

Visualizing a silicon quantum computer

Barry C Sanders^{1,2}, Lloyd C L Hollenberg³,
Darran Edmundson^{4,5} and Andrew Edmundson⁴

¹ Institute for Quantum Information Science, University of Calgary, Calgary, Alberta T2N 1N4, Canada

² ARC Centre of Excellence for Quantum Computer Technology, Macquarie University, Sydney, New South Wales 2109, Australia

³ ARC Centre of Excellence for Quantum Computer Technology, School of Physics, University of Melbourne, Victoria 3010, Australia

⁴ EDM Studio Inc., Level 2, 850 16 Avenue SW, Calgary, Alberta T2R 0S9, Canada

E-mail: bsanders@qis.ucalgary.ca, lloydch@unimelb.edu.au and darran@edmstudio.com

New Journal of Physics **10** (2008) 125005 (20pp)

Received 15 April 2008

Published 1 December 2008

Online at <http://www.njp.org/>

doi:10.1088/1367-2630/10/12/125005

Abstract. Quantum computation is a fast-growing, multi-disciplinary research field. The purpose of a quantum computer is to execute quantum algorithms that efficiently solve computational problems intractable within the existing paradigm of ‘classical’ computing built on bits and Boolean gates. While collaboration between computer scientists, physicists, chemists, engineers, mathematicians and others is essential to the project’s success, traditional disciplinary boundaries can hinder progress and make communicating the aims of quantum computing and future technologies difficult. We have developed a four minute animation as a tool for representing, understanding and communicating a silicon-based solid-state quantum computer to a variety of audiences, either as a stand-alone animation to be used by expert presenters or embedded into a longer movie as short animated sequences. The paper includes a generally applicable recipe for successful scientific animation production.

⁵ Author to whom any correspondence should be addressed.

Contents

1. Introduction	2
2. The silicon quantum computer	3
2.1. The building blocks: phosphorus donors in silicon	3
2.2. Qubits, control and readout	5
2.3. Overcoming decoherence—quantum error correction	6
2.4. Scale-up	6
3. Need for the visualization	6
3.1. Outreach	6
3.2. Funding	7
4. The team	7
5. Converting science to visualization	8
5.1. Storyboard	9
5.2. Animatic	9
5.3. Wireframe	10
5.4. Render and composite	10
6. ‘Behind the scenes’	11
6.1. Motivation	11
6.2. Zoom	12
6.3. Atomic scale reached	12
6.4. Quantum control of a single-electron spin qubit	13
6.5. Interaction between two qubits via a J-gate	14
6.6. Transport of an electron spin along a donor rail to a readout zone	15
6.7. Readout of the qubit state through a tunnelling event to a reservoir	15
6.8. Error-corrected CNOT	15
6.9. The ‘death star’	16
6.10. Algorithmic applications	17
7. Sound	17
8. Using the visualization	18
9. Future work	19
10. Conclusions	19
Acknowledgments	20
References	20

1. Introduction

The fast-growing field of quantum computation is a multi-disciplinary research area involving physicists, computer scientists, chemists, mathematicians and engineers, with the objective of making a computer that operates on quantum principles and runs on quantum algorithms. The reason for the growth of, and interest in, the field is that quantum computers can efficiently solve certain problems that are intractable on existing classical (non-quantum) computers, such as Shor’s factorization algorithm [1], and are provably faster than any classical algorithm in other cases, such as Grover’s search [2]. The technology for quantum computing is embryonic

with many candidate technologies, including nuclear magnetic resonance, ion traps, photons, neutral atoms, superconducting Josephson junctions, quantum dots and bulk silicon [3].

The field is both exciting and challenging, not only because of the promise inherent in the endeavour, but also because of the need to integrate a multidisciplinary team in a focused effort, educate and train students in topics that extend beyond their disciplinary training, communicate results to the public and media, and convey to funding agencies progress and rationales for large levels of funding. The challenge of communicating the concepts and research directions and goals to such diverse groups motivated us to develop a sophisticated, concise animation that shows how our project—a solid-state quantum computer constructed as quantum bits (qubits) corresponding to phosphorous atoms injected into a silicon substrate—will work. Here, we explain how the animation project developed, the need for the visualization, the choices that were made with respect to what was included and what was excluded, how to present the concepts, the sound accompaniment to the animation and the value of the animation that was produced.

Animation is expensive—both in terms of time and cost—to produce, and is used so rarely in science that its value is largely unappreciated. Furthermore, funding for visualization is not easy to find. A research grant primarily supports expenditures that lead to new research results in the discipline, and arts funding that would support animations is small and would be unlikely to be assigned for science or technology animation. Although the educational value of the project is apparent, funding for educational physics is quite limited. Thus, acquisition of sufficient funding became the first hurdle.

Fortunately, the visualization also informs the funders about their investment and how it would work. The need for an animation to make informed judgments about such a complex, interdisciplinary project justified the cost. With the backing and expertise of the Australian Centre of Excellence—Centre for Quantum Computer Technology (CQCT), and after viewing prior animation work by one of us (Sanders, who produced the animation) and his team, the Laboratory for Physical Sciences in Maryland underwrote the project. A scientific advisory team was assembled comprising the expertise of the Centre's physicists responsible for the full-scale silicon architecture proposal itself [4], based at the University of Melbourne, including a former student who joined the University of Waterloo as a postdoctoral research fellow.

Communication of the complex concepts inherent with quantum computing in a short animation is challenging; it requires careful consideration of how to visually represent quantum concepts and components, depict processes, explain the purpose and motivation of the quantum computer, all while maintaining an aesthetic that engages the interest of the viewer. Calgary-based scientific animation experts EDM Studio brought technical and visualization expertise to the project as well as project management capabilities, the latter was valuable to ensure timely delivery of the animation at the highest standard with full team consensus on each stage of the project. The project management and communication and decision processes are quite important in visualization projects.

