a scheme for the creation of cat states of the required size in collaboration with the Institut d'Optique, France. The idea employs post-selection performed on photon number states using homodyne detection – see Lance et al., Phys. Rev. A **73**, 041801(R) (2006). The resulting “squeezed” cat states had an effective coherent amplitude of 1.6, easily the largest yet produced. The protocol straightforwardly generalizes to cat states of higher amplitude. Figure 10 shows the quasi-probability distribution (Wigner function) of the experimental states – see Ourjoutsev et al., Nature **448**, 784 (2007).

Quantum Computing Theory Programs

Researchers in the Device Modelling and Algorithms Program at UM made important advances in 2007 at extreme ends of QC design. At the donor level, large-scale simulations of a 3-donor CTAP device were carried out by Purdue student Rajib Rahman in collaboration with Centre researchers at UM. In this calculation three P donors were placed 15 nm apart and 15 nm below a Si-SiO₂ interface, with four gates controlling donor levels and tunnelling rates. The triple-well was populated by a single electron. TCAD simulations of the potential landscape for a set of gate biases were computed (at UNSW) and coupled to a NEMO tight-binding quantum calculation of the energy levels. A large device volume comprising over 3,000,000 atoms ensured numerical convergence. The aim was to find the gate-bias control pathway corresponding to the CTAP protocol, as a guide to fabrication and measurement of 3-donor CTAP devices. As seen in Figure 11, evidence of left-right tunnelling along the CTAP path is quite clear. These results indicate that the CTAP protocol within a 3-donor system controlled by surface gates is a viable mechanism for spin-coherent transport in a Si:P quantum computer.

At the other end of QC design, we completed threshold calculations of both the logical T and CNOT gates required for universal fault-tolerant QC. As for the logical CNOT-gate threshold calculations reported in 2006, a 21-qubit T-gate schedule, including DiVincenzo-Aliferis quantum error correction (QEC) protocol over the 7-qubit Steane code, was compiled on the Si:P bi-linear qubit array. In the working approximation that all transports are successful (but may result in spin errors) an extended rectangle analysis was carried out to determine the failure rate at each level of concatenation for a given set of inputs. The results (Figure 12)

show the pseudo-threshold error at Level 1 of 2.6×10^{-6} rising to an asymptotic threshold estimate of 4.5×10^{-6} (infinite concatenation). These calculations set the scene for the next challenges in QC architecture design – achieving convergence between required threshold and quantum component precision.

The Centre’s program in Measurement and Control at Griffith studies devices operating on a quantum scale, with applications to QC. In 2007 they continued to investigate feedback control of quantum systems for state preparation. Previously they developed two protocols for rapid-state purification of qubits – see Jacobs, *Phys. Rev. A* **67**, 030301(R) (2003); and Wiseman and Ralph, *New J. Phys.* **8**, 90 (2006). During 2007 they considered the practicalities of such schemes, within the constraints present in a superconducting charge-qubit system – see Figure 13 and Griffith et al., *Phys. Rev. B* **75**, 014511 (2007).

Another 2007 highlight was publication in *Nature* of the first proposal for an experimentally achievable method of performing a Heisenberg-limited phase estimation, based on a QC algorithm. The theory (developed by Berry, Bartlett and Wiseman) was published together with the experimental demonstration – see Higgins et al., *Nature* **450**, 393–6 (2007). Measurement of optical phase is used in length metrology and myriad other applications. At the fundamental level, measurement precision is limited by the number N of quantum resources (e.g. photons) used. Standard measurement schemes, using each resource independently, lead to a phase uncertainty scaling as $1/\sqrt{N}$ – the standard quantum limit (SQL). However, it has long been conjectured that one should be able to achieve a precision limited only by the Heisenberg uncertainty principle, dramatically improving the scaling to $1/N$. We developed the first achievable Heisenberg-limited phase estimation procedure. The idea is to replace entangled input states with multiple applications of the phase shift on unentangled single-photon states. We generalized Kitaev’s phase estimation algorithm, using adaptive measurement theory, to achieve a standard deviation scaling at the Heisenberg limit – see Figure 14. This work was reviewed in two high-profile articles: Jonathan Dowling’s “Kittens catch phase”, in *Nature* **450**, 362 (2007); and Jeremy O’Brien’s “Precision Without Entanglement”, in *Science* **318**, 1393 (2007).

The program in Quantum Information Theory at UQ is directed towards general aspects of quantum information, such as entanglement and error correction, with particular emphasis on condensed matter systems and quantum optics. During 2007 Gerard Milburn, in collaboration with Bill Munro of Hewlett Packard (Bristol) and Kae Nemoto of NII (Tokyo) investigated a doubly-resonant optical cavity containing a Kerr nonlinear medium that couples two modes by a cross-phase modulation. One of these modes is driven by a single photon pulsed field, and the other by a coherent state. They found a new method to study such non-stationary problems that will have wide application in optical QC.

