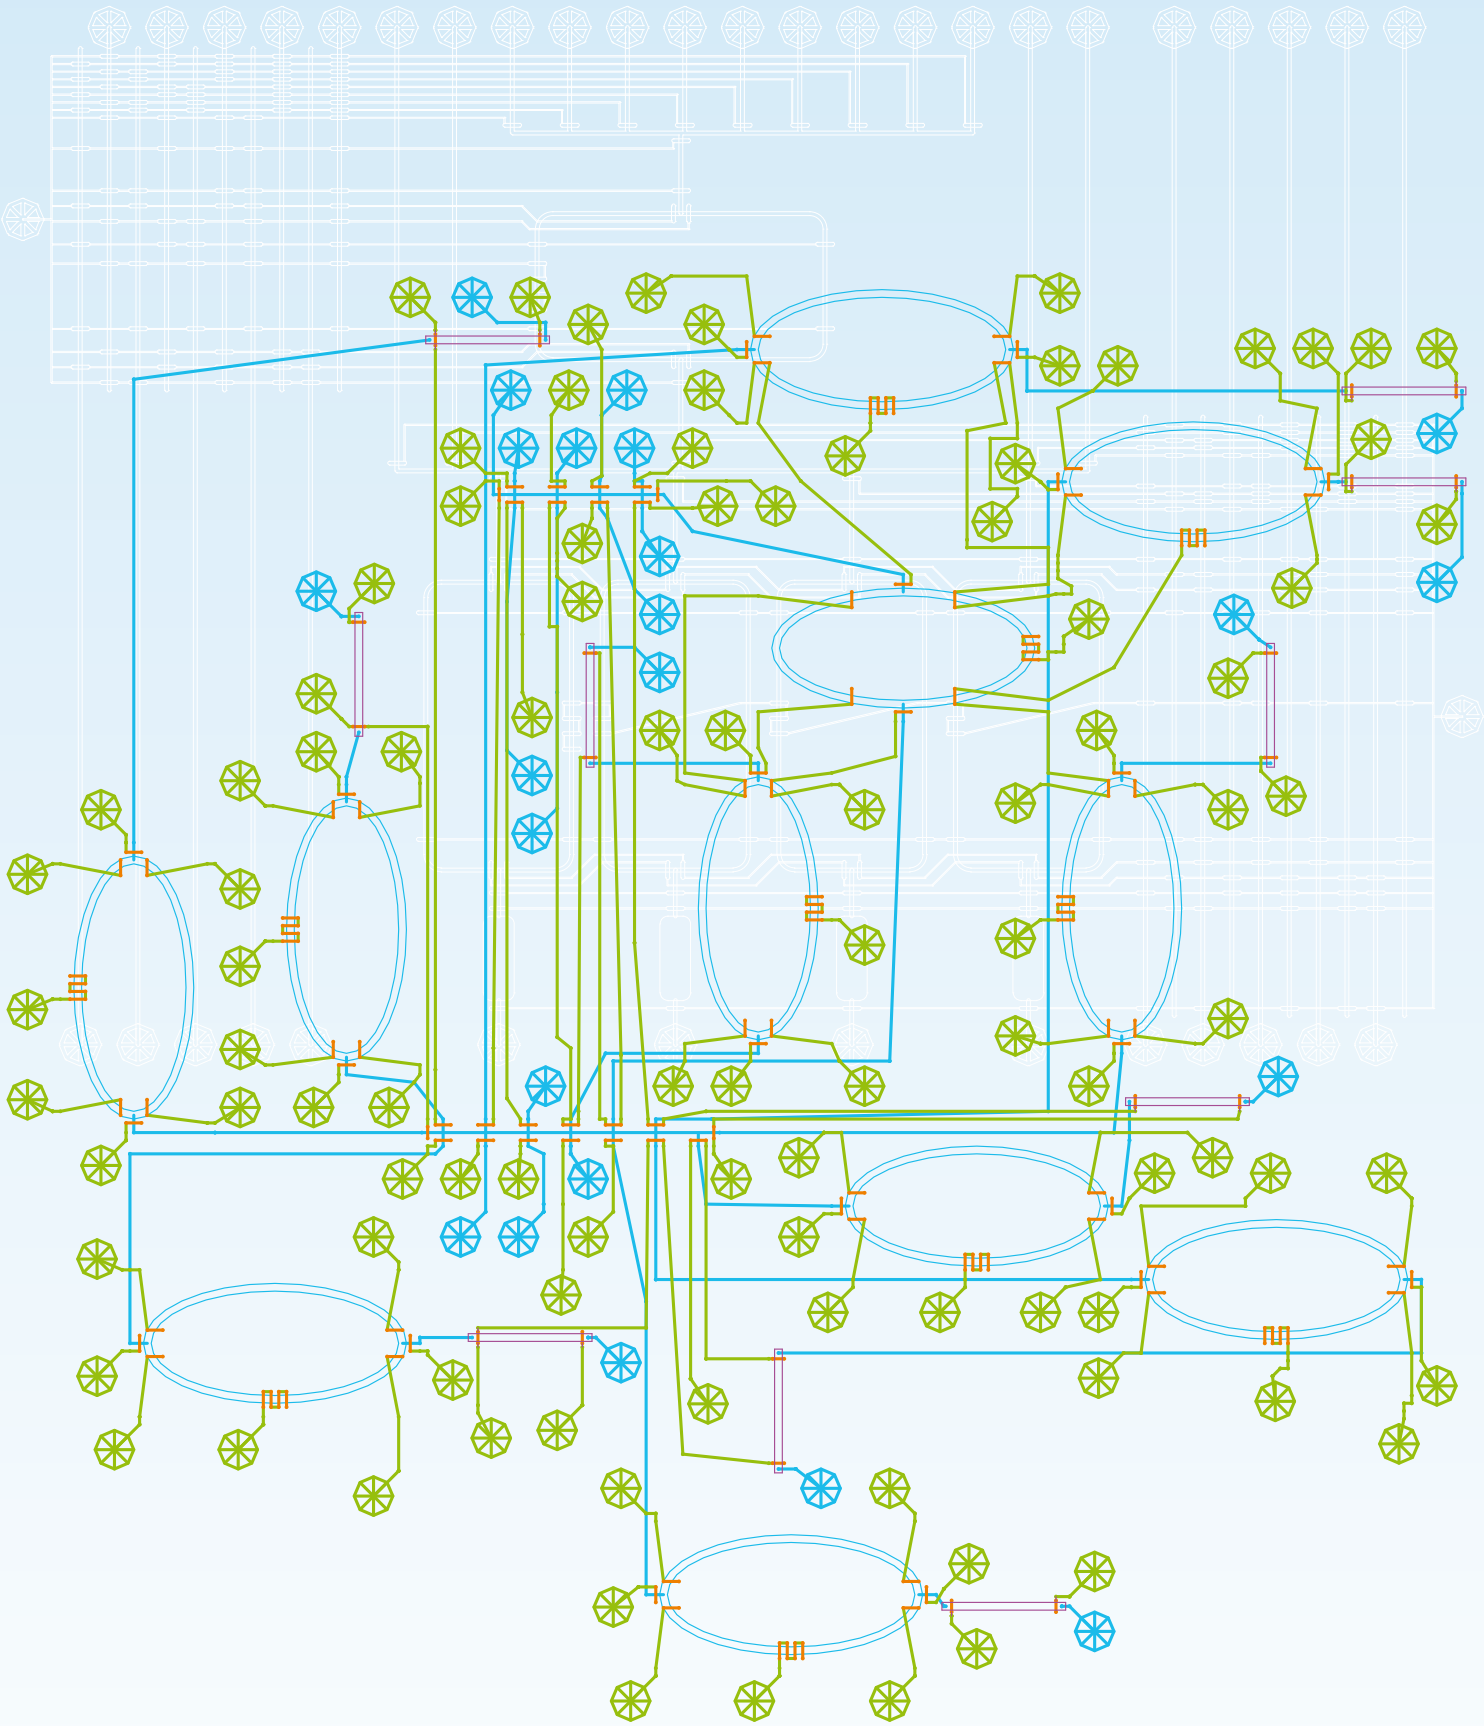

Microfluidic biochip layout from Columba automated design tool.