2. The silicon quantum computer

2.1. *The building blocks: phosphorus donors in silicon*

The system to be visualized spatially and temporally is a full-scale quantum computer architecture with elements comprising individual donor phosphorus atoms in a silicon crystal

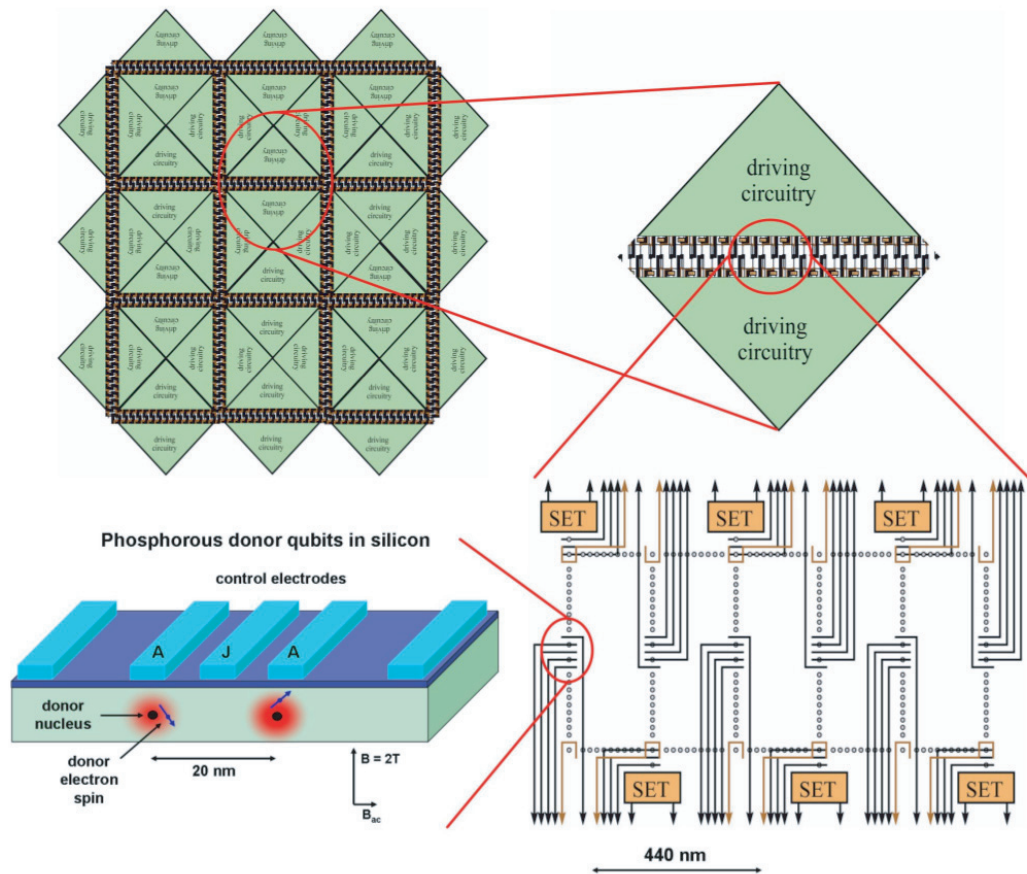


Figure 1. Highly schematic view of a silicon-based quantum computer showing all scales from the macroscopic classical control chip level down to the physical quantum core involving the control of individual electrons around phosphorus donor atoms.

(figure 1) [4]. The concept is based on the original ideas of Kane [3]. The phosphorus donor atoms are placed 20–30 nm apart in a regular array 20 nm below the surface of the crystal. On the surface, an array of nanoelectronic gates 10 nm in width provides addressing and control of the atomic states of individual phosphorus atoms. Such atomic-level fabrication is now achievable using ion implantation techniques [6], and to high precision (at the 1–2 nm level) using STM techniques [5], and progress has been made in demonstrating the required single atom control in these devices [7].

The physics of phosphorus donors, fundamental to conventional complementary metal oxide semiconductor (CMOS) device operation, is well known. Each phosphorus atom replaces a silicon atom in the crystal matrix. After four of the five valence electrons bind with neighbouring silicon atoms in the crystal, the donor atom effectively behaves as a single electron hydrogenic system with an atomic radius of 2 nm. Thus, the range of this bound electron is quite large compared to a silicon crystal spacing of 0.54 nm.

Remarkably, although the spatial extent of this phosphorus donor electron encompasses many silicon atoms and their respective binding electrons, the magnetic properties of the donor electron represented by the electron's spin states (pointing up or down with respect to a magnetic

field) are only very weakly affected by the crystal lattice. This proves to be a crucial property within the quantum computer design.

2.2. Qubits, control and readout

A quantum computer differs fundamentally from a conventional computer in the way binary information is stored and processed. In a conventional computer, a bit of information is represented by values 0 or 1. The quantum analogue of a bit is the qubit, which according to the rules of quantum mechanics can effectively assume both 0 and 1 values at the same time. In a silicon quantum computer, the qubit is most naturally represented by the two spin states of the donor electron. The fact that the spin states are largely unaffected by the complexity of the surrounding silicon lattice, a fluke of nature, means that donor electron spins offer a natural physical representation of a qubit.

In order to perform information processing on such qubits, one must be able to control the electron spins individually, have them interact in specific pairs and readout the information. The original idea of controlling donor spins by surface gate electrodes was put forward by Kane [3]. In a constant background magnetic field one can distinguish spin-up (1) and spin-down (0) states, whereas an oscillating field at the correct frequency can flip the spin states of the electron (physics similar to that employed in an MRI machine).

By applying a suitable bias to the A-gate directly above a specific donor atom in an interaction zone, the deformation of the electron charge distribution focuses the effect of the oscillating field onto a single target donor atom qubit. In this way, one can affect single qubit logic gates: a flip of the spin state or even manipulation of the quantum state of a specific donor electron into a superposition state of both spin-up and spin-down states. In binary representation, this corresponds to 0 and 1 states existing simultaneously. This simultaneity of binary representation leads directly to the massive parallelization possible when considering many qubits in a quantum computer.

Qubits interact according to the rules of quantum mechanics which again allows an entire two-qubit Boolean truth table to be effected in a single step, in stark contrast to a conventional computer. Such a two-qubit quantum gate is implemented in a silicon system using the J-gate between neighbouring donor atoms. A bias applied to the J-gate can draw the electrons closer (or push them apart) and hence control the interaction between their spin states. A specific combination of single qubit operations (effected by A-gates) and two-qubit operations (effected by J-gates) is required to build up a controlled-NOT gate [8, 9], an important universal two-qubit logic gate required for computation.

Read-out of the information encoded on the donor electron spin qubits requires a spin measurement with sensitivity down to the single individual spin level. Single electron transistor (SET) structures are required which are sensitive enough to detect the process of an electron tunnelling from a donor site to an electron reservoir. By suitable tuning of surface gates, the tunnelling process can be made spin-selective (e.g. spin-up only), so that the SET detection event corresponds to a specific spin state of the donor electron and hence the qubit state.

According to another twist of quantum mechanics, whereas the electron's quantum spin state may be in both up (1) and down (0) states simultaneously, a measurement (or readout) process will only detect one of the spin states with a probability determined by the specific quantum state prior to measurement. This stochastic access to the quantum information stored within the quantum computer is central to the way in which quantum algorithms work.

