

From 2009 to 2020: A history of developments in programmability

<http://www.embedded.com/columns/technicalinsights/219500635> By [Alan Gatherer](#)

[DSP DesignLine](#)

(08/28/09, 05:34:00 PM EDT)

Editor's note: This is the second of a multi-part 2020 Vision series outlining what the future may hold, as viewed by technologists within Texas Instruments. Click [here](#) for part 1 ("**Processor architectures: Where will we be in 2020?**").

Predicting the future is primarily an act of the imagination. However, digital signal processors are showing some strong trends and I think it is possible to predict what will happen in the next few years as we move towards the next order of magnitude increase in computational efficiency.

Here are my thoughts on the next 12 years.

2009: Multicore is here. With the increase in SoC architectures, single-core CPU devices have become more the exception than the rule.

2012: Network-on-Chip (NoC) arrives. A NoC is a high-performance device, which is really a grouping of processing islands connected by packet-based, point-to-point asynchronous communication highways.

2010"2015: Component-based software. The number of cores on a device is still fairly modest, and individual software components are developed for a single computational cluster by "component developers" and then "assembled" onto a multi-core system. Development tools for this methodology improve steadily as virtualization of hardware through middleware is driven by efforts such as the SCA (Software Communications Architecture) for SDR (software-defined radio). Auto generation of glue code between components becomes the norm.

2015"2020: Single program multiple data (SPMD). The component-based approach begins to fail as the number of cores reaches 32. Turning to techniques used in high-performance computing (HPC), the embedded software community develops the SPMD approach where a program can be compiled to run over multiple cores. While initially requiring explicit description of the communication flow, pragmas are now employed to enable the parallel nature of algorithms to be exploited by a variety of multi-core devices.

2015: The Death of the FPGA. An important footnote in the history of programmability is the demise of the FPGA. Small multi-core CPUs consume significantly less power as well as provide a richer set of mapping options for complex algorithms and communication patterns than does the distributed fabric of ALUs and LUTs that make up FPGAs.

2020: The CPU disappears. Spreading functionality across multiple CPUs drastically simplifies the silicon overhead on each CPU, and hardware-based OS support manages NoC traffic efficiently. Programmers are unaware of the communication between CPUs and can develop/debug code without having to know which individual execution units are involved. Programming follows more the overall flow of data than its individual parts.

The range of devices available in 2020 will be about the same as it is in 2009. In 2020, embedded DSPs will still be a heterogeneous combination of CPUs and accelerators. Even though programmers are unaware of the individual devices when programming, it will still be true that some devices perform certain tasks much better than others.

Since much of the value of SoCs is placed in the careful choice of peripherals, CPU and DSP manufacturers differentiate themselves by providing the best combination of different IP blocks and how they connect. In the end, the quality of development tools and application software support will determine the first-tier players.

About the author

Alan Gatherer is the CTO for the High Performance Multicore Processors group at TI and is responsible for all strategic development of TI's digital baseband modems for 3G wireless infrastructure. Since joining TI in 1993, he has worked on various digital modem technologies including cable modem, ADSL and 3G handset and basestation modems. In addition, he holds 60 patents and is author of the book "The Application of Programmable DSPs in Mobile Communications."

Processor architectures: Where will we be in 2020?

http://www.dspdesignline.com/guest_blogs/218600219 By Gene Frantz

[DSP DesignLine](#)

(07/22/2009 5:11 PM EDT)

I have come to the conclusion that too many of us have no clue where we are going with technology. Rather, we are just busily moving forward and don't know if we are even moving in the right direction. It would seem that with our extensive experience in traveling we would understand a basic concept " to travel to a distant place requires two points:

Where I am

Where I want to end up at

The same goes for technology " we need to know where we are going to move in the right direction. So, I have challenged several of our senior technologists to think about what the state of the art will be in the year 2020. You might say that we need to have 20/20 vision for the year 2020. I have invited a number of technologists to provide their point of view (POV) of what the state of the art in IC technology will be in the year 2020, and I'm interested to hear what you have to say on the topic. But, since this is my blog, I will have the first and last word on what the year 2020 will hold for us.

So, here are my first thoughts on the topic.

Processing elements will be single clock domains. After many years of assuming that Moore's law would give us faster and faster clock speeds, we have finally concluded that clock speed is no longer our friend. In fact, we should have noted that 15 years ago, but as we move forward, processing elements will be of the size that the CPU can communicate with all of its resources in one clock cycle.

Systems will be made up of multiple processing elements. Integrated systems will be made up of many heterogeneous processing elements, each being a "single clock domain" processor.

The processing elements will be arranged in a similar style as FPGAs today.

We will take advantage of the **third dimension**. Integration using stacked die techniques (SIP) will be just as common as fully integrated SoC.

All will be programmed with a **high-level language**. The development environment will have the ability to take into account all of the resources in the system. That is the microprocessors, DSPs, accelerators, peripherals, analog signal processors, analog peripherals, RF and other things I have forgotten about.

IC designs will consist of **smaller teams** (5 to 10 designers) taking a **shorter amount of time** (6 to 12 months) to do the hardware design. Reuse will be the norm. While I am at it, let me explain that there are two definitions of "Reuse":

1. I'll do such a good job on my design that everyone after me will use it.
2. I don't have time to reinvent the wheel, so I need to find something that is close enough to what I need to meet the schedule.

Unfortunately we use the first definition more than the second. Small design teams with short schedules will require us to use the latter definition. And, yes, there are companies already adopting this concept of reuse.

The bulk of the innovation will be in the software on top of the hardware.

Hardware will become part of the platform on which innovative designers will develop their ideas.

So, this is a sketch of how I see 2020. After a couple of POV papers from others at TI, I will come back with a conclusion. My colleagues will dive into topics such as programmability, tools and SoCs in the next few blogs. If you would like to share your view of 2020 with me, please comment or send me a private note.

[Gene Frantz](#) is TI Principal Fellow, Futurist and Business Development Manager, DSP.