Hot chips and cool codes

Intel's Chief Technology Officer, Justin Rattner, talked to *Nature Materials* about the impressive advances in microprocessors and their applications.

What projects have you been involved with?
I spent the first couple of years at Intel, from 1973, working directly with Ted Hoff, who invented the microprocessor. Later, in the early 1980s, I was looking for opportunities in the area of artificial intelligence. I didn't see any immediate applications for Intel but stumbled across high-performance computing. So in 1984 I founded an internal start-up company to build high-performance computers. We built the first sustained teraflop machine for Sandia National Labs. In December 1996, just a part of the machine exceeded the teraflop mark and then was the fastest computer in the world for four years, which is probably some sort of record.


INTEL

What is your opinion on recent efforts in artificial intelligence?
I felt then as I feel now that the rule-based approach to artificial intelligence was probably not going to lead towards the wide availability of intelligent machines. It was just a clumsy way of expressing intelligence and clearly not the way the brain was working. I was more a fan of the statistical approach. But this took fifteen years before it really began to produce good results. Now we don't even call it artificial intelligence, we call it machine learning. Google is a big example of statistics-driven techniques that produces almost uncannily accurate results. We are going to see many more examples of machine learning and this is what people will think of when they talk about artificial intelligence.

How have processors evolved since the 4004 in 1971?
We have a lot more transistors. It was amazing that Ted Hoff and Masatoshi Shima were able to construct a working processor, the 4004, with about 2,300 transistors. We have gone from having barely enough transistors to having more transistors than we can usefully apply in a single processor core. That's why we went to multicores because the only way to improve the performance capability was to put that transistor budget into more cores. But there are many other dimensions. Clock speed has gone from maybe 100 kHz in the 4004 to 3 GHz. On the architecture side, we have gone from 4 to 64 bits. Internally, the machines are many times more efficient. We used to talk about how many clock cycles it took to execute a single instruction. Now, we talk about how many instructions we can execute per clock.

How do you see the future of silicon technology?
This is a very interesting time to ask this question. We have been working from a common transistor architecture for 40 years. Just this past year we've transitioned to a completely new materials stack-up in our 45-nm technology node. We've gone to a hafnium-based high-*K* dielectric material for the gate insulators and from a polysilicon gate to a metal gate. It is almost the dawn of a new era. In some ways Moore's law ended sometime over the past year for the silicon gate transistor. We literally got to a stage where we couldn't make it any smaller and a new Moore's law has begun around the high-*K*/metal gate. In future, we will move from a bulk device to a surface device. Intel calls this tri-gate and Robert Chau is co-inventor of the tri-gate transistor. That will be another dramatic change, and we will see a lot of these important moves in the future.

Are there any alternatives to silicon?
The question we face is how long charge-based electronics is going to last. Transistors are very effective devices for operating on electronic charge and it is hard to find alternatives that actually do much better. We are putting serious effort in looking at other quantum effects. In the meantime, we are continuing to work the device architecture; making the transistors better as we shrink them and look for something fundamentally better such as spin. Also, the wires are not getting better, and that has motivated our work in silicon photonics. Light isn't particularly good for computing, but for communications it is ideal. Looking at alternatives for wires is as important as looking at alternatives to transistors between the wires.

What are the exciting applications multicore chips will bring?
We started with the theory about half a dozen years ago that, particularly on the client side, applications were on something of a plateau. There just wasn't the same excitement we saw in the previous decade as media and the Internet came to the PC. But there was probably a set of applications out there that just required vastly more computing power than we had. So we began an effort to find out what those applications were and to understand their performance requirements. And many of those applications, if not all of them, present themselves in these three categories that we call recognition, mining and synthesis, which most people interpret as visualization although there are other kinds of synthesis that don't involve making pictures. Some of these more entertaining applications are things like super-resolution, where we can take a video let's say from your cell phone and then clean up the poor quality of the recording. We have become more and more enthused by this and now have a production processor in design, codenamed Larrabee. Larrabee is, I think, the first of the Intel products that takes us to that range of 10 to 100 times of today's performance and we believe will usher in a new era of applications for the PC.

Interview by Joerg Heber.