PEARCEY ORATION

Michelle Simmons AO Melbourne, September 6, 2023

Trevor Pearcey was a true technological pioneer and a great Australian. CSIRAC, the computer he and his team built at CSIRO in the 1940s, was a triumph of ground-breaking electronics and ingenuity. For these reasons alone, it is an honour to deliver an oration in his name. But I can take extra delight in this moment, for Trevor Pearcey and I had more than a few things in common. We both started life in England and migrated to Australia early in our research careers. We both had ambitions to build a new kind of computer – in each case one that was at the frontier of what was possible. I believe we also shared the view that we should not underrate Australia's potential as a place for commercialising and industrialising computing technologies. These are all themes I will touch upon today.

One of the things that is fascinating about the CSIRAC – and I saw it today Melbourne's *Scienceworks Museum* – is its size. This computer was the state of the art in its day. Built using 2000 vacuum tubes, that period's version of our modern transistors, each big enough to hold in your hand, and stacked together in banks, the machine filled an entire room. I know I'm not the first to point out the contrast with our modern calculating machines, but it is astonishing in comparison to think that the current Apple M1 Max chip, manufactured by TSMC, manages to fit 57 billion transistors onto one little slither of an object about the size of an after-dinner mint.

Computers are not the only technology that has shrunk over time. The first motors were enormous – big, bulky, steam-powered machines originally designed to pump water out of mines. Now you can find miniature electronic motors all over the place: in refrigerators, dishwaters, vacuum cleaners, cameras, and of course in an ever-expanding range of transportation technologies. The first mechanical clocks were also rather grand and expensive; so much so that the earliest mechanical chronometers were largely to be found in significant public places like churches, cathedrals and town halls. Today, anyone can buy a watch and our whole world has become synchronised as a result. Or think about printers. Gutenberg's printing press was twice as high as a man, larger than a grand piano, and arduous to use. Yet now you can pick up an efficient laser printer from Office Works that's not much bigger than the size of a briefcase.

The miniaturisation of computing, then, is not an entirely unique story. One of the ways humans can enhance any technology is to try to miniaturise it. It is not uncommon that, having made something smaller, we find new uses for it, which end up transforming society in unanticipated ways.

The story of computing is unique, however, in scale if not in kind. And I say this not just because of the stark comparison that might be drawn between the CSIRAC of the late 1940s and the wild multitude of conveniently sized and extremely powerful computing devices we all have at our disposal today. What truly sets computing apart from other technologies that have miniaturised is the sheer speed and the scale of the miniaturisation, and the unbelievable fact – as I will explain this evening – that we have now brought this particular process of its ultimate limit, whereby the core features of our computing machinery can actually be reduced to the size of individual atoms.

I'm sure everyone here has heard of Gordon Moore. Gordon Moore was the co-founder of Intel. He's the one who first noted that the number of transistors on a microchip was roughly doubling every 18 months to 2 years – which essentially meant, in practical terms, that the smallest features on these chips were halving in size over the same timeframe. Moore published his data in the 1960s and projected the rate of growth into the future, observing that if any semiconductor company wanted to remain competitive over time it would have to keep cramming more and more transistors into less and less space. In so doing, he set a standard that the entire industry benchmarked itself to, and turned what was then an observation from only a few years of data into a self-fulfilling prophecy that lasted decades.

Thirty years later, in the late 1990s, when I was working in the Cavendish Laboratory in Cambridge and considering migrating to Australia, people would look at these same data on the number of transistors in a microchip and wonder about the forward projection. Because, back then, judging from the speed at which the semiconductor industry was innovating, it looked as if it would be only 20 more years – i.e. by the early 2020s – that these companies would find themselves making devices with features at the level of individual atoms.

In those days, of course, there wasn't any technology that would make such a thing possible. But there were microscopes that had been invented in the 1980s – scanning tunnelling microscopes – that enabled us to visualise individual atoms. I decided that I would come to Australia and see if I could use this technology to leapfrog the projections of Moore's Law and build electronic devices at the atomic scale on an even faster timeline.

