Australian Universities International Alumni Convention
Singapore, 10-13 June 2008

"Singularity – The Way Ahead"

Plenary lecture by Alan Finkel, June 13th, 2008

1. Preamble

In the early 1800s, Dr Dionysius Lardner, a professor of Natural Philosophy at the University College London predicted, "Rail travel at high speed is not possible, because passengers, unable to breathe, would die of asphyxia."

In 1883, Lord Kelvin, the famous scientist after whom the scientific temperature scale is named, proclaimed, "X-rays will prove to be a hoax."

Despite these spectacular failures to predict the future, I am going to give it a try. After all, if I am still around fifty years from now, if I am proven wrong I will be too delighted by my personal longevity to be upset.

2. Thesis

My presentation today is based on four main assertions.

The first is that over a period of fifty years, improvements in telecommunications technology will be massive and will have a huge impact on education.

If you think of the improvement from which you benefitted by upgrading from a dial-up modem to a broadband connection, this will give you a clue as to the kind of improvements you will enjoy during the next fifty years, many times over.

My second assertion is that the conceptual ***manner*** in which technology will be adopted for use in education will be dramatically different to the ways in which we have used technology so far.

These revolutionary shifts will be much more far reaching than the revolution that took children from playing outside in their parent's gardens to playing computer games indoors.

My third assertion is that artificial intelligence will advance to the point where computer programs will finally become intelligent, so much so that they will transcend human intelligence.

These superhuman-intelligence programs will be able to replace human teachers in our education system.

My fourth assertion is that the previous three will combine to radically alter the way that teaching is delivered and learning is experienced.

Giving consideration to how education will work fifty years from now is ***not*** an idle exercise, even if it has little chance of proving to be accurate.

It is important, because in order for universities to be at the forefront, in order to be ultimately successful, there is no time like the present to be preparing for that future and testing the waters.

3. Singularity

In their most extreme manifestations, the changes I have referred to will lead to a "singularity" in higher education. What, you may ask, is a singularity?

In the natural world, a singularity is a place where the usual rules of physics and mathematics collapse.

Consider the Big Bang at the origin of our universe. In a billionth of a billionth of a billionth of a billionth of a second, the universe expanded from a dot, to a fiery ball of matter 100,000 light years in diameter.

This, of course, is completely impossible according to the rules of physics that apply today. The Big Bang was a singularity.

In the realm of human existence, the term singularity was first popularized by the American mathematician Vernor Vinge.[i]

In 1993, Vinge predicted that the accelerating rate of technological progress will give us the means within several decades to create superhuman intelligence. He went further than this by stating, "Shortly after, the human era will be ended."

This concept has been echoed with slightly less scaremongering by many other futurists. The American author Damien Broderick stated that we are developing technologies that will fundamentally transform our existence.

Taking Broderick's lead, look at a graph of technological knowledge increasing with time. Assume, as many have claimed, that our technological knowledge doubles at regular intervals.[ii]

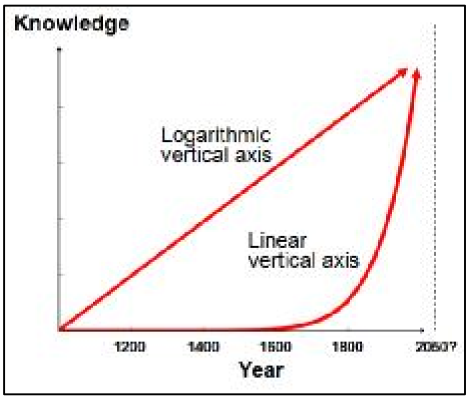

On this basis, the line starts out looking flat, not far above zero, but over the last few hundred years it breaks away from the baseline and curves steeply upwards at a dramatic rate.

Depending on one's perspective, the line appears to reach a vertical limit, corresponding to a singularity in which much of what we know and understand will be overturned.

This growth curve that we are looking at is plotted on a linear y axis.

More often, growth curves are plotted on a logarithmic vertical axis. That is, the information on the vertical axis is compressed so that more details can be seen. On a logarithmic axis the growth curve is a straight line.

Growth based on regular doublings shows up in many areas of technology.

For example, the growth in the number of transistors per microprocessor doubles every 18 months, following the postulate of the so-called Moore's Law.

You can see steady doubling in the growth of the phone industry, with the graph here showing the growth in revenues and phone calls in the United States.

4. Telecommunications

As I already prefaced, if there is to be a singularity in higher education it will be predicated on a number of converging factors.