Microfluidic Design Automation and Neuromorphic Computing




An emphasis on building new bridges between electronics and biology unites two TUM-IAS Focus Groups: Microfluidic Design Automation (MDA) and Neuromorphic Computing.

Both are hosted by Prof. Ulf Schlichtmann of the Department of Electrical and Computer Engineering at TUM. The principal collaborators in the MDA Focus Group are Prof. Krishnendu (Krish) Chakrabarty from the Department of Electrical and Computer Engineering at Duke University in the USA, a Hans Fischer Senior Fellow of the TUM-IAS, and Prof. Tsung-Yi Ho from the Department of Computer Science at National Tsing-Hua University in Hsinchu, Taiwan, who is a Hans Fischer Fellow. This Focus Group is also supporting two doctoral candidates, Yasamin Moradi and Chunfeng Liu. The Focus Group Neuromorphic Computing was established more recently in collaboration with Hans Fischer Fellow Prof. Hai (Helen) Li, who also is a member of Duke University's Department of Electrical and Computer Engineering.

The MDA Focus Group is addressing critical issues that could limit both the versatility and impact of microfluidics technology, best known as the basis for the rapidly growing "lab on a chip" market segment. Microfluidic biochips, microarrays, biosensors, and microreactors can be used in a range of applications including DNA sequencing, clinical diagnosis, drug discovery, and environmental monitoring. One of the keys to unlocking the technology's full potential will be computer-aided design tools similar to those that have helped the microelectronics industry master huge increases in the density and complexity of integrated circuits while also keeping reliability high and costs low. That is the main focus of the MDA collaboration.

In a way, it is the very success of modern computing that allows creative minds to imagine applications that expose its limitations – and that is where the emerging field of neuromorphic computing comes in. Interconnected trends in big data, embedded systems, and communications are giving rise to expectations that traditional computing architectures and devices seem unlikely to meet. Research in the Focus Group Neuromorphic Computing is aimed at providing an alternative inspired by the human brain and nervous system, through the investigation of novel devices and architectures that remove barriers between computer memory and processing.

The TUM-IAS conducted an interview on these two lines of research – as well as some areas where they intersect – with Tsung-Yi Ho (TH), Helen Li (HL), and Ulf Schlichtmann (US) in Munich and Krish Chakrabarty (KC) in North Carolina, who took part via videoconference.

Tsung-Yi Ho

Q: The promise of the “lab on a chip” has been heralded for many years already. What’s the status of real-world applications?

TH: There are several applications now, ranging from small-scale to large-scale. One of the key applications would be next-generation DNA sequencing. Companies are already using microfluidics for preparation of samples. There is even progress on “organ on a chip” technology. It is possible to simulate an organ on a microfluidics chip, and this has been used already in a clinical test. It shows a lot of promise.

KC: I would say, though, that we still are in the age of unrealized potential. There’s been a lot of hype and excitement, a lot of activity. But most of the work that these chips do is still the tedious work that people used to do on the benchtop, things such as pipetting into a test tube, shaking the test tube, doing mixing, transferring the sample to other media. Those things have been automated and have been demonstrated. There are many research labs that are using lab-on-chip to do these things, and there are a few scenarios in commercial practice, though not a lot, where this is done. But these chips still do not have intelligence in them. You still have a lot of human intervention, in making the decisions on how to interpret the data and what to do next. And I think this is where the next-generation lab on a chip can be very powerful, because we can embed real-time intelligence and decision-making capabilities into the systems.

Q: Is that something design automation can help achieve?

KC: That is one activity of our MDA Focus Group. We are looking at going beyond the regular tedious things that people do right now on the lab on a chip.

Q: And a prerequisite for that is reducing the human effort involved in designing the chips?

TH: We know a PhD candidate at EPFL who spent almost one month to design a microfluidic chip by hand, with 918 valves, the fundamental element. So think about it, if you have more than 10,000 valves – you can put that many elements on a chip, but it’s impossible to design it manually. You must go to automation.

