

EDP 2009 once again had an excellent technical program. Four main themes were presented: Network on Chip (NoC), the RTL ASIC Hand-off, Automation of Analog Design and Parallel Computing. Thursday dinner was held at the Monterey Yacht Club with Paul McClellan giving a thought provoking talk on EDA's future. Most of the presentations are available at <http://www.eda-stds.org/edps/>

NOC – WILL WE MISS THE BUS?

Will we Miss the Bus was a series of talks, and a panel, investigating the move from Bus-oriented architectures to Network on Chip. This was an especially good panel technically covering a topic that has become very hot recently. ASP-DAC this year showed the rapid switch from Bus architectures to NoCs in the Asian market.

RTL ASIC HAND-OFF

The good news on the RTL Hand-Off session was that Atrenta gave an excellent presentation and following discussion. Unfortunately they were the only EDA vendor presenting. If you are interested in my presentation on the decline of the Upper Mainstream companies, it is posted on the EDP website (www.eda-stds.org/edps). This decline is what is driving the move to the traditional ASIC hand-off, which is now being shifted from the Gate Level net-list to an RTL hand-off with the Silicon Virtual Prototype.

AUTOMATION OF ANALOG DESIGN

The fourth session, on Analog automation, was a bit of a surprise. Cadence, Magma and Ciranova presented, three of the four main players in the new set of analog tools that were introduced recently. What impressed me the most was that they all agreed that analog designers are the most conservative designers in electronics. Analog automation seems to still be at the stage digital automation was when Daisy introduced features such as copy command and rubber-banding. All in all not an impressive view of what is becoming the bottle-neck in SoC design.

AN HONEST LOOK AT PARALLELISM

Parallel computing took up one session and the two other Keynotes. One thing that is notable about the last year is the openness of the talks on parallel computing. Almost all of them have started out with the statement, "This is really, really hard stuff." That usually was followed with a synopsis of the dead ends we have encountered in programming multi-core/processor architectures. That's a far cry from what we were hearing just two or three years ago.

Chris Malachowsky, one of the founders of NVIDIA, gave a overview of what works (loosely coupled, multi-core systems such as server farms and cell phones), special applications that work, Embarrassingly Parallel programs (video and other data-streaming applications) and what doesn't work (everything else so far). He then went on to discuss CUDA, NVIDIA's approach to programs that have a large data-streaming content but do not fall within the Embarrassingly Parallel problem category. These all fit into their own Applications Specific category, yet range from oil & gas exploration, medical imaging, Financial applications and of course graphics. Another interesting aspect of CUDA is that it is language agnostic. It works with C and C++ now but will work with FORTRAN and OpenCL. Other languages can be added if need be. This is a recent trend with parallel computing. The actual language is no longer a large part of the discussion. On the topic of threads, NVIDIA's approach seems to be hard-wiring them as the safest way to handle them.

The "Are Threads Dead" session had advice such as "Threads – a Necessary Evil", you really need to know your algorithms, don't expect more than a 2x to 4x speed-up and threads only scale to 4 processors (actually we've seen 16 processor architectures that work). One recurring theme was the need for parallel programming development tools. Right now multi-core software developers are flying blind and, in the heavy weather that is parallel computing, that is a recipe for disaster.

Patrick Groeneveld, from Magma, gave the Friday Keynote. It was an honest look at the difficulties of writing parallel versions of EDA applications, specifically Logic Synthesis. Patrick presented a translation of Amdahl's Law in to an electrical circuit, which really helped out the hardware-centric audience. Personally it cleared up the emphasis being placed on the Concurrent Memory by the Super Computer community. That looks like one of the only ways to attack the initial resistance in Patrick's circuit diagram.

All in all EDP 2009 was once again an important workshop for advancing the state-of-the-art in electronic design.