2.3. Overcoming decoherence—quantum error correction

One significant obstacle to be overcome is the fragility of quantum information to stray interactions with the rest of the universe (usually called the environment), and the vulnerability to errors due to such decoherence processes increases with the number of qubits. Even though the quantum coherence of the donor electron spin state is relatively long-lived (60 ms) compared to the typical quantum gate operation time (sub-microseconds), decoherence processes will eventually pollute the information and render the computer useless.

To counter these errors, quantum error correction protocols have been developed that allow for fault-tolerant operation at the cost of massive redundant encoding of quantum information across many physical qubits. The encoding is arranged recursively, e.g. a first-level logical qubit is constructed from a group of physical qubits (14 in the so-called Steane code), which, when combined with readout over a subset of these qubits and feed-forward correction, provides protection against decoherence errors to the next order in the error rate.

Logical encoding over groups of qubits and the associated addition of qubit gates in the quantum error protocols requires some form of nonlocal qubit interaction. In a silicon quantum computer, the donor atoms are fixed in a silicon lattice and restrict qubit interactions to the nearest-neighbour local interactions. To circumvent this problem, the electrons are transported along rails of donors using surface gates in order to be placed into a specific interaction zone at a specific stage of the error-correction process [10]. The bi-linear array visualized here, with horizontal and vertical transport rails into interaction zones, allows for the necessary quantum error-correction protocols [11].

Concatenation of the encoding proceeds by replacing each physical qubit with a first-level logical qubit protected against first-level errors. This second-level logical qubit is now protected to much higher order against decoherence errors. The process will work in principle because, whereas the number of qubits grows exponentially with concatenation level, the protection against decoherence grows super-exponentially provided the fault-tolerant condition is met.

2.4. Scale-up

Taking all this into account, to perform a large-scale factoring problem for numbers of binary length 1024 or greater, a quantum computer operating fault-tolerantly (below the error threshold) would comprise a significant fraction of a billion qubits. It is clear from such estimates that in terms of assessing the feasibility of a quantum computer the effective density of qubits matters, as does the fundamental quantum operation time. A strength of the silicon quantum computer design is that qubits at the nanoscale are controlled by CMOS-type electronics, which in silicon has great potential for scale-up to the large component numbers required.

3. Need for the visualization

3.1. Outreach

Whereas ongoing developments in quantum computer technology captivate the public and is notable in receiving significant media attention, understanding the concept and purpose of quantum computation can be elusive. Given the widespread interest in quantum computation

and the large public investment, outreach is important, and visualization is an effective tool for this purpose.

3.2. Funding

Quality scientific animation, as the reader will hopefully come to appreciate from the remainder of this paper, is necessarily expensive: of the order of 10 000–25 000 \$/min to produce. Obtaining funding is thus a challenge. Funding for the animation was specifically targeted to the production of an animation that could be used for explaining the solid-state quantum computer especially to funding agencies.

Funding of this animation was enabled by the project underway for constructing the basic building blocks towards a silicon-based solid-state quantum computer. Although the animation is admittedly expensive, the cost is a very small fraction of the overall research and development budget, thus making it affordable in the larger context. Despite the communications value inherent in such an animation, other sources of funding are scarce as the visual tool does not fit into the usual project descriptions of various funding agencies.

It is important to note that the animation is not a research project, but rather a representation of such a project. As new animation techniques are not developed, the project would not be supported as part of computer graphics funding. Already scarce arts funding is also not a suitable choice for supporting a science-oriented animation. Educational funding would be appropriate, but such funding is often inadequate and highly competitive. Therefore, we believe that the best source of funding for a major animation undertaking is as part of the training, outreach and marketing budget of a research network with funding significantly greater than the cost of the animation.

4. The team

In order to prepare an animation that is physically accurate, subject of course to necessary aesthetic and time constraints, a team comprising quantum computation scientists and visualization experts was required. The expert technical team comprised Lloyd Hollenberg, Andrew Greentree, Ashley Stephens, of the University of Melbourne node of the CQCT and former PhD student from the University of Melbourne, Austin Fowler, since relocated to the University of Waterloo.

The technical team originated the concept, having expertise spanning all areas of a silicon quantum design, from quantum device modelling to quantum error correction. The technical visualization was undertaken by Calgary's EDM Studio. EDM's co-founder Darran Edmundson's combined background in physics and scientific visualization—including a PhD in computational physics and six years of experience at the Australian National University Supercomputer Facility Vizlab—played an important role in the project's success. In addition to offering technical expertise, Barry Sanders served as producer. The team was spread over three cities in two countries, and the diversity of expertise made communication a challenge. At EDM's behest, the popular web-based project management tool Basecamp^{TM6} was used for discussion threads, to-do lists, project milestones and file exchange.

