[Screen begins with **Professor Joyce Poon's** first slide – fades to **Professor Farid Najm** speaking]

Professor Farid Najm: My name is Farid Najm, I am a Professor and Department Chair for Electrical and Computer Engineering and I'm happy to welcome all of you back to the university and to the department.

So as you know we are a big department we have about 80 professors at this point and then about 20 other emeritus professors, meaning retired but possibly still doing research. So we have expertise in all areas of electrical and computer engineering technology. And on the research side we are growing still. This year we are hiring. We expect to hire three or four new faculty members. One or two of these will be in the area of machine learning whether this is from a computing perspective or a communications perspective will depend on who we interview. We'll also be hopefully hiring someone in the area of systems control with a focus on robotics. Of course the connection between robotics and learning systems is not far in the future. We're also searching in the area of power systems, which sounds like a traditional area and why would we be hiring there? Well with the integration of renewables —solar and wind and similar— the power grid needs to be better managed to maintain stability and so there's a need for expertise in the area of protection and control of the grid once you start hooking up all kinds of generation at the local level. We've actually been trying to hire people in that area but they are very much in demand at industry. So the industry has been vacuuming these people off the market. So it looks like in the pile that's in hand this year there's a couple of good files so we have high hopes that we'll get someone good. And not surprisingly it's someone from the US with an Iranian sounding name; many people are fed up and looking to leave.

Our professors are working on many areas of research and we were just chatting about machine learning back there so we have lots of people at this point looking and interested in machine learning. Again, from both a communications and a computing perspective, but also from a systems control perspective. We can talk about that offline if you're interested. And we're also moving in the direction of neuromodulation. We have a new institute of neuromodulation that's been approved and ready to start next year which will include membership from both engineering, neuroscience and surgery departments.

We are continuing to refurbish and improve our research and teaching labs so you will have heard me speak before, in previous years, about the power systems lab that was completely renovated. All the old equipment from the 1950s was craned out and disposed of, including a big generator that was a big as this table [gestures to large table in front of him]. I don't know how they pulled that out of the basement of Galbraith, but we have a new installation there that's really better than anything I've seen anywhere, in the US or in Canada as far as a power systems lab for undergraduates. This year we're adding a piece to that, we're adding a DC Mircogrid. Again because the renewable sources of generation are typically DC, such as solar. So there's a need to build up experience and expertise in that and figure out how you connect that to the rest of the grid. And so we're putting solar panels on the roof, wiring them down to a large battery in the basement of Galbraith and hooking up a DC Microgrid off of that in the lab from which we will power LEDs for lighting instead of the big, hot lightbulbs that we have down there. It will be both a research and a teaching installation and there will be an LED in a public space outside the lab that shows the power transfer in real time between the solar cells— the solar panels— the battery, and the rest of the system. We are very much excited about that. This will be supported in part with government support but also we are looking for other sources of support for that. So if you're interested in that, feel free to chat with me after this.

On the teaching front I can tell you that our students today are very good. I'm sure many of you are very good, were very good, but you'd be surprised today at the incoming averages from high school. With all the demand for computing and high tech, with the excitement of the app economy around us, everybody wants to write code, everybody wants to write software. If you look at computer engineering, the average incoming GPA from high school admitted this September is around 94% which is amazing, but it's also horrible because we're turning away people with 91 percentile. How do you tell someone that they can't be admitted? But what it means is that we have attention from a lot more students than usual, essentially, who want to do engineering, so we're able to be that selective. And students are, again, doing great and one of our students, Sandro Young, was last year chosen as the best student as far as academic achievement across the whole three campuses of U of T, not just engineering. So our students are making a splash and our alumni are making a splash as well.

This year ECE alumna Julie Payette was named Governor General, you will have heard about that. And last month, ECE alumna Vivienne Sze, who is now a professor at MIT received an Emmy award for her research on video compression technologies. She and the team she was heading received an Emmy award, which is rare for an engineer. So we continue to be firing on all cylinders, if you wish.