It was not a totally stupid idea. In fact, there had been a precedent for it. The scanning tunnelling microscope, which allowed us to 'see' for the first time that atoms exist, has a very fine metal tip that you bring down to the surface under voltage control. When you bring this tip close to a surface under vacuum, a tiny current begins to flow. To image a surface, you simply maintain a constant current and scan the tip across the surface which then deflects in height as it traces over the atoms, giving a height profile which when you raster scan thereby building beautiful topographic image of the surface. But in the 1990s, researchers at IBM had turned this idea on its head. Instead of just looking at atoms on a surface, they'd wondered whether they could use the tip of this same instrument to pick atoms up and move them around.

They famously formed the world's smallest logo – IBM – by applying a voltage to pick up metal atoms on a surface, and then pulsing the tip to drop them off again in a new location. Our hope was that we could start to build things in semiconductors with a similar kind of precision. However, while it is easy to pick up metal atoms on a metal surface, it is not easy to pick up atoms inside a semiconductor crystal. The bonds are just too strong. We had to come up with another, much more complicated, way to make atomic-scale devices in silicon.

Our atomic fabrication method, which is really a product of 20 years' work by my team, did end up using a scanning tunnelling microscope but it also mimics the lithographic process that you would find in a conventional cleanroom like an Intel cleanroom. It works as follows.

- We start by etching markers into the surface of a silicon wafer so that once we've made our atomic-scale device, we will have a mechanism for finding it again. We call these markers 'registration markers'.

- Next, we load our wafer into an ultra-high vacuum of the scanning tunnelling microscope and heat it up to >1000 degrees before cool it down slowly so that we can reconstruct the surface and form those beautiful rows of silicon atoms at the surface. We then introduce hydrogen to the system, which basically covers the whole of the surface.
- Now is the fun part. Using the fine metal tip of the scanning tunnelling microscope, we desorb selected hydrogen atoms from our surface in the same way that IBM did with the metal atoms of its logo. We literally 'write' on the surface to expose the underlying silicon. This gives us a mask that is mainly hydrogen but with gaps where we want them.
- We now want to bring in the atoms that will make up the active components of our device. In our case these are phosphorus atoms – an element with one extra electron than silicon. We supply phosphorus in the form of phosphine gas, which sticks to the exposed silicon, but not to the surrounding hydrogen.
- Some magical chemistry occurs at this point. The formula for phosphine gas is PH_3. When you heat the sample up, the phosphine starts shedding hydrogens. One by one, they go. Eventually, the phosphorus atom, deprived of its hydrogens, drops itself into the layer of silicon below and pushes a silicon atom out.
- We then take the device across to a silicon crystal growth system and using a technology called molecular beam epitaxy, we encapsulate everything with silicon at low temperature atomic layer by atomic layer to make sure that our phosphorus atoms don't move and will remain stable.
- If we want to, as a check, using the scanning tunnelling microscope, we can image the atoms beneath the surface to show that they are still there – exactly where we put them.
- Finally, we can take the device out of the ultra-high vacuum and move it through to a cleanroom, where we use those all-important registration markers laid down at the start, so that we can site contacts to the buried device below.

It took us a long time to figure out this process. It required some of the most sustained and systematic problem solving of my career. Despite having some outstanding students and postdocs working on this with me, it took us at least 10 years – and of course we are still refining things. But it was well worth the effort, because our approach to atomic fabrication culminated in our creating the world's first electronic devices, engineered from the ground up, for which the core components were designed and made using just a few or even single atoms.

Among the atomic-scale devices we have made is the 7-atom transistor, which was picked up by the Guinness Book of World Records, as my son discovered one day to his surprise in the school library. In fact, it was more than just a simple transistor. This was quite an interesting device because it gave us 7 different charge transitions as we removed electrons from the system.

We have also made tiny, tiny wires out of phosphorus dopants. In fact, we've shown that we can pattern a wire just 1.7 nanometers or 17 Angstroms – just a few atoms worth – wide and that it will conduct electrons as readily as metallic copper, proving that Ohm's Law is maintained down to the atomic scale. This is very important as one tries to build more complex atomic electronics, because we need to be able to control charge movement in and out of our systems with low loss and low noise.

Then, of course, there is the single atom transistor – where we designed and built a device with a single phosphorus atom situated between gates also engineered from phosphorus donors in silicon and loaded individual electrons on and off that single atom as if operating an on / off switch. That device was exciting not just because it's a technological landmark. It also enabled us to see what the wave function of an electron actually looks like! Somewhat miraculously, it turns out there's a way of using the scanning tunnelling microscope to do this.