⁶ <http://www.basecamphq.com/>

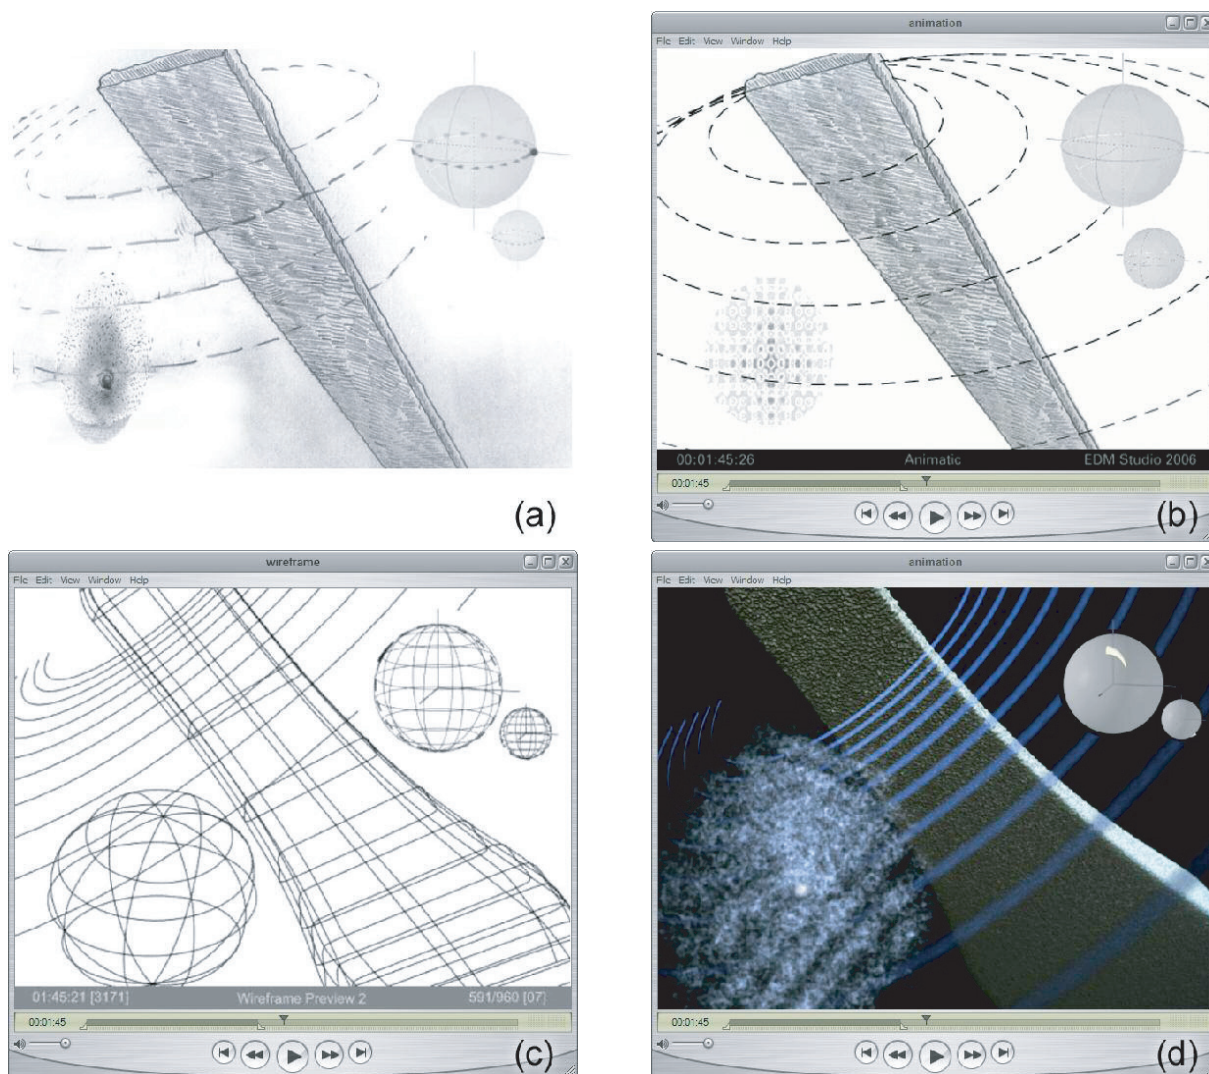


Figure 2. The major stages of animation production as illustrated by an excerpted frame from our ‘Quantum Computer in Silicon’ animation: (a) storyboard, (b) animatic, (c) wireframe and (d) render and composite.

5. Converting science to visualization

Despite the orderliness implied by the section title, the representation (conversion) of hard science as compelling visualization is quite difficult. The crux of this difficulty is the tension between the inherent complexity of any scientific story (as depicted in figure 1) and the harsh constraints of animation production. In our case, success resulted from ongoing communication between the two camps, scientific and artistic. The science team came to understand and respect the well-defined stages of animation production. The artists, meanwhile, came to accept that accuracy—a fidelity to the core science—was absolutely paramount, even above aesthetics.

Given the readership of this paper, it is worth explaining the animation production pipeline in some detail. Generally, the process is one of successive refinement, each completed stage being treated like a locked door that precludes back-tracking. Whereas approaches vary,

generally the pipeline comprises the steps shown in figure 2: storyboard, animatic, wireframe and render and composite. These stages are described in some detail below.

5.1. Storyboard

The first—and without question, the most important—step in creating a successful animation is the completion of a detailed storyboard. This chronological sequence of cartoon-like panels should be sufficiently fine-grained that a reader is able to connect the constituent snapshots into a mental movie. Our four-minute ‘Quantum Computer in Silicon’ animation began life as a 39-panel storyboard, the result of several hundred person-hours of studio and scientific advisor input.

Before any storyboarding takes place, however, discussions with the animation producer should address the following:

- Who is the target audience? Are they highly technical peers, undergraduate students, or members of the general public?
- Closely tied to the above, what is the intended use of the animation?
- What reference materials exist for educating and informing the creative team? This can include published papers, slides from presented talks, soft science papers (Scientific American, New Scientist), simulation results, television documentaries on the field, etc.
- Which elements of the science need to be included, and which can possibly be removed without compromising the overall story? Just as with the film version of a book, tendrils of the science plot will (by virtue of budget and time constraints) need to be simplified or perhaps omitted.

The scientists themselves, both as stakeholders and with a thorough understanding of the processes to be depicted, are a good starting point for storyboard creation. We proceeded as follows. For the key sequences identified during initial research, each member of the science team was asked to submit comic-book-style sketches to represent physics. It was emphasized that submitted diagrams should be limited to 3–5 rectangular panels, be extremely simple, and contain no mathematics.

The aim of this exercise was to

- (i) distill a given process down to a small number of key visual steps;
- (ii) identify a number of possible alternatives for visualizing each topic; and
- (iii) gain a clearer understanding of each physicist’s mental movie—if major differences of opinion among members of the science team are to emerge, these need to be addressed earlier rather than later.

From these panels, and in close discussion with the producer, the results were synthesized by the studio artist into a preview storyboard for further iteration.

Completion of the storyboard was formalized with a sign-off by the producer. While this might seem somewhat regimented, the effect was to reinforce that story issues need to be clarified at this early stage, not later when costly computer work has already begun.

5.2. Animatic

In stage two of the process, the static storyboard is converted into the simplest of movies. In this *animatic*, each panel of the storyboard is held on-screen for some duration, perhaps with

simple pans or zooms indicating desired camera movements. One typically finds that while the storyboard ‘works’, the relative duration of some scenes need to be lengthened or shortened. Again, completion of the animatic phase is formalized with sign-off by the producer and/or science team. From this point onwards, with very few exceptions⁷, the basic content and duration of each scene should remain unchanged. If the animation requires sound (music, narration or both), work can begin using the animatic as a reference.

5.3. Wireframe

In the next stage of production, work finally commences on the painstaking task of sculpting and animating digital models. Whether it is the representation of morphing electron clouds or changing magnetic flux, the benefit of having a well-developed storyboard and animatic becomes evident, allowing the artists to proceed in digital construction with confidence. We refer to this stage as *wireframe*, because the scene elements are rendered as polygonal outlines. This conscious decision to avoid surface properties like texture, colour, shininess, etc forces both stakeholder and creative teams to concentrate on the more primitive aspects of the budding animation.