The first will be a massive increase in telecommunications speed.

Like other technologies, you can see regular doublings in the volume of internet data traffic. There is every reason to think that this solid growth will continue.

Such improvements in technology are already having a significant impact at universities throughout the world. At many, or perhaps most, universities, lectures are available on line. This wasn't possible when I went to university, I couldn't even get a photocopy of the lecturer's notes let alone online access.

Today, video recordings of lectures are routinely available. We might take it for granted, but distributing lectures by streaming video is a revolution made possible by broadband internet.

One limitation is that streaming video today is grainy because the speed of the internet limits the video to half the resolution of standard broadcast television.

Here is a picture of Jane Goodall, the famous primatologist, captured from a video lecture that I downloaded from the Harvard University web site. As you can see, it is low resolution.

But not for long. Within a few short years in Australia, we will enjoy the benefits of a National Broadband Network that will make data speeds of at least *12* megabits per second available to 98 per cent of the population.

This is more than fast enough for real time streaming video of lectures at standard definition television quality.

Leap forward another ten years and there is likely to be a *100* megabit per second broadband network in Australia.

100 megabits per second is more than enough for real time streaming video of lectures in *high-definition* television format.

Video lectures will be better quality, providing a superior sense of presence, as can be anticipated from this photograph of Jane Goodall taken at the resolution of high-definition television.

In fact, 100 megabits per second is sufficient for providing holographic videos, in which the lecturer can be standing in three-dimensional glory in the house of the student.

This slide is an image of Ray Kurzweil, inventor of the music synthesizer, addressing an audience at a convention last year in Canada. The unusual thing is that Kurzweil wasn't really there, this picture is a three-dimensional hologram, projected in real time.

However, such increases in internet speed will ***not*** revolutionise teaching. Increased internet speed will merely improve the visual quality of what a student off campus receives.

Are there other, foreseeable improvements in technology that will improve the ***quality*** of teaching? According to many of the lecturers whom I have asked, the essence of good teaching is interaction with the students.

On this premise, if technological improvements can enhance interactivity they will help teaching.

My son agrees. He is currently studying engineering at Monash University. I asked him what makes it worthwhile to actually attend a lecture instead of viewing it offline.

The answer is that it depends on whether the lecturer interacts with the students.

Bad lecturers do not make eye contact. Instead, they monotonously write on the white board.

Good lecturers engage the students, they tell jokes and stories, they take questions from the students, they ask questions of the students and they proceed at a pace assessed by reading the mood in the lecture theatre.

Technological improvements can assist lecturers to maximize interactivity. For example, a recent enhancement in some lecture halls is the provision of keyboards or tablet computers with which students can send messages and responses to the lecturer's console.

If one student in a class of 200 sends a "don't understand" message the lecturer might proceed, but if fifty send a "don't understand" message a good lecturer will get the hint and try again.

In another example of today's leading edge, in some of Monash University's science and medicine teaching laboratories the benches are equipped with two-way video that enables the supervisor to see, at a central location, what is being magnified at every microscope. The supervisor can then advise the students on how to improve the image quality.

These interactive technologies can help even if some of the students are in an adjacent hall, because the students can respond to questions and post their own queries even though they are watching a television representation of the lecturer.

Similarly, interactive technologies can be used for non-campus students at home or in a remote community hall.

Nevertheless, there is a natural limit to how much activity a single lecturer can manage, and hence a limit to the maximum practical class size.

Perhaps for a given lecturer incremental advances in the technology I have described will help the lecturer to teach 300 students instead of 200. But no amount of technology will make it easy for the lecturer to deliver the same lecture more times, or deliver it at night after a long day of research.

Other advances in computer processing power might solve that problem, but before discussing them let's continue to pursue the issue of interactivity.

Considerably higher degrees of interactivity are possible, which brings me to my second assertion.

5. How technology is used will change dramatically

That is, the *manner* in which we use technology will change just as dramatically as the *capacity* of technology.

The signs of these pervasive changes have been there for a while. In testimony to the US Congress in 1995, Professor Chris Dede said that "the most dangerous experiment we can conduct with our children is to keep schooling the same when every other aspect of our society is dramatically changing."[iii]

Some of the more innovative university departments around the world are taking heed of Professor Dede's warning. They are actively incorporating the latest changes in our society into their teaching methods.

One of these changes is a virtual reality world called Second Life.

Second Life is an experience, in which participants are immersed in a magic world.

In Second Life, you can talk to vampires, angels, cat people and super heroes.