Now, we somehow managed all this a decade ahead of the timeline projected from Moore's Law. But let me tell you something astonishing. There was someone else, even before Gordon Moore's famous paper, who saw all this coming. Back in 1959, just ten years after Trevor Pearcey's CSIRAC was completed, at a time when commercial computers built with transistors were only just beginning to be sold, and still more than 25 years before the scanning tunnelling microscope was even invented, the Nobel Prize winning physicist, Richard Feynman, one of science's great visionaries, gave a lecture entitled *"There's Plenty of Room at the Bottom"*.

The subtitle of his lecture was *"An invitation to enter a new field of physics".* Right at the outset, he stated that what he wanted to talk about was "the problem of manipulating and controlling things on a small scale" and he asked why we shouldn't aspire to make circuits and machines at the level of perhaps 10 or 100 atoms. "I am not afraid," he wrote, "to consider the final question as to whether, ultimately – in the **great future** – we can arrange the atoms the way we want; the very atoms, all the way down!"

What Feynman grasped was that achieving such a goal was a problem of practicalities not principles. He said, "The principles of physics, as far as I can see do not speak against the possibility of manoeuvring things atom by atom. It is not an attempt to violate any laws; it is something, in principle, that can be done; but in practice, it has not been done because **we are too big**."

Not anymore. I first read this wonderful lecture years ago and forgot about it. It was only recently that I rediscovered it and realised just how well we have been delivering on Feynman's vision. Sometimes, what seems like science fiction, actually does come true – even if you have to wait 60 years!

There is one part of Feynman's vision, however, that we're still delivering on. Feynman was a quantum physicist. He won his Nobel Prize for quantum electrodynamics. He was not an experimentalist, but he deeply understood the implications of controlling things on an atomic scale, which is why, at one point in his talk, he says, "When we get to the very, very small world – say of 7 atoms – they behave like *nothing* on a large scale, for they satisfy the laws of quantum mechanics. So, as we go down and fiddle around with the atoms down there, we are working with different laws, and we can expect to do different things. We can manufacture in different ways. We can use, not just circuits, but some system involving the quantised energy levels, or the interactions of the quantised spins, etc."

In this, once again, Feynman was right. For using quantised energy levels and the interactions of quantised spins is our **next frontier**, and the twenty-first century's next frontier in computing. Having figured out, over the last 20 years or so, how to make atomic electronics reproducibly, and with exquisite precision, we’ve now started to bring all the various components together to control quantum effects and to see whether we can exploit the quantised spins of electrons on phosphorus donors in silicon to make a quantum computer.

For those of you who may not have heard the exciting news, a quantum computer, is a new type of computer that exploits the laws of quantum physics, so that instead of performing calculations sequentially like a conventional, digital computer, it works in parallel, looking at many possible outcomes at the same time, and producing an exponential speed up in computational power. The concept is as novel, relative to the digital computing machines that surround us today, as Trevor Pearcey’s computer was to the mechanical calculating machines of the late nineteenth and early twentieth century.

It is based upon the idea that if you can control quantum effects such as the spin of an electron, then you can encode information at the quantum level, in much the same way as a classical computer encodes information onto its transistors. We should note, however, that whereas information on a transistor is termed a “bit”, information encoded in a quantum system is called a “quantum bit” or a “qubit”. These different terminologies are used because there are some fundamental differences between the two paradigms, as I will try to explain.

Consider for one moment an electron. An electron is a quantum particle – and a wave – and one of the quantum properties of an electron is its spin. To a physicist, spin has no analogue in the macroscopic world. However, we often talk about an electron being either spin ‘up’ or spin ‘down’; and, in this respect, we can think of spin as responding to a magnetic field somewhat like a tiny bar magnet. Indeed, it is this property that enables us to use spin to encode information. For example, it can align with a magnetic field, giving the ‘zero’ of classical information; or it can flip to align against the field, in a higher energy state, giving the ‘one’ of classical information.

But there is more to the quantum world than ones and zeroes. The quantum world is far more sophisticated than that. Because of the quantum property known as **superposition**, it turns out that an electron’s spin can be in some fraction of its ‘up’ state and some fraction of its ‘down’ state at the same time. A quantum bit or a qubit, in other words, encoded on an electron’s spin, does not just hold information as a ‘one’ or a ‘zero’ as is the case with a classical bit: it encodes information as a mathematical vector with the fraction specifying ‘zero’ and the fraction specifying ‘one’ providing two additional bits of information.