5.4. Render and composite

In the final stages, the animation team applies material properties (the so-called shaders) to the previously wireframe forms. Added realism and, where necessary drama, is created through the careful placement of lights and inclusion of shadows⁸. Conversion of these scene descriptions into actual images is known as *rendering*. Highly computer-intensive, it is not uncommon for a single frame to take an hour or more to complete. Considering that playback occurs at 30 frames/s⁹, a minute of finished animation might represent a month of single CPU time. The consequences of this are twofold. Firstly, many CPUs must be brought to bear on the problem to keep render times tractable¹⁰. Luckily rendering, with one image per CPU, is a trivially parallel computation. Secondly, frames are typically rendered in a dozen or more layers with separate *passes* for the specular, diffuse and shadow components of every object. A given frame of final animation is then assembled through composition of these multiple layers. The flexibility of this approach allows for huge changes to the final ‘look and feel’ without necessitating re-rendering. And, if new renders are required, often only a small number of passes need to be recomputed.

The reason for this staged approach should now be apparent. Unlike movies where entire scenes frequently end up on the editing room floor, in animation we cannot afford to waste the large effort inherent with digital modelling, shader writing and rendering. The benefit of developing a comprehensive storyboard—and more importantly, sticking to it through the rest of the production—cannot be overstated.

⁷ There is a natural tendency, typically at the behest of stakeholders, to consider expanding existing or adding new scenes during later phases of production. Such ‘scope creep’ inevitably leads to time and budget over-runs. Our team made a conscious decision to leave such upgrades to a future version. See section 9 for details.

⁸ Unlike photography, in computer animation shadows are not automatic but have to be computed.

⁹ Actually, 29.97 frames/s for NTSC television (North America and Japan), 25 for PAL television (Europe and Australia) and 24 for film.

¹⁰ Loosely networked clusters of machines dedicated to creating CGI imagery are known as ‘render farms’. Commercial render farm services exist allowing individuals and small teams to produce substantial animations in a timely manner and at low cost without the commensurate investment in computer hardware.

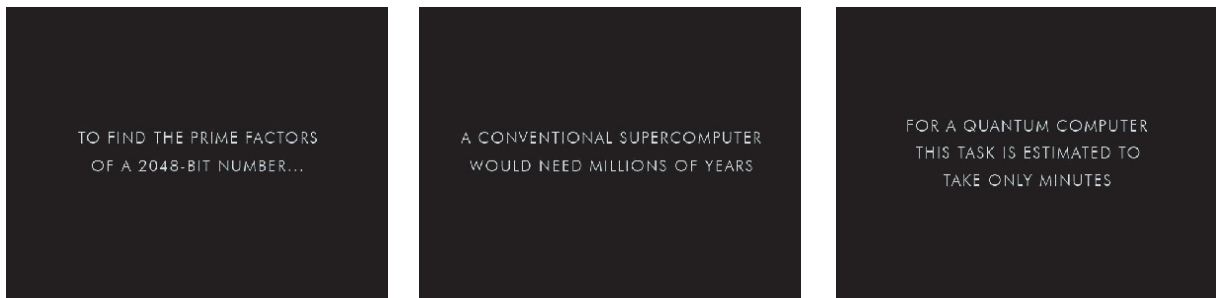


Figure 3. The animation begins with a simple yet effective comparison statement between conventional and quantum computers.

6. 'Behind the scenes'

The quantum computer architecture spans some ten orders of magnitude in both spatial and temporal scales. Length scales range from the sub-nanometre (crystal lattice spacing and donor nuclear spin effects), nanometre (lithography limit for electronics), micrometre (on-chip CMOS processors), to millimetres (fan-out of electronics) and metres (microprocessors, dilution fridge). As complex as the various regimes are spatially, this aspect of the visualization is quite straightforward compared to the challenge of conveying the quantum operation of the device to both technical and non-technical audiences.

Early in the storyboarding process, the flow of the visualization was decided as follows:

1. factoring problem motivation for quantum computing;
2. zoom from human to atomic scale;
3. representation of single-electron spin qubit;
4. single-qubit control;
5. interaction between two qubits;
6. transport of qubit along a donor rail;
7. readout of qubit via tunnelling to reservoir;
8. pull-back to reveal complex circuit corresponding to an error-corrected logical CNOT gate;
and
9. figurative explanation of factoring algorithm.

6.1. Motivation

The potential for exponential speed-up of a quantum versus conventional computer is conveyed with a simple comparison statement based on the factoring of a large (2048-bit) number. Figure 3 shows this 10 s sequence of 'white text on black background'. Inexpensive yet effective, it allows finite production resources to be redirected to more complicated aspects of the animation.

Within the science team, the exact phrasing of this sequence was the subject of extensive (often mathematical) discussion. This debate over 'a few words' came as a surprise to the animators but is precisely the type of cultural differences that can cause tension between collaborative teams of artists and scientists.




Figure 4. Elements of the zoom from human to atomic scale: (a) storyboard panel showing motherboard of the futuristic computer, QC chip at centre in standard IC packaging; (b) rendered frame of the same and (c) mixed image showing both storyboard and final rendered versions of the plunge to atomic scale.




Figure 5. Phosphorus atoms embedded in a silicon lattice with probability cloud of loosely bound electron: (a) science team description; (b) wireframe version and (c) final rendered version.

6.2. Zoom

The physical aspect of the animation begins with a ‘top-level’ artist’s view of a QC chip indicating the solid-state nature of the device and the link to conventional silicon-based CMOS fabrication (figure 4). The actual details of how the quantum computer core would be housed and connected to microprocessor control in the real world are a matter of current research. We begin with a zoom into the chip, settling on a bird’s-eye view where the camera is hovering over the tiled array showing fan-out of nano-electronics. At the scale of a few microns across, this is a view of a section of the quantum core itself. The atoms hosting the electrons, which themselves carry the quantum bits of information, lie just beneath the surface in the crystal lattice. In this context, briefly hinting at the complexity below, the camera plunges into the bi-linear array and silicon crystal.

6.3. Atomic scale reached

At this maximum magnification, figure 5, we see the electron probability cloud localized around a single phosphorus donor¹¹. The nuclei in a silicon crystal are shown to provide a reference cue to the solid-state nature of the environment.

¹¹ Only the probability distribution for the loosely bound electron is shown, the covalently bonded electrons that bind the phosphorus into the surrounding silicon lattice are omitted for clarity.




Figure 6. Transition to a rotating reference frame where otherwise helical spin dynamics become a polar great circle: (a) initial explanatory drawings from science team; (b) storyboard version and (c) excerpt from final render.

The electron probability cloud created by the animation team, crafted based on reference images provided by the advisory team, nonetheless possesses two characteristics important to the science stakeholders: (i) the extent of the cloud roughly matches that obtained from simulation; and (ii) the cloud density is not space-filling but rather reflects locations of silicon nuclei. In later sequences, where the lattice is omitted for clarity, the probability cloud still reveals an internal structure commensurate with the solid-state environment.

