

SOFTWARE DRIVEN DESIGN & CLOUD COMPUTING

Software is Driving Today's Hardware Designs

Interestingly enough, you generally hear from hardware engineers that software is driving today's designs. This is basically missing the point. Some of this comes from the fact that system modeling is often done using the C programming language. They miss the fact that the architects are modeling, not programming. If they modeled their designs using Legos, would you call it a Lego Driven Design?

Actually, I would not be surprised if, in the not so distant future software engineers will be complaining that hardware is driving their software development. In this world of power constraints, they would have a point.

What is going on is that we can no longer meet our power budgets unless both hardware and software engineers have a fairly deep understanding of the others' requirements. We can no longer build walls that separate the two functions. What we are starting to do is develop Software-aware Hardware/Hardware-aware Software. It just keeps on getting more and more interesting.

Cloud Computing for the Technical Community

The Power problem has now moved down to the server farms. In this case it's all about money, not battery life. The cost of electrical power to run and cool these servers is skyrocketing. It's getting to the point that the building needs to be specifically architected for the server farms. Every watt of power is becoming a big dollar investment. All of this means compute efficiency is becoming a must.

Presentations at ASP-DAC this year made the point that homogeneous computing has pretty much maxed out. Today's designs are almost always heterogeneous. The only argument at the conference was whether to use hardwired accelerators or applications specific processors (i.e. GPUs). In reality we will be using both. What we are seeing is servers becoming

applications specific, or at least domain specific (see the work at Berkeley on the **13 Dwarfs**). Not that there won't be a place for general purpose servers (i.e. Amazon's Cloud), just as the general purpose PC isn't going away. However for compute intensive applications, which EDA fits close to the top of the list, the servers will be tuned to the algorithms used. We are already seeing Google designing their own servers, and we are seeing a concentrated effort to develop servers tuned to EDA applications, primarily for graph based and dense linear algebra based algorithms. (I hope you don't expect technical cloud computing cost to be covered by advertising.)

We can't expect general purpose clouds to add these application specific servers to their cloud. Therefore, companies will rely on their existing computer farms. There is a market for vendors to build their own application specific clouds; however this would probably only be of interest for the casual users.