Now consider what happens when we bring two electrons together, so they become entangled – another quantum property. On its own, each electron will have had its individual spin state, but when they become entangled, they create a new, combined state based around four possible spin pairings: up-up, down-down, up-down, and down-down; or 11, 01, 10, and 00. Once again, though, because of superposition, the entangled pair can encode information as some fraction of each of these four states at the same time. Mathematically, this means that the spins of two entangled electrons, will now hold four additional bits of information.

And this pattern is repeated every time we add another electron. Now imagine that these electron spins are used to store information as quantum bits or qubits in a computer. In computing terms, every time we add another qubit to the system, we will essentially double the amount of information that the quantum processor is able to compute. So, with 2 qubits there are 4 states, with 3 qubits there are 8 states, with 4 qubits there are 16 states, and so on – until by the time we get to 36 qubits, our system contains more information that a classical computer with 57 billion transistors!

Nor is this the only thing that is marvellous about a quantum computer. It should also be said that any computational operation we perform on an entangled state like this will be performed on every entangled element at the same time. This is totally different to classical computing, where we must laboriously work through all calculations sequentially, one after the other. **It enables parallel computation to an unparalleled degree.**

So, Richard Feynman was absolutely right. You can indeed use the interactions of quantised spins to “do things differently”. The quantum phenomena of superposition and entanglement create the potential for a completely novel and uniquely powerful mode of computing; and if you can reliably initialise, control, and measure the spin of individual electrons on individual atoms, you have a superb platform for doing this.

This is where our atomic electronics comes into its own. Our now-proven ability to engineer placement of phosphorus donor atoms into a silicon matrix with sub-nanometre precision creates a perfect platform for exploiting the power of quantum physics in information processing.

Phosphorus has one extra electron and proton compared with its silicon host. It is consequently on the spin of a phosphorus electron or on the spin of a phosphorus nucleus that we encode our quantum information to create our quantum bits. At the same time, silicon is an ideal host environment for our qubits, since it naturally forms a low-noise, crystalline structure that keeps our phosphorus donors protected and in place.

Significantly, too, the technology we have developed not only makes the qubit out of phosphorus atoms, but all the electronic components in the processor too: from the control electrodes that create entanglement; to the sensors that initialise and read-out the information; as well as the amplifiers, capacitors and inductors. Everything, all together on one monolithic chip, is made of just two elements: phosphorus and silicon.

This brings huge advantages because it means we do not have the material interfaces and imperfections that are well known to create defect states and to cause qubits to lose their information or decohere. Working with just two kinds of atoms is simple, elegant and clean; and it enables us to form stable, high-quality qubits.

Indeed, let me share with you two quantifiable examples of just how advantageous an atomic system based on silicon and phosphorus is for creating stable, high-quality qubits.

First of all, with our atomic devices, we can read out the spin state of a single electron in ~1µs, with greater than 99.9% fidelity. (Fidelity is essentially how accurately you can do this.) This is the world record for the fastest, highest fidelity read out of a single electron spin. Atomic engineering is what enables us to achieve this.

- **Part of the advantage here** is that the intrinsic confinement potential of phosphorus atoms in silicon naturally confines the electron spin without the need for additional metal surface electrodes to create our qubit. Just by putting the atom there we create the qubit. As already discussed, this lack of materials interfaces means we can detect the behaviour of a single electron with minimal interference.

- **But there's more to fast read-out than this.** The closer a sensor is to a qubit, the faster it can read the spin state. With atomic manufacturing, we can bring our sensors (which are also made of phosphorus in silicon) very close to our qubit.

- **Finally, being able to engineer our sensors with atomic precision** to have a very high signal means that in detecting the spin we can rely on a process called Pauli spin blockade (based on Pauli's exclusion principle) – whereby we initialise one spin and try to bring the other close to it. This is the process that has enabled us to achieve such exquisitely high fidelity.