Magnetic fields, which cause the electron's quantum spin state to evolve, are switched on. To represent the quantum state, a Bloch sphere is shown—a representation of the qubit encoded into electron spin. For the non-technical audience, the south pole of the sphere is visually identified as the binary zero state, the north pole the binary one state. Elsewhere on the sphere, the electron is in a combination of one and zero simultaneously—a quantum bit, or qubit. As the electron's quantum spin state evolves in the magnetic field, the system traces out a helical path on the sphere passing through the pole binary states zero and one.

The spin-up scene, figure 6, is a representation of a purely mathematical operation transforming the viewpoint of the quantum evolution to a rotating frame where the otherwise helical path becomes a simple trace around a great circle. While this is merely an artifice that experts use for visualizing the mathematics of the quantum dynamics, it makes for an effective sequence in the animation (see *rotate.mov*, available from stacks.iop.org/NJP/10/125005/mmedia). In this rotating view, the quantum information encoded on the electron spin transforms smoothly between zero and one poles—the quantum analogue of a NOT gate.

6.4. Quantum control of a single-electron spin qubit

In the next scene of the animation, figure 7, a single qubit is controlled by an electric field. To reduce visual clutter and help maintain the viewer's focus on the Bloch sphere, the source of the field is not shown. The electric field distorts the electron cloud, taking the electron's spin state to a point on the equator of the Bloch sphere, representative of a balanced combination of logic zero and one. This is an idealized Hadamard operation, a component of the universal set of quantum logic gates. Note that in the storyboard version shown here, the equatorial spin dynamics on the Bloch sphere are incorrect. Fixed during a subsequent iteration, this is a good example of the need for frequent content reviews by the science team.

We then repeat the Hadamard operation, changing the camera view to reveal the surface nano-gate controlling the target phosphorus atom. An excerpt from this sequence is shown




Figure 7. Quantum control of a single electron spin qubit: (a) storyboard version; (b) excerpt from wireframe version and (c) excerpt from final render.




Figure 8. Interaction between two qubits via a mediating J-gate: (a) initial explanatory drawings from science team; (b) early storyboard version and (c) final render. The change in camera angle apparent between the storyboard and final render occurred in the early stages of wireframe modelling.

in figure 2 (referenced earlier in the context of the various stages of animation production). Notice that a second smaller Bloch sphere has been added. As timing plays an important role in the proposed quantum computer, this reference Bloch sphere acts as a clock with which operations must be commensurate. Close attention to the movie clip of this second Hadamard (hadamard.mov, available from stacks.iop.org/NJP/10/125005/mmedia) shows that the operation takes place in two distinct steps that leave the spin state stationary on the equator after one polar cycle of the reference spin. This is a good example of the attention-to-detail that is of great importance to scientists wishing to use a visualization for professional purposes. Animation teams, especially those not familiar with scientific visualization, need to recognize that such accuracy comes before aesthetics. A corollary to this point is that the development of scientific animation proceeds significantly slower than aesthetics-driven commercial projects.

6.5. Interaction between two qubits via a J-gate

Using the single qubit as a building block, the animation moves on to a two-qubit logic operation, the controlled-NOT (CNOT). In figure 8, a surface J-gate positioned between two donor atoms mediates coupling between qubits. The quantum-mechanical CNOT is complicated, requiring a well-defined sequence of both single qubit- and centre-gate operations.

Initially, and as seen in the early storyboard panel, the team envisioned showing the spin dynamics of the constituent operations. By the animatic stage, however, it became clear that




Figure 9. Readout of the qubit: (a) initial explanatory drawing from science team; (b) excerpt from animatic (note time-code bar at bottom) and (c) final render.

a full Hadamard-like description of the CNOT would be long, complicated and ultimately detracts from the overall narrative. As can be seen in the final rendered version (and the excerpt `cnot.mov`, available from stacks.iop.org/NJP/10/125005/mmedia), the Bloch spheres were omitted. This allowed an otherwise hard-to-understand quantitative description to become an enjoyable qualitative experience. Importantly, reducing the visual complexity focuses the viewer's attention on the important teaching point—that the CNOT involves a sequence of single-qubit operations interspersed with coupling operations mediated by a central gate.

6.6. Transport of an electron spin along a donor rail to a readout zone

During the operation of the quantum computer, the electron qubits move in and out of interaction and readout regions according to a predefined schedule. In the animation, we show transport occurring in a distinctly quantum fashion (a controlled quantum tunnelling event) along a pathway of ionized phosphorus atoms, literally reappearing at the required terminus in the foreground, in this case a measurement region.

6.7. Readout of the qubit state through a tunnelling event to a reservoir

At the interface of the classical and quantum worlds is readout (figure 9). In a process dependent on its quantum state, the animation shows the electron selectively tunnelling to a reservoir structure with the delocalized electron cloud appearing inside the reservoir volume. The detection of this tunnelling event signals a particular qubit state, and is detected by hypersensitive single-electron transistors constantly monitoring charge motion in the core.

Early investment in the development of custom animation tools allowed the creative team to quickly realize this shot. Specifically, the digital asset used for visualizing the localized (deformable sphere) electron cloud shown previously was leveraged to create the large rectangular cloud inside the reservoir. Tool and animation pipeline development is an important part of digital production.

6.8. Error-corrected CNOT

With the various components of the quantum-core illustrated now, the camera pulls back to view the bi-linear array with all attendant nanoelectronics in full operation (see `schedule.mov`, available from stacks.iop.org/NJP/10/125005/mmedia). The electron spin qubits are shown as