Speaking of firing on all cylinders, I'm very pleased to introduce **Professor Joyce Poon**, who will be our speaker for tonight. Joyce is one of our top notch researchers and professors in the department. She holds the Canada Research Chair in Integrated Photonic Devices. She got her PhD and Master's from Caltech, in 2007 and 2003 respectively. Her bachelor's is from EngSci here at U of T in 2002. This year she received an ECE Department Teaching award and was named as a Fellow of The Optical Society. She has a McCharles Prize for early research distinction in 2013, and the IBM Faculty award in 2010 and 2011. Please join me in welcoming **Professor Joyce Poon.**

[Applause]

Professor Joyce Poon: Thank you. Thank you Farid for the very kind introduction and thank you all for braving the weather to be here tonight. For the next 30 mins or so I hope to take you on a journey through the heart of electrical and computer engineering.

I would say that for the past half century or so the most significant contribution of electrical and computer engineering has been the information age and at the heart of the information age is this little device called a transistor. The transistor is a voltage controlled switching element from which you can construct micro-electronic circuits. In 1962 [coughs] —sorry I have a bit of a cough so my apologies—in 1965 Gordon Moore, who was the founder of Intel— who was previously at Fairchild, which became Intel— he predicted that the number of transistors on a microchip would double about once every two years. And at the same time that the sizes of these transistors, which is denoted by the gate length, the length of the gate of the transistor, would correspondingly shrink. This trend indeed did happen for the past 50 years or so, the number of transistors that are integrated on a chip kept growing exponentially, and the transistors themselves also got smaller and smaller. Today the length of these transistors is only about 10s of nanometres. So these transistors have become extremely small so it became difficult to manufacture them.

'More Moore' and 'More than Moore'

So what has happened is that now we are not really sure what to do, we have some fundamental limit with these transistors with only 10s of nanometres in width become extremely difficult—it's difficult to imagine what could happen, what would be necessary to keep moving the field forward. In about mid-2016 the International Technology Roadmap of Semiconductors (ITRS), which is a consortium of semiconductor experts, they produced a final report to explain or project transistor scaling in the semiconductor industry. So they forecasted how transistors would be up to the year 2025, 2030 or so and then beyond that they don't know what to do. So the fact that they issued the very last report after decades of work is really signalling the end of Moore's Law as we had known it.

So what to do? What does this mean? In the traditional way of shrinking transistor sizes, this is called 'More Moore' scaling, so the transistors get smaller and smaller and if we wanted to go to even smaller dimensions or to come up with something different then we really have to think of devices to replace the transistor. This type of More Moore scaling, looking beyond 2025, 2030 or so is called beyond CMOS type of scaling where we have to think of new materials, new devices that are on the nano-scale. At the same time we can use existing technologies but address different types of applications so rather than coming up with a transistor replacement we could maybe use older generation nodes, but address different types of applications that were not previously addressed with just microelectronics. This is called More than Moore scaling, it's about diversifying the semi-conductor industry. Basically More than Moore scaling means that we want to use the existing manufacturing infrastructure, which is extremely expensive —this is the infrastructure that is used to make microchips— make new types of chips that could address different application spaces. And if we think really broadly about what beyond CMOS means this is also kind of beyond CMOS, but in this case thinking about beyond CMOS microelectronics. The third way in which the industry could sustain itself is moving along this axis to combine this More than Moore, More Moore scaling to make new types of microsystems, system on-chips, system in packages, and if we think much more broadly about what this could mean, this could mean creating brand new types of microsystem, new types of computers, new paradigms for computing. And this is also kind of beyond CMOS, in this case going beyond digital computing. So in the next 20 minutes or so I will talk about how the research in my field and the research that my group is doing is going along these three axes: More Moore scaling, more than Moore scaling, and also thinking about new types of micro-systems.