US: Design automation for integrated circuits got started when they became more complex, so that the human designer had difficulties in handling the complexity in the first place and in getting anywhere near optimal solutions. Today in microelectronics we are dealing with systems that have billions of components. In microfluidics, we’re not there yet, but they are getting more complex, in terms of the number of operational units that you have on such a chip, how they are interconnected, and the assays that you want to perform. We’re getting to the point where, once you have a certain assay and want to design a specific biochip for that, you get there faster using automation techniques, and I think in many instances you will get a better solution than the human would get.

TH: We also want to do additional optimizations – to minimize the area or minimize the volume of fluid used – and these things all require design automation. Also, most microfluidic chips currently handle very basic operations, steps that traditionally need to be done separately. To do sample preparation, cell lysing,

Tsung-Yi Ho, Helen Li and Ulf Schlichtmann

extraction, and amplification, for example, and transfer the result of each step to the next, you might need at least three separate chips, or three components, or three robotic devices, and there may be some loss of material along the way. With microfluidics it is possible to integrate them so all the procedures can be performed on only one chip. But designing more highly integrated biochips by hand is impractical.

Q: An obvious difference between electronic integrated circuits and microfluidic chips is that you're dealing with different physics. What kinds of knowledge, technology, and experience can be successfully transferred from microelectronics to microfluidics?

KC: The similarities are more in terms of the fabrication. The lithography techniques are common. We make chips in pretty much the same way, using masks and lithography as in electronic circuits. Even bottom-up self-assembly has been attempted for both types of technologies.

US: Another similarity is that you have operations that need to be executed in a certain sequence, and the result of one operation sometimes feeds into the next operation. You need to decide when you do these operations, and on which of your chip's functional units you do them.

TH: The properties of fluid samples are totally different, of course. So there is a different delay, and fluid dynamics, when you design a lab on a chip. That's the big difference. And currently these chips are still designed bottom-up. That means that every component is designed by hand, and they can put it together, integrate the components, and then realize it doesn't work. They will iterate many times. For integrated circuits, the design automation framework is already very mature, and it's more top-down. So that means given the specification and the goal, we can design for the result using customization globally.

KC: When you look at electronic circuits, you have to first ask the question of what is the science, how do you explain electronic behavior. And it comes down to electrons and holes, the physics of the devices, and it has been very well explained and understood over many years. In microfluidics, that is not the case. There

Tsung-Yi Ho, Helen Li and Krishnendu Chakrabarty

are still a lot of open questions, and a lot of gaps in explaining how these devices work. The physics is much more complicated. It involves fluidics, sometimes optics, mechanical parts. And also in semiconductor physics, the underlying science is the same for most of the technologies. In microfluidics, you have different physics for different technologies. So we have continuous flow in channels using pumps and osmotic magnetics; that's one set of physics. Then you also have dielectrophoresis, where you have very high voltages and cages, which are being moved around under electrical control. There's electrowetting, and many other types of esoteric mechanisms. So that makes it very difficult.

I would also point to a problem that requires automation in microfluidics, something I call volume management, which we don't see in electronics. If you look at semiconductor integrated circuits, the carriers are effectively infinite. You can take as many electrons and holes from the substrate as you want, you can get your currents, and you can send them back to the substrate. It's a closed loop. In microfluidics, we're dealing with a finite amount of liquid volume. The stock volume is limited. We talk about having a reservoir with a certain milliliter volume. So any reaction that you do, any chemistry, you must be aware of the volume limitations. You can't set it up so you run out of one reagent while you have other media left over that you can't do anything with. So a lot of the design automation involves looking at how you can optimally utilize these volumes and finish the required chemistry in the shortest possible time.

Q: Where do you see microfluidics technology heading, and what are the other critical hurdles?

US: As Krish mentioned earlier, we are very interested in this fusion between microfluidics and microelectronic intelligence, cyber-physical systems as they are called. You have the microfluidic chip on the one hand, and you have electronics on the other hand monitoring the outcome of steps that are happening on the chip. What we're dealing with on the chip, after all, is chemistry and biology, and the outcomes of certain reactions depend on variables such as temperature and cannot be completely predictable. So you need to observe them using, for example, some optical system, and then decide whether you repeat an experiment or what is the next step. And if you do

that using electronic intelligence, you might be able to handle fairly complex biochemical assays and analyses on a single chip, without human intervention. That potentially might bring us closer to the promise of doing complex analyses at the point of care, even if there is no doctor or trained technician present, for example in a rural village in Africa.