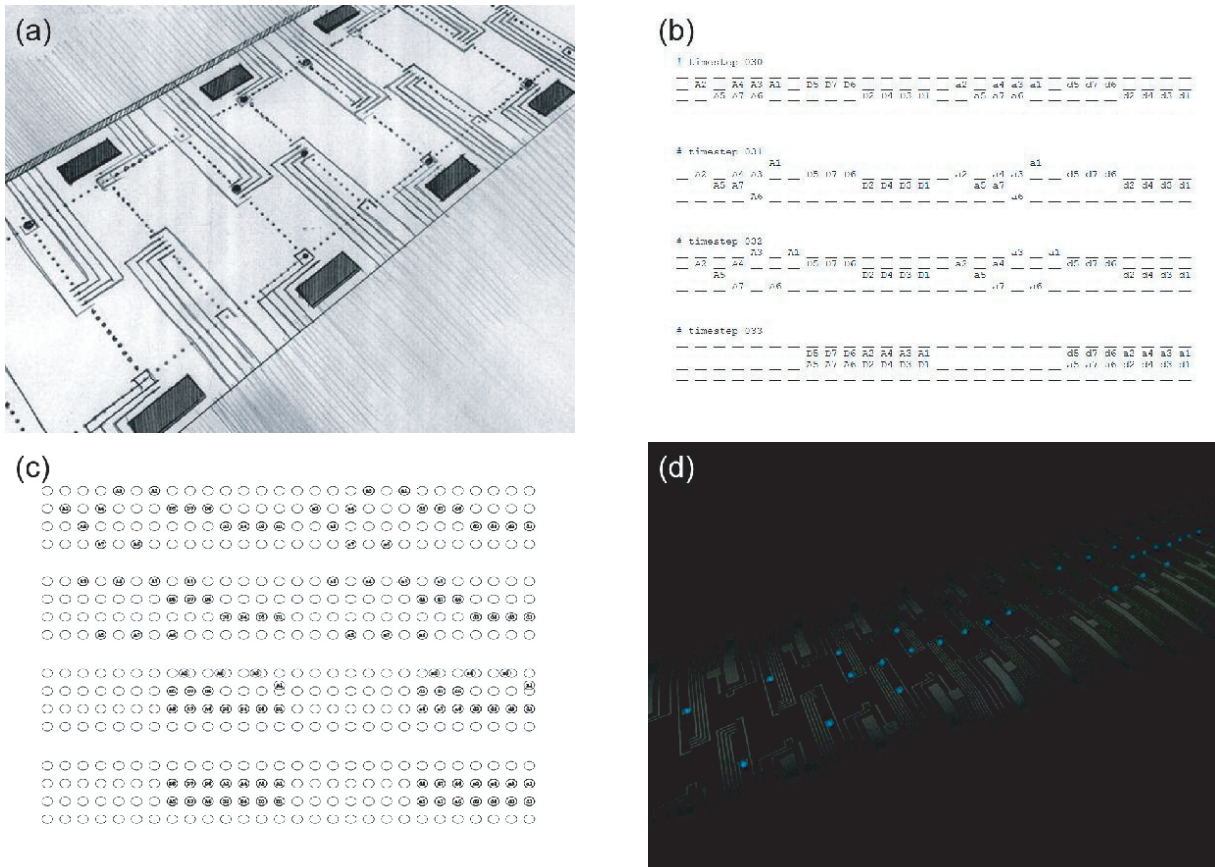


Figure 10. Error-corrected CNOT operation: (a) storyboard version of schedule; (b) machine-readable schedule submitted by scientists; (c) preview animation for verifying correctness and (d) final rendered version.

blue clouds moving along the bi-linear rails below the surface. At this level 14 physical qubits are grouped together to form a single logical qubit in order to protect against the constant errors induced by the environment.

Importantly, this ‘hero shot’ is faithful to the complex series of electron qubit motions, readout and control required by fault-tolerant quantum error correction for carrying out a first-level logical CNOT gate. Getting this right required close collaboration between the science and animation teams. Rather than animating qubits manually, figure 10(b) shows a parsable text description invented for encoding qubit movements. A custom python-language parser used this information, plus a built-in knowledge of allowable qubit moves, to automatically output animation paths for all 14 qubits. This data was rendered in a preview form (c) for validation by the science team prior to final rendering and composite with the surrounding gate structure (d). Anticipating the need for iteration, this semi-automated approach allowed rapid changes to the schedule animation.

6.9. The ‘death star’

In our final look at the quantum computer proper (figure 11), the camera pulls back to show large numbers of qubits moving according to the quantum error correction schedule. These

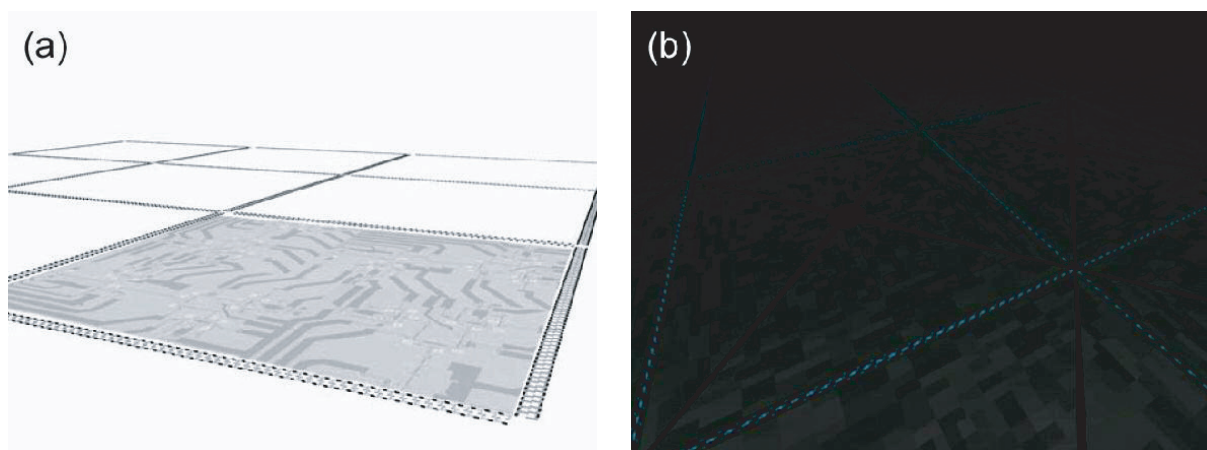


Figure 11. Tiled rails of error-corrected CNOTs border islands of classical support circuitry in this shot internally referred to as the ‘death star’: (a) storyboard version and (b) final rendered version.

rails of qubits, bordering giant square islands of classical support circuitry, were briefly visible during our initial plunge to the atomic scale. With the benefit of the intervening component descriptions, our hope is that the audience is now aware of the immense but ordered complexity that lies below.

This sequence has the potential to be extremely emotive, a micron-scale analog of the futuristic cityscape seen in many movies. Time and budget constraints, however, precluded the creative team from expending all but a token effort on this shot. Diverting finite production resources away from the visualization of core science to more aesthetically rewarding shots is an easy trap for animators to fall into—if and when such a decision is made, it should only be with the informed consent of the producer.

6.10. Algorithmic applications

At the highest logical level, the quantum algorithm dictates the sequence of quantum gates needed for performing the task at hand. This brings the animation back to the factoring of large numbers. Such an application is, of course, most useful in the encryption of classical data. The actual details of Shor’s quantum factoring algorithm are extremely difficult to understand, let alone present in a visualization. In agreement that a figurative approach was best, the creative team proposed the storyboard panel of figure 12(a). This shows potential factors impacting a 1D ‘wall’ notched to allow the correct answer through. Subsequent discussions with the science team revealed that this depiction was not even figuratively correct—catching this problem prior to the commencement of digital production avoided the waste of precious resources. The revised sequence shows four classical buffers (b) that subsequently are used for generating probable factors of a large binary number (c).