More than Moore

The first example I'd like to give is a More than Moore scaling strategy and that has really been very exciting in my field of expertise, which is photonics. I'd like to introduce you to very new developments that are happening along the lines of foundry silicon integrated photonics. So what is this? In foundry silicon integrated photonics we're trying to use the existing semi-conductor manufacturing infrastructure to make photonic devices and photonic circuits. This way of implementing or making photonic devices and circuits really emerged really I'd say in the past 10 years or so, this is quite new. Over the past decade there have been some R n D foundries—foundries are chip-making factories, giant cleanrooms that process these wafers—so the R n D foundries that are popping up around the world: in the US, Singapore, France, Belgium. There are also commercial scale foundries, some might recognize these names like Intel, Google Foundries (which used to be IBM), also TSMC in Taiwan, they're all using their equipment, which was used to make microelectronic chips, in order to make photonic optical devices and optical circuits. What is the opportunity here? The exciting part about trying to make silicon photonic chips using a foundry is that you could make devices and micro-photonics chips on very large silicon wafers. So these wafers are eight inches or 12 inches in diameter, you can make many, many, many chips on one of these wafers. So that's a way of reducing the cost of photonic chips. And in the past, in my world of photonics, we don't use wafers this big, we use a lot of exotic materials; we use compound semi-conductors, we use lithium niobate, we use all kind of weird stuff and these wafers are just not this big, they will be three or four inches or so in diameter. And we make very few devices, very few chips at a time. But being able to make photonic devices and components in silicon allows us to make, literally, tens of hundreds of dies in one single wafer and that was really un-thinkable before. Another really interesting thing about trying to make these photonic devices in silicon is that the waveguides—which are the structures that actually confine the light, they're like the wires for light—in these platforms in silicon they are very small. The width of one of these optical waveguides in silicon is only 500 nanometres (nm) or so, and the height is only about 200 nm. This is about 100 times thinner than a hair, really tiny, and because it's so small light is very much confined in nanostructures you can have very high density of integration. So you can make very sophisticated photonic devices and photonic circuits and you can bend these light guides very tightly and you can build very complex chips all in a very small footprint. We could also imagine integrating photonic chips now that are made on silicon with other technologies such as microelectronics, maybe fluidics and so on. And we could also use —because there is now so many different chips that can be made in every single wafer run— we can really imagine bringing photonics in a really very wide range of applications, in to new domains in which we do not currently use integrated optics.

The motivation for these foundries, especially the commercial ones, to get in the game, to actually try to produce photonic chips is that there is a great demand for bringing optical communication to shorter and shorter distance links in to computers. So in the past (or even present day) the backbone of the internet was a fibre optic communication network. So optics is only used for very long distance communication, say between continents or between cities and so on, like hundreds of kilometres, thousands of kilometres, we use optics there. The reason we use light is that light has the highest information carrying capacity compared to any other means of communication that we know but it's very expensive to build optical chips. What has happened is that in the computing world, whether we're talking about cloud computing, datacentres, or supercomputers, or very high end microchips, is that there is also a lot of data that has to be shuttled around, say between servers and the datacentre or between a microprocessor and the memory and so on, that there is a great need to find more energy efficient ways to carry this information at a very high speed. So there is great interest to replace connections between computers that used to be electrical cabling with some fibre optic cables instead. And in these types of systems, in datacentres or maybe even chip-to-chip communications, the board, maybe even in the future a microchip, a computer processor itself with an optical communication network, in all these applications you will need a lot of optical communication links much more so than the ones that links up our continents. So there's a great need to find a way to manufacture, at very low cost, mass production of these photonic devices that can generate light, detect light, and modulate light (turn light on and off) and that's why there's a great interest to use silicon photonics to address these upcoming needs in computing.

In my group we've done a number of demonstrations and worked on a number of projects related to silicon photonics for communication, data communications and I just want to highlight a few example projects that I did with some of my colleagues in ECE in this department. In this first example, this is a collaboration with Professor Peter Herman who is in the Photonics group and he is an expert in laser machining, laser manufacturing and together we demonstrated an integration strategy to put a very special kind of laser, called Vertical Cavity Surface Submitting Laser, on to a silicon photonic chip and to have it quite low loss coupling between this light and the chip as well as having a means to kind of stabilize the polarization. In another example we collaborated with Professor Hoi-Kwong Lo, who is an expert in Quantum Information, to demonstrate the world's first electro-optic transmitter for quantum key distribution. So he' a cryptography expert, particularly using the properties of single photons in order to generate a secure key between a sender and a receiver in a communication link. We demonstrated the world's first quantum key distribution transmitter for light, for photons on silicon.