KC: So clearly the challenges also come down to the cost of the products. These are mostly throw-away devices. You don't use them over and over again. And anything that you throw away must be cheap. If you do not have design automation, you really cannot explore the entire space of implementation, and you won't get to the lowest cost point very easily. If you have the tools, as Tsung-Yi was saying, you can use your mouse on your computer, click on different options, and get different layouts. And you can do it in a loop, considering larger numbers of possibilities, or use some mathematical frameworks for doing optimization – down to the smallest form factor, the smallest package, the lowest cost. It's very important.

We worked with a start-up company that came out of Duke, founded by one of my postdocs. They were selling to hospitals in Illinois. Their initial chip was manually designed. It was done by hand. They had come up with a layout with a small number of input/output pins, but the pins were a problem, because they took up a lot of the package area. Using mathematical tools and design automation techniques, we were able to decrease the number of pins by a very big number. The second version of the chip used our optimized layout, and they were very pleased.

Q: Has the framework of the TUM-IAS Focus Group helped you bring together complementary areas of expertise, extend your own capabilities, or expand your reach?

TH: Definitely. When we formed this Focus Group and started doing research on microfluidics, of course we started by focusing on reliability, because we have relevant expertise here. Ulf is the expert in designing integrated circuits for reliability. And very quickly we were able to publish findings on automated design for microfluidics, particularly considering reliability,

Ulf Schlichtmann

at a very good conference and also in a very good journal in our area. We have research results, and we also have a plan for a business. And in July 2017, thanks to the TUM-IAS funding, we had the first Munich Workshop on Large-Scale Microfluidics Design Automation. We've reached out to people designing biochips and also applications, in industrial as well as university research, including the European Microbiology Laboratory in Heidelberg. So we have already expanded this collaboration.

US: I'm in some sense the new kid on the block here, because Krish and Tsung-Yi started on microfluidics much earlier. Krish is the father of design automation for microfluidic biochips. And I was introduced to this topic first by Tsung-Yi when he was here as a Humboldt Fellow. When we realized that some of the techniques we have been using for electronic chips are also applicable for microfluidic biochips, that got me and my team really interested in this topic, specifically in terms of reliability. Some biochips are throw-away, as Krish was saying, but some of them might be used repeatedly, and the valves used in these chips are mechanical in nature, so they have a tendency to wear out. Traditional biochip designs were using the

valves very unevenly, and we had some ideas about how we could even out the usage of valves to increase the lifetime of the chip. This is one area where we brought some specific research that originated in electronics into the microfluidics field.

KC: Ulf is being extremely kind. I don't want to be the father of too many things. I have my own kids at home, I'm their father, and that's it. But otherwise Ulf has put it very well. We share the same excitement.

In the past I have looked at a whole range of design automation problems, from synthesis to placement, layout, routing, and chip design, also looking at failure mechanisms, how these chips fail, what kinds of faults they exhibit, how we can test them properly at low cost, make them reliable. But of late I have started moving more toward the biology. I want to understand the applications better, and what are the various biochemistries that we can run on these chips. So the motto that I follow now is "realistic microbiology on simple chips." I would love to see these chips simple enough that the user isn't intimidated, but the applications that would run on these chips would be complex, would be realistic. And that's where design automation is key, because it can allow complex chemistries to run on chips that are, by themselves, pretty dumb. But we have software, we have intelligence layered on top of the actual circuitry, that allows realistic chemistry to run.

So I think that is the complementarity. Tsung-Yi is one of the most respected persons I know in the design automation space. And I'm now looking more at how can I bring that to the biologists, what is the bridge we can form.

Q: How would you describe the bridge that connects these two TUM-IAS Focus Groups, Microfluidic Design Automation and Neuromorphic Computing?

HL: It is a two-way street. My research, especially where it concerns novel electronic devices, can benefit directly from both their design tools and their expertise. In the long run, I hope what I bring in terms of machine learning and a new computing paradigm might also contribute to advances in microfluidic design automation.

Essentially, neuromorphic computing aims to utilize what we know as VLSI, very-large-scale-integration circuits and systems, to mimic biological nerve systems and then to achieve more functionalities – including cognitive functionalities – and even further, to enable the system to learn by itself and maybe achieve some self-awareness in the future.

