

Future Shock – Semiconductors and Software

By
Gary Smith

One of the more interesting consequences in the move to Electronic System Level design, has been the Semiconductor Industry's reaction to the challenge. Keep in mind we have been putting microprocessors into ASICs since the late 1980s so the software connection isn't new. What is new is that the responsibility for embedded design has shifted.

In 1998 we moved to the 180nm semiconductor node, which provided the engineer enough available gates, 35 million, to design true systems on a chip. These designs became known as SoCs. Bryan Lewis, of Gartner Dataquest defined a SoC as an ASIC including at least one processing element. Prior to that we had been designing chip-sets in which one of the chips contained the processor. The other ASICs completed the system.

The real change occurred in 2006. By that time Semiconductor vendors were supplying SoCs to the Chinese consumer OEMs that were selling into the US market. These Semiconductor vendors were also supplying demo kits that gave the manufacturer the board layout and just about everything else needed to complete the electronics for many of the consumer products; the main example was the MP3 player. That year John Barber, of Gartner Dataquest, reported that Wal-Mart had a return rate of 65% of the MP3 players that they had sold for Christmas. The culprit was a confusing user interface. The result was that the Chinese OEMs requested that the US semiconductor vendors supply the software needed to run the MP3 players. All of a sudden the Semiconductor world was being held responsible for the SoC's embedded software.

About the same time it became clear that in their search for a more power friendly SoC architecture that the Semiconductor Industry had abandoned the high frequency Von Neumann computer architecture in favor of a slower multi-core architecture. What it took them a while to understand was that they had also abandoned the sequential Von Neumann Model of Computation, which had driven computer architecture from the late 1940s, to a multi-core computer architecture that had no Model of Computation. The initial response, from the semiconductor vendors was, "Not my job." Unfortunately that was also the response from the Software Industry. The Embedded vendors proclaimed that parallel computing had been tried, and failed, for over 25 years and that they didn't intend to waste R&D resources trying again. Microsoft claimed that a

jury rigged version of C, F#, would work just fine. That's when the microprocessor vendors started getting worried.

2007 was another pivotal year. The ITRS, after much debate over whether software was part of their charter, took a look at the software issue. The result was that software was included in the ITRS Cost Chart, an auxiliary part of the roadmap. The results showed that developing the software for an SoC had passed the cost of designing the silicon during 2007. It also showed that unless we developed some type of concurrent software development infrastructure by 2012, that the cost of producing a SoC would have a negative impact on the continuation of Moore's Law's curve. That resulted in the inclusion of software in the main body of the roadmap for 2008.

Bottom line is that we are in trouble. We recognized the problem five years too late and then were slow getting into action. We have four years to come up with a concurrent infrastructure and we have yet to come up with a Model of Computation. There is a lot of questioning the resources being spent on turning C into a concurrent language. We are spending billions of dollars solving the back end lithography issues for IC manufacturing, we are spending billions of dollars designing SoCs, and we are spending millions of dollars on the concurrent software problem. Something has to change.