In our final example, this was done with Professor Sorin Voinigescu who is right here today, he's an expert in high-speed electronics and together we demonstrated a three dimensionally integrated transmitter where his electronic chip is bonded on to our photonic chip and together this little assembly was able to produce optical modulation at 340 gigabit/second— very high speed—with a record on to off ratio (1 to 0 light ratio) of 6.4 db for this class of microchip (a very high end CMOS note, 20 nanometre fully depleted silicon on insulator technology).

Moving away from that, of course we really enjoy addressing these data communication types of application but my group also has great interest to come up with a strategy to implement extremely large-scale photonic integrated circuits. These would be photonic circuits that have hundreds, even thousands of photonic devices and that would be considered to be very large in my field. And these types of very large and complex photonic circuits with hundreds of thousands of components could be useful, for example, in optical switching, which I will show you in a moment. This type of project/goal has some limitations, it becomes very difficult to accomplish if we look at generic silicon photonic platforms. Generic platforms mean the common stuff that's available today. If you just signed up to make chips with a foundry this is what you might see. This is a cross-section of a silicon photonics platform, light is confined in the silicon layer and it's surrounded by silicon dioxide. There's only one single layer of waveguides, one single optical layer. So if you wanted to build a very complex photonic integrated circuit the problem is that you only have one level in which light could be and so what if these waveguides, these lightpaths, have to cross? You might get some losses, so just having one layer becomes very difficult to build large scale photonic circuits. In the electronic domain what happens is that you might have hundreds of millions or billions of transistors in the lower most silicon level but they are actually many layers of electrical interconnects, many metal layers that would connect up these transistors in order to give you all kinds of electronic circuits. So one layer of transistors but there could be tens of layers of metal. But we don't have that in optics.

So we drew some inspiration from this and we developed our own platform in collaboration with a foundry where we had one layer of silicon but on top of it we introduced additional layers of waveguides, optical layers, made in a material called silicon nitride. We've done up to two layers of this optical waveguide material silicon nitride, on top of silicon but you can add more. This type of platform in fact is the most sophisticated type of monolithically integrated silicon photonic platform you will find in the world today. In these types of platforms the light can go up and down between these levels, they can climb up through these types of tapered types of transitions. We can then make what we call overpass and underpass crossings where the topmost waveguide can pass over many bottom layer waveguides or a bottom layer waveguide can pass over many waveguides that reside on the top so you don't have to cross in the plane, you can go up and down and you can shuttle light around and you can add these bridges and tunnels for light on the chip.

This was done in collaboration with a foundry, as I said, and the foundry was A*STAR IME in Singapore and this was on an 8 inch wafer and these are just some scanning electron microscope and transmission electron microscope of the cross-section and of a waveguide tip. We've demonstrated very low loss crossing, these overpass underpass type of crossings. We've also demonstrated very neat new kinds of devices that use multiple layers together. Most people can't do this because they only have one layer so we could be really inventive and come up with new devices that take advantage of these multiple layers together. One example is this kind of bi-level grading, in which we can use this type of device to take light from fibre and scatter that in to the chip. So to connect these chips to a single mode optical fibre. This type of design with uses multiple layers can have much higher efficiency and at the same time be very broadband compared to having just one material. So you can gain something new by having these two materials that you couldn't get if you only had one material. We also developed a new type of p-n junction that is used to modulate light and without going in to the details this is a very specially designed p-n junction and we are able to get efficiency, this modulation efficiency that is 10 times better than if you use a standard, run-of-the-mill p-n junction that is available in a generic platform.