The Quantum Algorithms program at Macquarie aims to theoretically develop new ways to process information using quantum systems. In the practical operation of QCs, executing quantum gates in a time-optimal fashion is the most preferred option. Shorter gate times lead to little degradation due to decoherence and also shorten the overall run-time of the entire algorithm. The study of time-optimal quantum control is a rapidly growing area and researchers have developed a variety of methods to analyse these problems. Centre researcher Alberto Carlini and colleagues have developed a unique analytical method of re-casting the problem into a variational setting, and devised an action, when varied subject to whatever constraints one chooses to make, yields the time-optimal method of generating a target unitary operation – see Carlini et al., *Phys. Rev. A* **75**, 042308 (2007). He has expanded this formalism to encompass the effects of quantum noise and surprisingly finds that certain operations can be more quickly executed if one can control the noise and steer the quantum state through the interior of the Bloch sphere to its destination – see Figure 15.

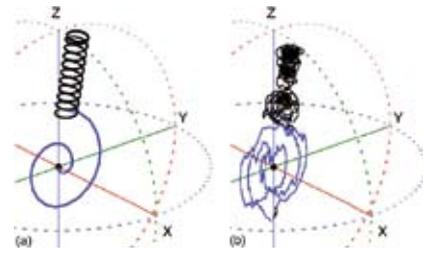


FIGURE 13 Path of the Bloch vector of a CPB for controls designed to approximate Wiseman and Ralph’s rapid purification protocol. (a) Intended path: the feedback control is initially off, so that the Bloch vector grows and rotates and around the x -axis. Once the Bloch vector has exceeded a threshold, a high frequency z -rotation is applied which locks the Bloch vector to the measurement axis. (b) An example trajectory showing the effect of measurement noise. The overall shape of the stochastic trajectory is similar to the conceptual path of (a).

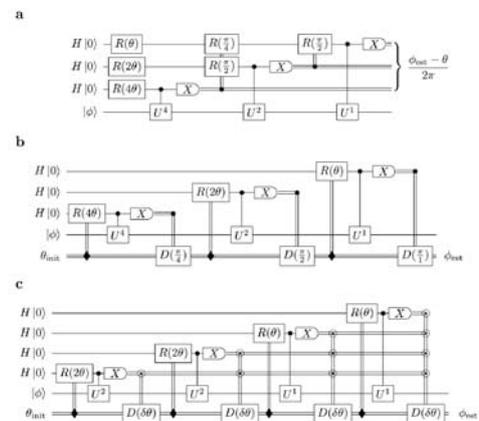


FIGURE 14 Quantum circuit diagrams of Kitaev’s phase estimation algorithm and our generalization. (a) Kitaev’s algorithm with the inverse quantum Fourier transform implemented with measurement and classical feedback and a random initial phase estimate θ . In general, $K+1$ qubits yield $K+1$ binary digits of precision; here $K = 2$. (b) As in (a), but here we implement θ (now called θ_{ini}) and the feedback operations by coupling the qubits to a common element, the “feedback phase” (the lowest rail). (c) Generalisation of the circuit to include $M > 1$ qubits for each binary digit; here $K = 1$ and $M = 2$. The last procedure gives Heisenberg-limit phase estimation for $M > 3$, where the unknown phase ϕ is encoded in the relevant eigenvalue $e^{i\phi}$ of U .

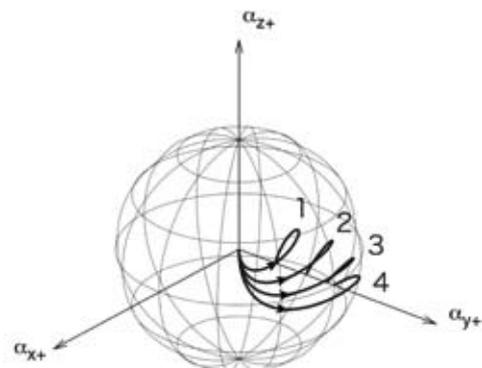


FIGURE 15 From the work reported in [A. Carlini et al., *Phys. Rev. A* **75**, 042308 (2007)], which plots the optimal trajectories in the control parameter space required to generate an entangler gate which takes the initial two qubit state $|00\rangle$ and generates the final state $\cos\phi|00\rangle - \sin\phi|11\rangle$, for $\phi = k\pi/8$, $k=1,2,3,4$.