In Second Life you can fly, you can build machines that defy the laws of physics, you can indulge in commerce or you can enjoy an exotic vacation.

Just as its name indicates, Second Life is a fantasy world that is an alternative to the real world.

Second Life is not a computer game because there are no measurable objectives. In a game, the goal is to beat an opponent. In Second Life, the goal is to participate.

The excitement for participants in Second Life is that they get to interact with other people, who, because they are represented by fantasy figures known as Avatars, have extroverted personalities that make them more exciting than their real-world counterparts.

At the Monash University Faculty of Pharmacy, some of the lecturers have purchased an island in Second LifeThey have named their island “Pharmatopia”.

Together, they are building Pharmatopia as a place where students can have a virtual, interactive educational experience. The students will immerse themselves in activities such as running a pharmaceuticals manufacturing plant or heading up a clinical practice.

In Pharmatopia, the students will mix chemicals without fear of bodily harm if the mixture explodes.

In Pharmatopia, the students will dispense prescriptions to ailing patients without fear that their first attempt might end in tragedy.

In Pharmatopia, the students will be able to miniaturize themselves and be injected into a patient so that they can see firsthand how the medicinal molecules interact with the proteins in the patient’s body.

The immersive system of education possible in Second Life will enable students to appreciate the look and feel of science, technology, history and the arts in ways never before possible.

Imagine if you are studying Ancient Rome and you have the opportunity in Second Life to mingle with senators in the Forum, to sit in the stands at the Colosseum, or to watch Julius Caesar lead his Legions into battle.

Imagine if you are studying architecture and you have the ability to build a virtual house, walk its corridors and invite your fellow students inside for a cup of virtual coffee.

Monash is not the first in this kind of endeavour, although it is on the leading edge. Edinburgh University has been using Second Life since 2006 to teach informatics, in two virtual buildings that they have designed themselves.


THE FACTS
New university block comes to Second Life

The cost? A few thousand US dollars, equating to several million Second Life dollars, and a substantial amount of programming time by staff and students.

The biggest problem for the lecturers in a Second Life class is not that the engineering students will be throwing paper aeroplanes. In a virtual world, the uncooperative students will be creating a disturbance by flying *themselves* around the lecture hall.

In a virtual world, tutorials can be held outside, around a campfire.

In a virtual world, the students, represented by their avatars, can sit at a table talking to tutors sitting in the real world.

And, somewhat bizarrely, in the real world, the real-life tutors can sit a table teaching student avatars who are sitting in the virtual world.

The biggest mistake of all would be to dismiss Second Life as a toy.

In a recent article, Bill Joy, the founder of Sun Microsystems, said, "We, sadly, spend much more of our collective energy and focus on virtual reality for *entertainment* than for *education*..." [iv]

Corporations have noticed the growing interest in virtual worlds and many large companies have paid to establish a corporate presence in Second Life.

These virtual worlds are an opportunity, not a distraction.

With these kinds of changes to the *conceptual manner* in which we use technology, the experience for the student and indeed the lecturer can actually improve.

Can you imagine the new ways of using technology that educators will develop over the next fifty years? It is a certainty that students will be privileged to enjoy an immersive educational experience that is far better than what you and I experienced and indeed, far better than what we can currently imagine.

This will make a difference to on-campus and to off-campus students.

Take the latter. Today it is well known that students resist distributed learning, it is too lonely. However, if they can immerse themselves in an alternative world, where they can interact with lecturers and other students in a way not currently possible the loneliness will be diminished.

Still, these changes do not alter the fundamental limits to the student-to-staff ratio. A lecturer is a human being, no amount of technology can enable him or her to deal interactively with more than a few hundred students at once.

Which brings me to my third assertion.

6. Artificial Intelligence

This is, that artificial intelligence will advance to the point where computer programs will finally become as intelligent as human beings, and then more so.

This will open up the possibility that intelligent computer programs will be able to replace human teachers in our education system. At that point, education will be totally and dramatically revolutionized.

There are many ways that advanced computer intelligence might develop, but the only one that I will mention today is based on learning from mother nature.

That is, we might be able to find the secret to true artificial intelligence by studying in exquisite detail how the human brain works. This approach intrigues me firstly because my background is in neurosciences, and secondly because this approach holds out significant promise.

By studying the human brain we will learn enough that we will be able to write programs that will mimic the human brain in all its glory: intellect, creativity, emotions and the supreme characteristic of all – human consciousness.