So the main purpose of this work was to demonstrate an optical switch. So an optical switch is a type of device that would have N inputs and N outputs and you can switch from any input to any output, this is a switch fabric. And the biggest chip we did was a 16 input, 16 output optical switch and this is the layout, this is the picture of it. It's a very big chip, it's as big as a loonie. This is our package, we just did some very simple wire bonding and it has more than 300 active devices: switching elements and also photodetectors. Our collaborators, who also partly sponsored this work, were Huawei Research Canada in Ottawa. And they did a bigger version of this chip, you can see the kind of similarities, you can see the patterns because it's just a scaled up version of ours. So they did 32 input 32 output and it's really a tour de force in terms of measuring everything, putting everything together, the packaging and so on. This has more than a thousand elements and these types of chips, this one (Huawei's) especially is the world's biggest silicon integrated photonic chip and we should be very proud that it's all done in Canada.

More Moore

Switching gears a little bit, so we talked about More than Moore scaling so let's talk about More Moore scaling and what I find to be very interesting there is that we can really consider materials that change properties and my group has a strong interest in phase change materials. In particular we've been looking at a material called Vanadium dioxide (VO_2). VO_2 is a material that can go from an insulator state (electrically insulating) to a metallic phase. This is the resistivity plot, like so, and you can get this material to reversibly change between insulator phase and the metallic phase by heating it and cooling it, for example. When you heat it the resistivity drops by a hundred or 1000 times and it can come back up again. It's pretty interesting because it's like a superconductor but its not because superconductors have a phase transitions at cryogenic temperatures, a few Kelvins but this is very much close to room temperature, so a very accessible phase transition. You can also initiate these phase transitions through shining light on this material or by applying a voltage or current. Accompanied with this very large resistivity change is a very large refractive index change and we've taken advantage of this large refractive index change to build extremely small electro-optic devices. So this is VO_2 that is deposited on a silicon waveguide and we can build an optical switch—that turns light on and off—with a ratio of 10db or greater with only about a micron of length. So these are extremely tiny devices.

Through this process of learning about the optical and electro-optical properties of VO_2 we gained a lot of insight into this phase transition and a lot of material properties of VO_2 . I think using this phase change materials could be very interesting for some kind of transistor replacement. We could imagine, for example, VO_2 being in the back end di-electric here, integrated with the back end metal. You could make maybe silicon memristor and neuristor types of devices this way. I don't think it's a good idea to put VO_2 as the channel material yet for transistors because the mobility is not so high but we can probably put it in as part of the gate or in the the di-electric layer for the transistor so that when this thin layer of material goes from being an insulator to a metal you can get a very steep turn on of the transistor.

New Microsystems: Light and Brain

Finally, in the final ten minutes or so of the talk I would like to talk a bit about light and the brain. We think about new types of computing systems, new types of paradigms—computing paradigms—I think we can draw a lot of inspiration from the human brain. The brain—we all have one, it's the world's most efficient computer. Its very energy efficient, very robust, it can do a lot of things that a regular computer cannot do; it can learn, it can adapt, it can heal, the power consumption is very low (10 miliwatts). There are about 100 billion neurons in the brain and about 1000 times more synapses—connections—between them. Our brain is constructed very differently from a microprocessor. As many of you know, machine learning, deep learning, which is all the rage these days, draws some inspiration really from how to brain works, crudely. So there is a lot of excitement around these machine learning methods but they're really based around how the brain works. So if we can understand the brain better perhaps we can think of even newer algorithms that could fuel machine learning. At the same time, so little is understood about the brain that perhaps with better neural networks, or better artificial intelligence types of systems we can interpret the data that is being measured from the brain in order to make sense of what is happening. I really see an interesting intersection between biological intelligence and artificial intelligence and that could be coming up. And the field of neuromorphic computing describes the opportunity of having brain-inspired computing and that could be at the device level and it could also be at the algorithmic level.