Neuromorphic computing is different from traditional computing. In the architectural perspective, traditional computing separates the processing from the memory. Processing can go really fast, but memory becomes the bottleneck. Essentially what we have in our human brains is a large volume of nerves and synapses that form network structures. We want to reform the computing architecture by mimicking that structure and then realizing it with VLSI circuits and systems. For about 30 years, efforts to mimic biological systems have focused mainly on the software level. Neural networks, machine learning, so-called deep learning – the main emphasis in these areas is on making algorithms functional and executable on existing computing systems. What we are trying to do is to put it completely into the hardware levels, so that execution and implementation will be more efficient.

Q: What role does the so-called memristor play in this?

HL: It lets you do away with the wall between processing and memory at the most basic hardware level. The memristor device was originally predicted by Leon Chua back in 1971. Stan Williams, at HP Labs, led the group that was first to realize a physical device representing the expected characteristics. HP's device is an oxide-based device. It's a thin-film structure. When you apply voltage or current through the device, charge across the device – which can serve as a basic unit for computational logic – is represented by its resistance states. And the good thing is when you remove the charge, the resistance states will remain. This behavior is very similar to the synapse in biological systems.

Another very important similarity between the two is the potential for large connectivities. So even though, as I said, 30 years ago people were talking about neural networks, the huge volume of connections is an essential requirement if you'd like to get useful applications from it. What we can do with the memristor, since it's only

two-terminal thin films, is form what we call crossbar structures. Essentially you have horizontal wires and vertical wires, and at each crosspoint we are able to make one device. So you can see the connection density is extremely high.

Q: Potentially billions on a chip?

HL: Potentially billions on a chip is not a problem. We might have to partition that into small arrays or groups in order to get flexibility in the design, in the functionality.

Q: Is this still purely a research topic, or are there products based on memristors?

HL: I do not think there are memristor-based neuromorphic computing system products yet. The technology itself, which is often called ReRAM or resistive memory technology, has been under development for several years. In general, any memory technologies using resistance states to represent logic for information can be categorized as resistive memory. One of Leon Chua's recent efforts in fact is to prove that many materials, including spintronics and phase change, are part of the memristor technology umbrella. Memory products based on these technologies are available now. And we have devoted a lot of effort to developing a new computing platform and realizing reasonable applications using memristor technology.

Q: Where do you see the biggest advantages this approach might offer?

HL: By bringing information storage and computation together, simultaneously, a neuromorphic computing system is expected to be especially good at handling perceptions, cognition, and learning, and at capturing and maybe extrapolating from existing conditions. For some applications, such as scientific computation for instance, existing computers might still be more advanced than a human being or a system that mimics a biological system.

In the Neuromorphic Computing Focus Group we are primarily working on the architecture but also circuit design, and working to bring those design concepts to the point where they can be implemented and realized.

Helen Li

We put a lot of effort on the design side, and we also put a lot of effort into trying to connect with application levels. Initially we were working on the design concepts and a simple demonstration. But right now we're trying to move to larger-scale systems and make the approach really useful to potential users or customers.

Q: How would you define large-scale in this context?

HL: The demonstration originally is only dealing with very small things, like recognizing characters, for instance, to tell whether this is A or B or C or D or something else. But when we progress toward really useful applications, this actually requires a lot of computing resources. We have to make the hardware bigger and bigger to accommodate our requirements. This is what I meant by large-scale systems. But when we eventually shift toward commercialization, we will face the quality and reliability problems.