To understand how our brains work we need to learn how every single brain cell is connected to its colleagues, and how each brain cell operates internally and communicates with its neighbours.

In the world of engineering, the process of discovering how a complex piece of existing technology works by examining its microscopic structure is called reverse engineering.

The goal of reverse engineering the human brain is to create the wiring diagram, and a map of the types of connections between cells, for the whole of the brain.

The task is enormous because of the complexity of the brain.

The human brain weighs a mere 1.3 kilograms, yet it contains 200 billion brain cells.

Each cell is tiny, about one hundredth of a millimetre in diameter, much too small to be seen without a microscope.

Each brain cell connects on average to 5,000 other cells, therefore the brain contains 1,000 trillion connections between cells.

As Irish zoologist Lyall Watson once said, “If the brain were so simple we could understand it, [then] we would be so simple [that] we couldn't.”

To reverse engineer the brain, scientists commence by dissecting out a small piece of brain tissue. The cube in this picture is a gross exaggeration of the size of the tiny piece that they start with, given the limitations of today’s technology. The cube’s volume is approximately a billionth of the volume of the full brain.

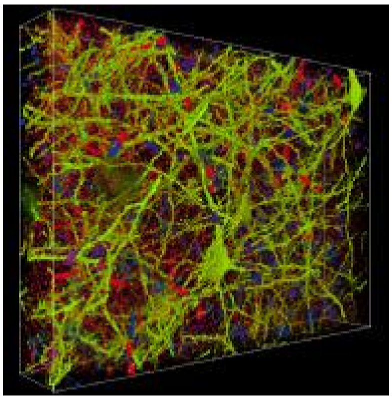

Using advanced dissection and imaging techniques, the scientists can create a computer image of the cellular structure and connections in the piece of brain tissue.[v]

As you can see, it is extremely complex.

As powerful as supercomputers are today they are limited to analysing these extremely small pieces of brain tissue.

In their test case, the tiny pieces of brain tissue that researchers are working with only contain 10,000 cells.

To create a map of all the connections and run it as if it were an electrical circuit, scientists are using an using an array of four IBM supercomputers, each of which costs US$2.5 million.

To do the same thing for the whole brain would require 200,000 of these supercomputers. This is clearly impractical, but according to Moore's Law, a *single* computer with this much power will be available within 25 years.

At the same time, the volume and resolution of brain tissue that neuroscientists can image is also increasing at a rate that seems consistent with Moore's law.

Nowadays, MRI scans of the brain can resolve components as small as 0.1 mm, but just thirty years ago scans were limited to several millimetres.

These kinds or rapid improvements in biological technology are not limited to neuroscience.

Here you can see that the amount of gene sequence data stored in the major database for scientists – Genbank – continues to grow at a rapid rate, and has been doing so for more than twenty years.

With this kind of computing power and associated neuroscience and genomics knowledge, it will be possible to reverse engineer the entire human brain.

If you look at a graph of computer calculations per second, it becomes apparent that computers will surpass the notional number of calculations per second that a human brain can perform by the year 2030.

But hardware improvements alone are not enough to create superhuman intelligence. In addition, the software has to be written. The advances that will derive from reverse engineering the human brain will guide the development of the software in the same timeframe as the hardware.

I cannot think of anything more difficult than the task of reverse engineering the human brain. It is the biggest challenge facing scientists for the foreseeable future.

The big question is, once we know all the connections, and the chemical and electrical activity underlying every connection, and we simulate all of this in a computer program, will superhuman intelligence emerge? I believe that it will.

The knowledge from reverse engineering the human brain will help us to heal minds and make truly intelligent machines. It will give us an objective insight into the mysteries of life, a tangible understanding of intelligence, consciousness and those other wonderful characteristics of our minds that are the essence of what makes us human.

The knowledge about artificial intelligence derived from this effort will also have a profound impact on education. How?

7. Combining Telecoms, Virtuality and AI

Well, as already stated, the biggest limitation on the number of students in a lecture is the ability of the lecturer to interact with the students.

Even with the interactive technologies I described earlier in my talk there is a clear limit to how many students a lecturer can manage.

A practical limit might be 300 students. Let's say an existing lecturer is lecturing to 300 students, 200 of whom are in the lecture hall and 100 of whom are participating off campus.

How do the students off campus actually know that there is a real lecturer presenting the material and answering their questions?

Consider now the advances that will come in artificial intelligence. What if the lecturer were replaced by a program?

A program that has all the skills necessary to teach engineering, or multimedia or even psychology. A program that has the ability to ask questions and answer questions.

