

Hardware Catches Up

Dr. Matthew C. Putman

An AI-enabled factory will be more productive and dramatic than the computer-enabled factory was, and will again allow hardware to stay one step ahead.

Six years ago as we were just starting our current company, Nanotronics, I and my co-founder (who happens to be my Dad and my business partner for 17 years) were on a plane seated next to a programmer from Nvidia, the premiere makers of Graphical Processing Units. GPUs are highly specialized computer chips that were originally designed to accelerate the processing of video information in a computer. The conversation we had led to a discussion about the future of computer hardware, and I became more excited than ever. Our new friend had us over to the Silicon Valley headquarters of his company, where we walked into a room that had the very latest demonstrations for their GPU processors. We saw video games with graphics so realistic and impressive that it made everything I had seen previously look simplistic and cartoonish. Just as impressively, the games reacted to my control button pushes instantaneously. After the tour we ended up in the company cafeteria where employees were playing the 1980's arcade games "Asteroids". This contrast struck me as funny, but the experience left me thinking about Nanotronics and reinforced that one of our primary goals should be to make semiconductor inspection as fast reacting and with a resolution jump as high as those amazing video games we saw in the demo room. It also made me realize something about the state of computer hardware and software in 2010. Advanced GPU hardware was just waiting for software to catch up. Games were cool and impressive, but certainly there were other things that such powerful hardware could and should be put to use. Were there medical, material, aviation and electronics advances that were on hold as we waited for software scientists and engineers to leverage this amazing hardware technology?

I did not have any real inside information, except for my personal experiences of having worked in factory instrumentation on and off since I was 8 years old. In 1982 my father, John Putman, started the company where I eventually went to work. Tech Pro did for the rubber and chemical industries something similar to what I became all the more inspired to do with our new company, Nanotronics. In 1982 video games like Asteroids were indeed one of the more impressive uses of that era's computer hardware. 1982 was the year that I got an Atari 2600 as a Christmas present. It was also the year that our new family business started putting the first PCs on the factory floor. In many ways it should not have been novel, but it seems that entertainment applications often lead and inspire pragmatic applications when it comes to hardware innovation. This was of course Tech Pro's advantage. My father watched me play Space Invaders and even with such a simple game it must have given him the confidence that computing power was ahead of software manufacturing applications in factories. Moore's Law was, as it still is today, also a wave for all of us to ride. The processors that impress us today will be superseded by even more impressive ones in 18 months. Moore's law has continued due to a thirsty public and an impressive response from the semiconductor industry. However, looking at the pace of software innovation it is apparent to me that software in no way obeys anything close to a Moore's law. In many software applications it is hard to even recognize a relevant improvement over

the course of a decade. Is Microsoft Excel exponentially better now than it was 10 or even 20 years ago? Even incredibly impressive technologies have very little improvement that is not hardware driven. A Google search was not any better than the hardware that enabled it in 2000 and in 2010.

While it is an exaggeration to say that general purpose, CPU processors were waiting for Tech Pro to come along to be as useful in factories as they were in games, I think that GPUs really did only recently find a killer app, and, maybe more importantly, this app may actually allow software to keep pace of hardware improvements. Artificial Intelligence (AI) is a decades-old concept that is now a very present reality, at least in narrow but powerful ways. Our lives are already driven by it, and in most cases machine learning techniques such as Deep Learning have pushed the capability of traditional CPUs to the limit. Interestingly the first proof points for superhuman AI also happened in the realm of game play. A recent example was made by the company DeepMind (now a division of Alphabet) that learned to play those same 1980's Atari games without explicit domain specific programming or supervision. This was only made practical due to the development and repurposing of fast GPUs. Now, these specialized GPU chips have stepped in and become the backbone for processing everything from facial recognition to instant language translation.

To me this brings up two questions.

How do we keep progress going if the pace of semiconductor technology advances slow, causing Moore's law to in turn slow or stop?

I'll elaborate on this in a future blog posting, but the answers to that could be a number of things from new materials and processes to developing reliable quantum computing architectures. The second question though brings me back to a place analogous to 1982 when factories were not yet using advanced data acquisition and computation that games were using.

Right now, why are factories not yet using AI to improve quality and incorporate rapid feedback to iterate design?

In some ways this second concern, and the one that I am interested in helping to solve, may actually assist in solving the first problem. Moore's law has advanced because people alone have pushed the process of making ever smaller transistors. But by using AI, it might actually become possible that the machines that we have created will also help us avoid the otherwise end of Moore's Law. I can't say precisely how this will happen, as that is the very nature of how a combined intelligence operates. It will be smarter than I am now. Still, we can imagine that mathematical problems that now require academic modeling and years to complete could be solved in real time.

We can also imagine doing this on any factory floor, promoting ever faster, better, and more flexible manufacturing processes.

It would be not be unimaginable to picture the AI expectations consumers have had from the perspective of the 1980's. Even in early Star Trek episodes much of it was part of the action. This is exciting. Right now my kids are asking our Amazon Echo questions, and because of the wonderful advances in voice recognition and many other machine learning areas it understands them in a way that the Enterprise's computer recognized Kirk's and Spock's voices. What we have not thought enough about is the intelligent factory that will make the next generation of these very computers and devices. I look forward to putting those in place and seeing what an AI-enabled factory looks like. It will be more productive and dramatic than the computer-enabled factory was, and it hopefully will again allow hardware to stay one step ahead.