Along these lines, what can optics do? One potential way in which optics can contribute is to do some kind of processing that is done on existing algorithms in the optical domain, so known algorithms, say some deep learning algorithm, convolution, neural networks, what have you, there are certain steps that have to be done, matrix multiplication and so on. Instead of executing it in the microprocessing, electrical domain can we do these types of steps in the optical domain instead? The advantage of doing optical processing instead of running it in a run-of-the-mill, standard computer is that this can be done at higher speed, you can have parallel processing, and there are also some phenomena that happen naturally in the optical domain— non-linearity and so on— that might make some of the steps, some of the computations that are done much more natural in the optical domain compared to the electronic domain. Because of this there have been very recent reports, all within the year, the first is out of MIT in which they implemented matrix multiplication that was used for artificial neural networks in the optical domain, in a silicon photonic chip. Actually, this chip is smaller than the ones we have demonstrated at Huawei and in my group but it's an example, a proof of concept demonstration.

Another maybe direction that has come up to implement artificial neurons in the photonic domain. In a lot of these types of morphic computing systems there's a need for having spikes and in an electrical domain you actually need to design circuits to create spikes but in the optical domain you can create spike trains just by injecting light in to lasers, so this just happens naturally due to the laser dynamics. So there's interest in creating photonic artificial neurons and maybe these types of implementations can be faster and more power efficient than actually doing it in the electrical domain. If we just want to execute well-known algorithms in the optical domain then there are integrated photonic circuits and devices that can be used for implementing some kind of switching, spiking, maybe thresholding function.

I am much more interested in some fundamental questions about how the brain works. And maybe by understanding how the brain works we can inspire even newer computing strategies. What has been very exciting in the neuroscience community over the past 10 years or so is that there have been new ways of trying to investigate neural activity with light. Optical studies of the brain is much more low noise it could be much less noise and specific than electrical studies of the brain because there is no background optical activity in the brain. This has really spurred a lot of research into using light to map neural activity. One example of this is to do functional imaging with optical reporter molecules. These neurons have been genetically engineered so that they can express a protein so that when the neurons fire they would emit light. You can create really beautiful types of brain-bow, they're called brain-bow images like this. The problem there is the brain is full of stuff so light cannot really penetrate very deeply in to the brain because the light would just scatter away. The brain tissue is like milk so it seems opaque to light. You need very large microscopy set-ups to do this and you couldn't penetrate very deep, maybe a millimetre or so. In another domain, the filed called optogenetics has emerged. This is a field in which, again, these neurons have been genetically modified such that you can turn on and turn off the neurons by shining light on them. This is a prototypical example experiment in which there will be a fibre optic implant stuck in to the brain of a mouse and we turn the light on and some behaviour can be observed, turn the light off and something else would happen and this is optogenetics.

We wanted to create new devices to further the paradigm of using light to investigate the brain and because we are integrated photonic engineers we know we can do much better in terms of integration, in terms of the complexity of devices, than what the neuroscientists have come up with. So in collaboration with Caltech, the group of Michael Roukes, we’ve been developing integrated neurophotonic probes. This is a schematic of some of these probes, its been done in silicon there are these shanks that get implanted in to the brain, light can come out and be collected along these shanks so we can use these chips in order to monitor and stimulate neural activity. This work is also done in collaboration with Roman Genov, who is a bioelectronics professor in the department, Andres Lozano and Taufik Valiante who are in the department of neurosurgery in Toronto. This type of architecture with these probes will allow us to overcome the scattering, the limit of light in brain tissue. We are trying to make all these chips in a foundry process, in a silicon photonic foundry so we can mass manufacture these dies and then distribute it to the neuroscience community. And because we can have a high degree of integration we can have many of these light-emitting pixels, and light collecting pixels, to enable massive parallel interrogation of neurons.

We have made some progress in this direction, all of these chips will have to work within the visible wavelength range because this is where these light sensitive proteins will operate. Luckily, we started working in silicon nitride waveguides and they are transparent in the visible range. We’re developing this in partnership with the foundry in order to make these multiple waveguide layer types of structures so that we can have that high degree of integration that is necessary to create these brain probes. Here are some example devices: this one emits red light, we could have very rapidly switching pixels, here in the blue. We also have developed some ways of using phased array to steer a beam without any moving parts. You can see the beam being moved on the surface here and there are no mirrors or moving parts to do this.