Run the program on a powerful enough computer and there is no reason why the program would not be able to teach five hundred students. On an even more powerful computer, the program could teach five thousand students, or five *million*.

It would not even have to teach all these students simultaneously. It could interactively teach them 24 hours per day, 365 days per year, responding to the study habits of students in any of the time zones of the world. Some students might study at night because they have day jobs, while others might study during the day because they are full-time students.

The students taking the course might be working from home. If they are lucky, they will have a private study cubicle where all four walls are video screens, so that they can see their virtual lecturer in front of them and their virtual student colleagues to the left, the right, in front and behind.

They will be learning in a simulated world, possibly called Third Life, Fourth Life or Fifth Life. They will be connecting to their virtual education world using the technology of the day, possibly called Web 3, or Web 4 or Web 5.

Imagine a student in a total-immersion study cubicle, wherever it might be. The subject being taught is economics. She is struggling with a statistical analysis. Virtual students on either side of her are working on the same problem.

The virtual lecturer walks up to the front of her desk, asks how it's going. She says she is not sure how to proceed, so the lecturer offers some advice. She thanks him for his assistance and she completes the assignment.

She hands in her work. Seconds later, the marked paper is returned to her.

She reads the comments from the virtual lecturer. The lecturer's remarks constructively criticize the style of her answers to the essay questions, pointing out where her arguments are weak, complimenting her for the quality of her writing, and suggesting how she might do better next time.

"It's wonderful getting instant feedback", she thinks to herself.

For many, this would be a quality experience.

For others, the lack of real human contact would be <u>intolerable</u>.

During the next fifty years, the world will split into those who find virtual teaching acceptable for reasons of cost, convenience or temperament, and those who demand human contact despite the cost and the commitment.

For these reasons, I imagine a future in which there will continue to be thousands of conventional university campuses world-wide.

These universities with a physical presence will cater to the students who have less need to work in jobs and are socially oriented. These universities will also be the seat of university research.

But for hundreds of millions of students worldwide, a virtual university will suffice. A mere handful of universities will each provide high quality education to millions of students.

Education will be provided in the style and at the time that suits the students, provided by artificial intelligence programs masquerading as caring, intelligent, stimulating and interactive lecturers.

Universities of today have to decide whether in fifty years from now they want to be one of the thousands that operate in the real world, with physical campuses, or alternatively whether they want to be one of the handful that will each offer global, virtual, computer-driven teaching to millions of students.

Some universities, of course, will opt to be both.

Given that computers will be so intelligent and so capable, what kind of jobs will there be for the graduates? Why bother going to university, real or virtual? The answer is twofold.

First, there will always be jobs, we simply are not in a position to imagine what they will be in a future world where machines have superhuman intelligence.

Second, humans will always need to learn because learning for curiosity is a human need, and learning in order to be more competitive is another human need. Whatever our lives might be like, we will want to study even if we are taught by computers with superhuman intelligence.

I don't believe that human life will become irrelevant. Vernor Vinge is wrong. Not about the explosion of superhuman intelligence that will constitute a singularity where life as we know it will cease to exist. In that he is right.

Where he is wrong is in his statement that this will constitute the end of the human era. Humans as a race strike me as being infinitely adaptable. Our needs will change. Jobs in fifty years will be in fields we cannot imagine today. Teaching will be inconceivably different.

Despite the fact that we will have built a race of computers with superhuman intelligence, the role of humans will remain meaningful.

The challenge for each of you in the audience today is to make sure that fifty years from now the role of your university will remain meaningful.

To ensure success, there is no time like the present to start taking baby steps towards that future.

Thank you.

---

[i] Vernor Vinge. The coming technological singularity, 1993. VISION-21 Symposium sponsored by NASA Lewis Research Center, and Whole Earth Review, Winter issue.

[ii] Thomas Fuller, 2007. http://newsfan.typepad.co.uk/does_human_knowledge_doub/increase_in_knowledge_workers/index.html

[iii] Chris Bede, Testimony to the US Congress, House of Representatives Joint Hearing on Educational Technology in the 21st Century Committee on Science and Committee on Economic and Educational Opportunities, October 1995, http://www.newhorizons.org/strategies/technology/dede1.htm

[iv] Bill Joy, "The Dream of a Lifetime" in MIT Technology Review, August 2005, http://www.technologyreview.com/Infotech/14695/page5/

[v] Brain image reconstruction courtesy Stephen Smith, Stanford University