Tsung-Yi and I already had several interactions in these areas several years back, dealing with a lot of placement and routing issues. In that case we borrowed Tsung-Yi's intelligence to solve the problem. And it went really really well. We actually reduced the design size more than 50 percent, and the routing sizes more than 40 percent. That work was published in 2015 and was nominated for a best paper award. In addition to the design automation, we would further need reliability controls. All emerging technologies such as nanotechnology-enabled neuromorphic

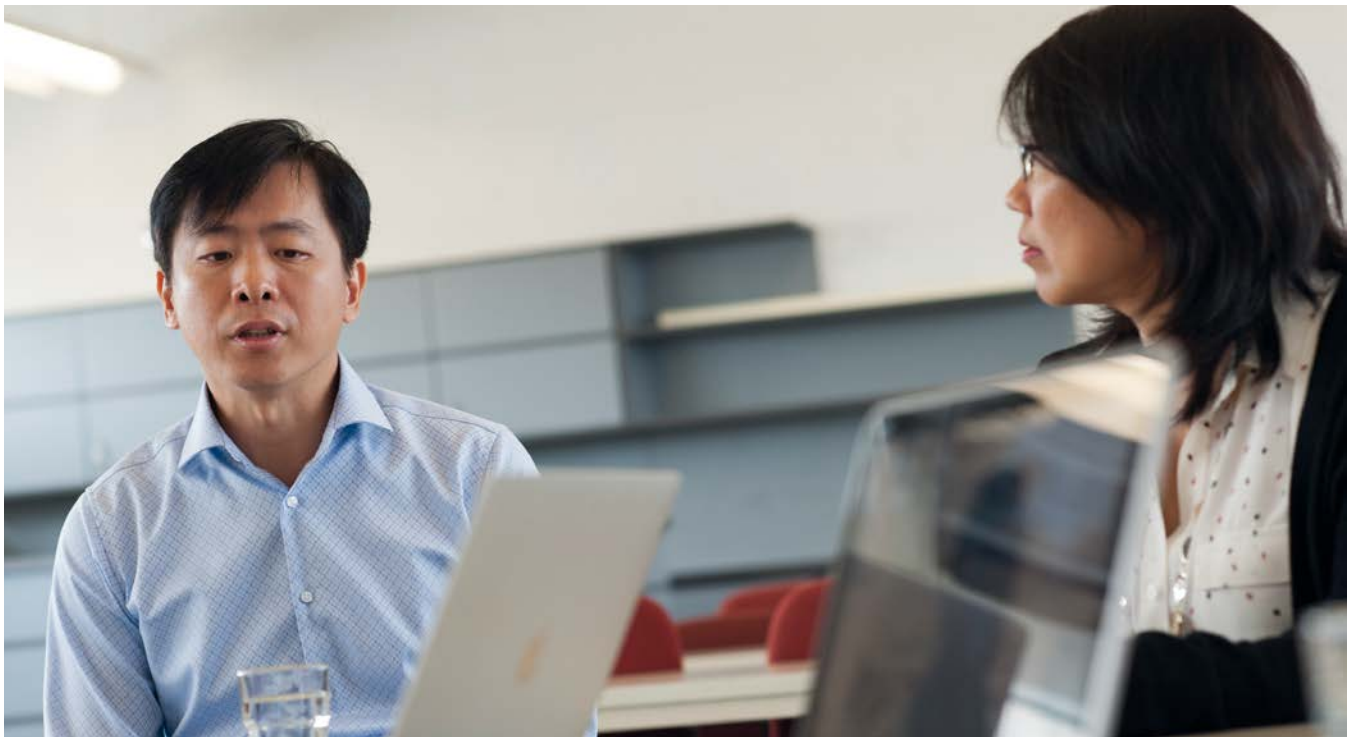

Tsung-Yi Ho and Helen Li

computing will have significant reliability issues, for example defects, huge variations, temperature and ageing effects, and so on. We'd really like to work with Ulf on these reliability controls. And further, once we are getting ready to deliver a product, we will need to get it validated and tested, one of the special strengths Krish brings to the group. Thanks to our collaboration, I feel free to focus entirely on the neuromorphic computing.

KC: Helen's work is very exciting. She's a new faculty member at Duke, and we're very excited that she has joined us. Now Helen is going to tear down the memory wall that has been such a long-standing problem in computer architecture, where the processor is starved and waiting for data. I can only cheer her on: Professor Li, tear down that wall!

Q: There, in a nutshell, is one breakthrough you'd like to see. What are some other major outcomes you have in mind?

HL: I am hopeful about the computer's ability to learn by itself. We are already in the big data era, and if the computing systems could be able to learn from those huge amounts of data and extract information targeted for different applications or customers, I think there will be a very dramatic improvement in this area.