As for a second result, which also confirms the advantages of atomic engineering, we now hold the world record for the fastest two-qubit gate. A key operation in computer logic, when we are running algorithms, is to be able to swap information between two bits. This is a reversible interaction called a SWAP gate. In 2019, we demonstrated that by placing our phosphorus atoms ~13nm apart and close to our high-fidelity read-out sensors, we could swap an electron spin from one atom to the next in 0.8ns – the fastest two-qubit gate in silicon qubits. This was achieved whilst being able to initialise the electron spins on each qubit individually, even though they are so close together. Nicknamed *"the Fastest 2-Qubit Gate in the West"*, my students made a fun YouTube video to explain the result – although I insisted that they also publish it in the scientific literature. It eventually came out in the journal (indeed on the cover of) Nature.

These results may seem technical, but they absolutely matter. There is currently a race on in the world to build a quantum processor. Different groups and different companies are all bragging about the number of qubits they have managed to string together. The number of qubits, however, means nothing if you don't also know their quality. What will count in the long run is not just whether you have 50 qubits or 5000 qubits but whether the qubits you do have are of sufficient quality to be error corrected.

Now it might be a surprise to you to know that the world is not perfect and that every computer – classical or quantum – relies on a process of error correction. Even TSMC and Intel, with their high-yield processes, don't guarantee that every one of the 57 billion transistors in a microchip is going to work! To get around this, all computers rely on a process of error correction. In classical computers, this means that processes are run in parallel throughout any computation, with multiple copies adding redundancy. The computer then looks to the answers, discarding those that are different, assuming them to have incurred an error. This process is essential for reliable computation.

In the quantum realm, error correction is even more important, since errors do not just affect whether we are spin ‘up’ or spin ‘down’ but can change the direction of the vector that describes the superposition of our spin state. Quantum systems are fundamentally far more sensitive to errors than their classical counterparts. Moreover, the impact of errors may be greater in some contexts too, given the highly parallel nature of quantum processing. All of which is to say that qubit quality matters – and that the speed and fidelity of our atomic systems will likely prove powerful advantages in the long run.

So where are we now? If you now believe, as I do, that atomic manufacturing is a revolutionary capability, and that this has opened the door to an exciting route for building an error-corrected quantum computer, I am hoping, at this point, that you’ll want to know two things. How far have we got – **what’s our biggest result to date? And what’s coming next?** In the minutes remaining, I will answer these questions.

First, how far have we got? In 2021, less than a decade after our team’s 2012 declaration that we had fabricated the world’s first single atom transistor, we realised our next watershed result: the realisation of **the world’s first integrated circuit manufactured at the atomic scale**.

To make this device, we had to manufacture small dots of 25 phosphorus atoms apiece, so that their energy levels aligned, and electrons could easily pass through them. To give you a sense of how remarkable this is, these dots had to be situated with sub-nanometre precision so that they were close enough but remained independent for the quantum-coherent transport of electrons across the chain. We also had to figure out how to tune the energy levels of each dot individually, and of all dots collectively, so that we could control the passage of quantum information among them. All of this represents a significant feat of atomic engineering.

The really wonderful thing about the device, however, is that it actually calculated something: we used it to run a quantum simulation algorithm that accurately modelled the quantum states of a small, organic polyacetylene molecule. This was a major breakthrough. Today’s classical computers struggle to simulate even relatively small molecules (greater than ~20 atoms) due to the large number of possible interactions between the electrons and atoms within the system. By placing phosphorus atoms within silicon with atomic precision we were able to mimic the single and double carbon-carbon bonds of the polyacetylene molecule and simulate current transfer through the molecule.

I should add that, in designing this device, we deliberately chose a small enough chain of the polyacetylene molecule that we could still classically simulate the expected current through the device in order to validate our results. For those in the audience whose organic chemistry might be a little scratchy, a polyacetylene molecular has alternating single and double bonds. What the quantum simulation on our device showed was exactly how the movement of electrons through the molecule varies depending on whether the molecule is bounded by double or single bonds. In this sense, our achievement wasn't just evidence of our artful engineering, it was proof that atomic circuits can be used to simulate real-world complex systems. It's an approach that may, with a little more development, be used to model novel pharmaceuticals, materials for batteries, or catalysts.

So, what's next? We are currently the only company in the world that can manufacture electronic devices with atomic precision. We have nearly a hundred granted patents with ~50 more pending. These cover each stage of the manufacturing process, different error-corrected architectures, the operation and control of the qubits, and near-term processors from quantum analogue simulators to AI accelerators. And we are now scaling our quantum hardware to take on heavy duty computational tasks that cannot be performed by traditional computers.