I am part of this program, this initiative which is call the Neurotech Alliance which Michael Roukes has started out of Caltech and you can go to this website to learn more. Our goal is to use foundry capability to manufacture these neural probes, they could be optical, electrical, chemical or combinations thereof. This is also done in collaboration with Ken Sheppard who is a professor at Columbia and I’m the optics person on the team. And we have foundry partners CEA LETI in France and A*STAR IME in Singapore. The idea is to develop these chips and to develop systems that are completely open source so anyone who is interested can see them, use them and can know how to build them and we’ll give them out for hopefully free, or maybe ... we’re definitely not-for-profit. We also have a set of Alpha doctors who are neuroscientists or neuromedicine researchers that will test these chips and give feedback to us.

I think it will be an exciting time when we can use nanotechnology to make some sort of neuro-interface because the integrations that is afforded nanotechnology allows us to put multiple forms of sensors and actuators on a very small form factor. For example, I only talked about light today but you can certainly integrate electronics on the same chips, you can imagine having fluidics, micro-electromechanical and nano-electromechanical systems all together. The types of experiments that this could enable from a neuroscience perspective will be phenomenal because now you can actually input multitudes of combinations of signals whether it's the voltage, temperature, the light, or what have you. You can monitor many different kinds of physical signals at the same time, you can also monitor things that are much more macroscopic, like behaviour, and so these types of tools will be very interesting and could really open up a different kind of paradigm for neuroscience.

Summary: Beyond Moore's Law

Summarizing, I think computers in whatever form they will take; whether it's a datacentre, something wearable or eventually in our brain, in our head, are very important for solving a lot of important questions in society phases: in sustainability, health, transportation, and so on. Every time there is some kind of new computing paradigm it's always been enabled by very specialized hardware. Whether we're talking about AI, which was enabled by GPUs (Graphic Processor Units), mobile age, cloud computing, every time this has happened it's been enabled fundamentally by some new hardware. Maybe even quantum computing one day will need specialized hardware to make that computing paradigm a reality.

Photonics, as I hope I've shown you, is a very critical part to this future of computing. I've talked about More than Moore, enabled by foundry 3D silicon photonics, I've talked about More Moore scaling in which we're looking in to phase transition materials to create new types of devices, and then I've talked about brand new things that are well beyond Moore in which we are looking at new interfaces to the brain. I think in the future, what is the goal, it could be 30-40 years from now, I think that computers will be cognitive, meaning that they can think. And I don't just mean at the algorithmic level, this could be at the device level and the physical level that the kinds of computers that we'll build would change and think and be smart and will be intelligent.

With that I'd like to give thanks to my students, to my group, to recent alumni, and especially to Wesley Sacher, he did a lot of the work related to silicon integrated photonics. Thank you very much for your time. (Applause)

Audience Member #1: The work that's going on, is it looking at any way of delivering chemicals in to the brain. For example, I have Parkinson's and I'm very interested in that.

Professor Joyce Poon: Yes. So because we work directly with these silicon foundry partners so we have all these plans for these types of platforms that we will create and indeed one of the directions that we will take includes integration of fluidic channels so these will be microchannels that can carry, for example, drugs. It can also sample chemicals in the brain as well. So it's definitely part of our scope. We realize that drug delivery can be an important application.

Audience Member #2: The modelling that you have described to us, that you outlined the partners with whom you work, how much of that is micro-modelling, what kind of facilities do you have to do micro-modelling at the scale you're talking about on site?

Professor Joyce Poon: My group has several work stations that are fairly powerful, that have a few hundred gigs of memory. When we are very crunched for computation time and resources we also buy time. We've bought time on the Amazon cloud computing service as well. And the type of modelling we do is a combination of electromagnetic simulation, so we'd actually simulation Maxwell's equation in space and time, time stepping and space stepping, so it's a direct solver for Maxwell's equation. We also solve for the electronic properties that carry density and so on, so that's done through a TCAT simulation. In that domain we also do the simulation of the fabrication which includes the implantation energy and so on in order to form these p-n junctions. Finally we also do multiphysics modelling because we also have thermal heaters. We also do mechanical simulation because, you can see the weird probes, we have to understand the mechanical stresses, whether they will break, whether they'll be straight or not after the bottom has been removed. So we also do mechanical modelling and when we do these types of simulations, you're right it's very tricky, because we use a lot of commercial software to do this but we have to interface them. We may simulate the electronic properties first, then tick that result and transport it over to an optical resolver to understand how that would affect the optical field. Same with heat, you do a heat transport solver and then we would mesh it with an optical solver.