7. Sound

Audio, whether it is voice-over narration, sound effects, or music, should be considered an integral part of the final animation. In our ‘Quantum Computer in Silicon’ animation, Melbourne



Figure 12. Figurative explanation of the factoring of large numbers by Shor's algorithm: (a) initial storyboard proposal, subsequently changed; (b) filling buffers in the final version and (c) 'cracking' the code in the final rendered version.

composer and sound designer Tim Kreger was involved from the animatic stage onwards. In addition to custom sound effects for such items as the spinning electronic cloud and controlling magnetic/electric fields, Tim crafted a compelling ambient music track that propels the animation forward towards the so-called 'hero shot'—an accurate visualization of the full error-corrected CNOT circuit discussed in section 6.8.

As our animation was intended for presentation by a technical expert in a play-pause-discuss scenario, it does not have a voice-over narration. With a year's worth of experience in showing the animation to various audiences, writing an effective narration—part of the plan for future work—will be relatively easy.

8. Using the visualization

The primary purpose of the animation is to act as an important aid in explaining the details of a full-scale quantum computer to both technical and non-technical audiences. Technical audiences include the CQCT team, conferences and workshops, and collaborators. Non-technical audiences are targeted in outreach and educational activities. Funding agency reviews are also important, and the animation is used for these purposes, with an audience that generally does not have expertise in quantum computing *per se*.

In the first year, the animation has been shown to audiences ranging from public lecture level through technical non-experts to technical experts. Even amongst the CQCT team, the animation is highly appreciated, particularly by incoming students and research and academic staff; it presents—for the first time—an easily digestible and comprehensive view of a silicon quantum computer.

While the animation is rich with technical detail, aspects such as control sequences or timing and scale issues are not overemphasized—the animation can be appreciated without being overwhelmed with information. However, the scientific content of the animation is useful for technical audiences. Our approach to using this visualization for technical audience or postgraduate education and training is to display the animation along with an accompanying PowerPoint™ presentation. This presentation provides detailed calculations, caveats on the claims, and accurate depictions of certain components that were necessarily abridged on resource or aesthetic considerations (e.g. the detailed structure of the electron cloud).

9. Future work

While there are many ways of improving the animation, resource limitations requires further improvements to be strategic and costed in advance. Most tempting is the opportunity to improve the aesthetics of the animation. For example, at present the electron cloud representation gives some indication of the spatial extent of the electron, but could be made much more realistic in terms of the effect of the underlying lattice periodicity. The structure of the electron cloud is known from theoretical models, and this knowledge could be directly incorporated in the rendering process.

In fact, the state of the animation need not be static. Over time, as demand for the animation continues, it is possible to improve the animation on an ongoing basis given the resources to do so. Moreover, the current research developments should be accompanied by animation updates but should use as much pre-existing animation as possible to optimize resources. One priority for us is to improve the representation of the control chip as new ideas are developed and the nature of this control is better understood.

We are considering the best approach to disseminating the animation. With an expert presenter, the animation is a valuable tool in communicating the nature of solid-state quantum computing. Without the expert presenter, the animation can be confusing. For the technical audience, a complete set of technical notes catering to various audiences could be provided, in addition to a scripted voice-over of the existing version. Beyond that, the structure of a presentation on a silicon quantum computer may incorporate the animation into a 15 min show complete with animation, laboratory scenes, and expositors and interlocutors.

Animating the solid-state quantum computer has provided the scientists involved with valuable experience in creating successful animations for educational, training and outreach purposes efficiently and on schedule. Future visualizations in quantum information and beyond are being considered.

10. Conclusions

We have created a four-minute animation that depicts a solid-state quantum computer in silicon with an emphasis on simplicity, accuracy and aesthetics, and this animation has been used numerous times to acclaim. The success of this visualization depended on effective planning, communication and negotiation. We share these principles here for others planning to create scientific visualizations.

Each stage of the animation process is time-consuming, expensive and nearly irreversible. The stages of development and schedule for these steps needs to be agreed in advance and adhered to. Once the process and schedule have been decided, the first step needs to be development of a comprehensive storyboard. The storyboard process brings out the essential science and leads to a shared understanding of the essential science and the representation of this knowledge between the animators and the scientists.

It is important to appreciate the cultural differences between the science and animation communities. Ongoing communication is important. Given that the team was scattered across two continents, appropriate project and task management software was essential for success, and BasecampTM served the purpose for us.

Creating this animation was a rewarding experience for all team members. The interdisciplinarity of the team provided an interesting challenge, and appreciating the different

cultures of the team helped the project proceed smoothly. Consequently, the project was produced on time and on budget. In this paper, we have shared this experience and the important lessons along the way.

Acknowledgments

We highlight and thank the members of the technical team, A Fowler, A Greentree, and A Stephens for their critical contributions to the animation, much of which underpins this paper. The animation was directly funded by the Laboratory for Physical Sciences, based on original research funded by the Australian Research Council, the Australian Government and the US National Security Agency (NSA), Advanced Research and Development Activity (ARDA) and the Army Research Office (ARO) under contract no. W911NF-04-1-0290. BCS has also received support from iCORE and CIFAR. Finally, we also thank R G Clark and G McWright for their support and advice in the making of this animation.

References

- [1] Shor P W 1994 Algorithms for quantum computation: discrete logarithms and factoring *Proc. 35th Annu. Symp. on Foundations of Computer Science* (Los Alamitos, CA: IEEE Computer Society Press) pp 124–34
- [2] Grover L 1997 *Phys. Rev. Lett.* **97** 325
- [3] Kane B E 1998 *Nature* **393** 133
- [4] Hollenberg L C L, Greentree A D, Fowler A G and Wellard C J 2006 *Phys. Rev. B* **74** 045311
- [5] Schofield S R, Curson N J, Simmons M Y, Rueß F J, Hallam T, Oberbeck L and Clark R G 2003 *Phys. Rev. Lett.* **91** 136104
- [6] Jamieson D N *et al* 2005 *Appl. Phys. Lett.* **86** 202101
- [7] Andresen S E S *et al* 2007 *Nano Lett.* **7** 2000
- [8] Fowler A G, Wellard C J and Hollenberg L C L 2003 *Phys. Rev. A* **67** 012301
- [9] Hill C D and Goan H S 2003 *Phys. Rev. A* **68** 012321
- [10] Greentree A D, Cole J H, Hamilton A R and Hollenberg L C L 2004 *Phys. Rev. B* **70** 235317
- [11] Stephens A M, Fowler A G and Hollenberg L C L 2008 *Quantum Inf. Comput.* **8** 330