US: I personally am not looking for one big breakthrough because, as they say, it's hard to make predictions, especially about the future. What I'm hoping is that our research will be used by other people to advance their fields, whether that means understanding biology, curing diseases, or something else.

TH: Currently, most of the people using microfluidic chips use them for just one dedicated purpose. They may not think it is even possible to create multipurpose integrated microfluidic chips. But with our tool, it is possible. I hope with our technology they can have a general chip that can integrate everything together.

KC: I'm hoping that in a few years time, we can make fundamental breakthroughs in cancer research, with customized drug regimes for individual patients. We can do that through using microfluidic biochips as a key enabler for large-scale distributed experimentation, collaborative experimental research. I can envision a large group of labs and scientists running experiments in a very collaborative fashion, helping each other through microfluidic-enabled biochips and

the cloud infrastructure such that, let's say, two plus two would be five, not four. So that they can realize treatments, pathways for disease, in a distributed, collaborative manner.

Q: All of you have traveled widely and have spent significant time in various visiting positions at leading international universities. How does your TUM-IAS experience compare to other visiting stays?

TH: Krish is the real world traveler, but I have spent the most time here, starting from 2011. I first came here and met Ulf, and we started collaborating with support from the Humboldt Foundation. I think Humboldt provided a very good start for us to initiate this project, and it went well, so we applied for support from TUM-IAS to continue it. The main difference is that the Hans Fischer Fellowship here provides support for one doctoral candidate. I don't need to worry about any internal financial issues during my stay at TUM, and this is very good. Our doctoral candidate is doing very well, and he is learning about working in different cultures, Taiwan and Germany. That's something very special about this program.

KC: I would add that my experience has been that in many places where I visit, I am more like an outlier, in that there isn't much of an internationalization effort at the host university. I am brought in on the initiative of a professor there, my collaborator, and it's mostly inward-looking. At TUM there is a structure and an organized effort to reach out, and I am no longer an outlier. There are many people like me on campus and in the TUM-IAS who are from other countries and who are working closely with other researchers. So I find that internationalization is in the genetic code of TUM and the TUM-IAS, and that comes through all the time.

US: In science, often you meet somebody, you think you have some great ideas, then you get back to your office and there's tons of the usual stuff to do, and the cooperation never really happens. But here there really is quite a bit of common work, not just with jointly supervised doctoral candidates, but also extending beyond that. There is collaboration with a top PhD candidate of Krish's, who by the way took the opportunity of another TUM program, the Research Opportunities Week, to come here last spring and to interact with people here. Tsung-Yi has been very generous in giving advice to quite a few other doctoral candidates, and this has resulted in a very significant number of papers. And we have really been successful in targeting the top venues with our publications.

KC: There's also the opportunity of bilateral exchange visits. In the other places, it's one-way. So I visit, but there isn't much opportunity to invite researchers back to Duke. In the TUM-IAS it's much easier, because there are programs and funding to facilitate that. I think it's also very good, and possibly unique, that our doctoral candidates are expected to travel and spend time at the other institution. It's also easier to do research here because English is very widely used and very widely spoken. In many places you run into a language barrier, beyond the immediate close circle. Of course Munich itself is a very international place, so that makes it easier for visitors to stay here and to work here.

HL: I would stress one thing Krish and Tsung-Yi already highlighted, that doctoral candidates hired at TUM are expected – and supported – to go to Duke and join our groups for a period of time. Such a two-way communication will really help them learn the culture in the United States and then bring that back to TUM. Also, they are able to build a solid connection with my team. Extensive communication across TUM is also a big advantage. I've never seen such "luxury" situations in other universities.

US: I think the Alexander von Humboldt Foundation, which Tsung-Yi mentioned, is one of the really great science institutions that we have in Germany, and very few countries have anything comparable. They're really great at bringing scientists together and getting international scientists to Germany. And the TUM-IAS was built on the model of the Humboldt Foundation, but they actually went beyond that – I'll say it again – by providing support for a doctoral candidate, and also by making an effort within TUM, across the different disciplines, to bring the Hosts and Fellows together through things like the General Assembly and many other events. I think Krish put it very nicely. It's in their genetic code to foster this bilateral and cross-disciplinary international research.