Virtual Worlds

Future systems will have to rely on programmability.

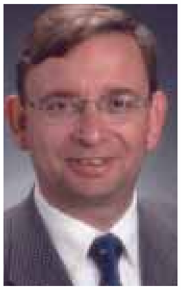

by Ivo Bolsens
CTO
Xilinx, Inc.
ivo.bolsens@xilinx.com

The wave of digitization is all around us. None of us has a crystal ball to predict the next killer application. But we will be surrounded by intelligent systems that monitor our health, protect our security, and increase our comfort. Thousands of sensors and actuators will generate and manipulate real-time digital signals to control our personal environment.

In less than 10 years, we will partly work and live in virtual worlds comprising an online alternate universe several times larger than today's Internet. These virtual worlds – like "Second Life" and "World of Warcraft" – require a massive amount of 3D real-time compute power.

The pace at which these trends are developing is astonishing. In September, for the first time in history, 9,000 IBM employees in Italy organized a strike in the virtual 3D world of Second Life. Driven by massive virtual reality data sets, complex 3D rendering models, and real-time animation, we will witness a data explosion and a need for a digital processing capacity that makes current data centers pale.

Mass Market of One

In the mass market of one, customers will expect standard products to be tailored to their specific user profile and requirements. For example, a programmable remote control will configure itself to deal with an individual's specific home network, comprising video and audio devices, security systems, and domestic appliances. FPGAs are ideal for such products: they are standard, they are programmable, and they offer performance that exceeds any other programmable platform.

Moreover, FPGAs have the advantage of getting increased performance by exploiting both the increase in clock rates as well as the increase of extra available silicon hardware as transistor sizes decrease. Processors are relying on increasing performance by boosting up clock frequency. However, at a time when the available number of transistors is outpacing the capability to use them efficiently, adding more hardware and using parallelism and pipelining to increase performance makes much more sense.

This parallelism allows FPGAs to react instantaneously on interrupts from the outside environment. It is clear that processor architectures are running out of steam with respect to further performance increases. FPGAs still have a very promising road in front of them.

The Case for FPGAs

The explosion of digital data to be processed creates a data-transfer bottleneck for traditional processor architectures. It is getting very hard to move data from one large global data or instruction memory to the processor data path. Increasingly complex caching schemes and pipelining try to deal with this issue. Today's modern processor data paths occupy less than 10% of the silicon die. The rest of the die is covered with caches and controllers to keep up with the data-transfer problem.

In contrast, FPGAs rely on a distributed memory architecture that supports the concept of keeping the data as close as possible to arithmetic and logic, while guaranteeing memory bandwidth orders of magnitude larger than traditional processors. This makes FPGAs especially well suited for high-end DSP and network processing applications.

In combination with software and tools, FPGAs will bring programmability, performance, connectivity, and energy efficiency to every electronic device.