Audience Member #3: I'm just wondering what degrees of collaboration you have with EM group in regards to EM simulation or metamaterials.

Professor Joyce Poon: We haven't really collaborated with the EM group on this. Our department is very strong in electromagnetic simulation. I did actually collaborate with Professor George Eleftheriades, not on the simulation part, not in terms of developing algorithms for simulation, but he was interested in having a metasurface that was highly tuneable in the optical domain so together we worked on a paper that describes a tuneable metasurface implemented in vanadium dioxide. Because of its very large change refractive index you're able to have a very wide degree of tuning. So we're collaborating but maybe not in a way that you were thinking. We're not developing new algorithms for simulations.

Audience #4: My question has to do with the foundries. You said that you collaborated with A*STAR foundry. So from what I understand, starting up a foundry for a new kind of process is very, very expensive. So I'm just wondering about the economics of that, how that really works.

Professor Joyce Poon: Great question! We also collaborate with LETI, IBM. You're right, the cost that they have to spend to run the services is very high, and bringing in equipment is very high. We are very fortunate because we would come up with new ideas and the foundry thinks, oh this is really new, and it could be useful for us down the line. So we went in to these collaborative models in which we work with them, and this is not something that they are offering out to the public yet. This is the phase of the foundry doing its internal R n D before it would e offered to the public. So you're right, it costs them a lot. We also pay a lot to work with them, but they probably put in more than what we pay them. They will generate some IP, because they have to develop new processes, for example if we want to have these multiple levels then they had never done something so complicated before so they are also generating their own internal IP. So they see it's worthwhile so they really trust us to work together. We've been very successful in partnering with them and finding good results. Sometimes now they even approach our group to develop something new.

Audience #5: I just want to get some clarification. Sometimes I read about quantum computing. In this talk that you've just described, where does quantum computing fall in to this? Which of the axes and stuff like that?

Professor Joyce Poon: I think you should ask Professor Voinigescu who has a deep interest in quantum computing. I would think that quantum computing is actually a little bit along the diagonal because you would still need extremely small transistors to hope to see any quantum effect. So you need to have something very small going down the More Moore scale but its also using the existing infrastructure to do something new, new applications so that's kind of horizontal. And at the same time you could create new computer architecture, computing paradigm, so that's a new type of microsystem. I think if it was to be done, in silicon, it would likely have to be at cryogenic temperatures. The computer will not look like something on a desktop, it will have a lot of other stuff attached to it … [Professor Sorin Voinigescu interrupts] you don't think so? You think it could be room temperature? I don't think so.

Professor Sorin Voinigescu: It's going to be room temperature or it's not going to happen.

Professor Joyce Poon: It can happen if it was to be room temperature.

Professor Sorin Voinigescu: Well, it happens today but that's not the real application.

Professor Joyce Poon: So anyway I think it belongs to the diagonal.

Audience #6: I heard about D Wave. Is there any collaboration with that group?

Professor Joyce Poon: I personally do not work on quantum computers and we do not collaborate with D Wave so I can't comment on that.

Professor Farid Najm: Maybe you can chat about that after. So let's thank Joyce again for her wonderful talk. You'll find on your chair there's the ANNUM publication which you are welcome to take with you and there's also a small postcard for CONNECT, which is our platform for remaining in touch with alumni and for alumni to remain in touch with each other and mentor students. It's a very nice way to stay in touch and I encourage you to join. If you haven't, the website is there on the card.

Please enjoy the networking and the food and drink. Thank you again for coming. [Applause]