To manage this, we must face an entirely new suite of problems – problems of the kind that Trevor Pearcey would have understood intimately. The silicon processor I have been describing, manufactured with atomic precision, will be the fundamental core of our quantum computer and the basis of quantum speed up. It is the most important layer and without a competitive advantage in this layer the computer will not do anything! But above this, and around it, there is an entire stack of associated capabilities that must also be developed.

These include read-out and multiplexing capabilities, which will be provided as a separate cryogenic hardware chip used to initialize and read out information onto the qubit hardware. We will also need high transmission wiring with noise filtering and shielding, connecting the qubit hardware to the read-out chip and to room-temperature control electronics. This routing and I/O hardware is essential to get information in and out of the computer quickly.

Beyond all this, of course, we are also building the necessary high-frequency, low-latency control electronics at room temperature that are required for accurate real-time control and read out of qubit hardware. There is dedicated control software for operating the classical electronics that control the quantum system. There are the all-important software protocols that will be used for minimizing errors, which amounts to our error-correcting layer. And, above all, there are the quantum algorithms (~70 have been invented to date) that will produce our eventual applications in optimization, machine learning, quantum simulation and integer factorization.

I know that sounds like a lot for a small Australian start-up to be taking on. But the prize is large. The Boston Consulting group (BCG) has forecast annual quantum computing revenues at $90B to $170B USD globally by 2040. McKinsey has independently estimated that quantum technologies could create USD $700B annually across all industries. And is our ambition in building a full stack quantum computer so very different from Trevor Pearcey's undertaking with CSIRAC?

I am firmly of the view that Australia can compete at every stage of the global race to manufacture quantum hardware. The atomic-scale manufacturing technology I've described today, developed through the Centre of Excellence for Quantum Computation and Communication Technology, and now spun out through Silicon Quantum Computing, is globally unique. Yes, a lot of other alternative methods are being pursued in the quantum computing field – ion traps, superconducting, diamond, organics and Majorana systems to name a few. But none have yet been manufactured at scale; and with world record speeds, fidelities, spin lifetimes, low gate densities and device stability, we are already creating the highest quality quantum processors here on shore.

We're in this incredible position because we got in early. We now hold key patents and know-how, creating barriers to entry for others. This mimics America's twentieth century advantage in classical computing, where they invented the transistor and the integrated circuit, and then controlled the early industrialization of these products.

As happened in the early phase of classical computing, the hardware is where the greatest value will lie and where the spill overs will be greatest. That's because success in hardware manufacturing will attract an associated industry of component manufacturers, control system providers, and app developers – whereas the opposite will not apply. With a globally leading Australian onshore manufacturing capability, we can build a vibrant and multifaceted quantum computing industry here; while without such a capability, the Australian quantum computing sector will end up just a satellite to an industry that is largely based elsewhere.

There is also a window of opportunity that works for us right now, but which won't stay open forever. In their first incarnations, quantum computers will be high-end, high-value, highly specialized products, manufactured with high precision in low volume and at high margins. This can be done from Australia. It is not a mass-market consumer product requiring high volume, cheap manufacturing and economies of scale.

I am also confident that with the right product we can find customers despite our distance from global markets. As with early mainframe systems in classical computing, users will access quantum computing hardware remotely – in our case via 'quantum as a service' from the cloud, thereby eliminating the tyranny of distance from which so many past Australian technology operations have suffered.

Finally, in 2023, I shouldn't need to add that we have the people to make it happen, because that should be obvious. In the quarter century that I have worked in Australia I have never had a team so good as the one working with me in Silicon Quantum Computing and at UNSW today. They've come from all over the world, and they are extraordinarily talented and dedicated. And for the latter, I can't blame them.

Together, we are working at the very frontier of what is possible. We have figured out how to manipulate individual atoms and how to store information on a single electron. I have to pinch myself every time I think of that! We are also part of a grand quest, one that involves some truly great scientific forebears: people like Richard Feynman and Wolfgang Pauli and, yes, Trevor Pearcey. And if we can succeed with our ultimate goal –

the error-corrected quantum computer – **well, not just Australia but the whole world will benefit.**