

Broken Design Flows and Point Tools

Broken Design Flows and Point Tools

by Dick Selwood, IC Design and Verification Journal

Where do you go for help when your design flow is broken? **Wally Rhines** of **Mentor Graphics** wants it to be to him and his company. He feels that the EDA tool chain breaks every 2.5 process nodes (and has some convincing PowerPoint slides to back his case), and that 45 nanometre is the next inflexion point.

Stemming from this he argued, when giving a **Globalpress Electronics Summit** Keynote, it takes a broken tool chain to get engineers to adopt new tools. And who can blame them? Apart from the cost of purchase, it is hard work changing to a new design tool. You have to learn not just how to use the tool but also how to use it to take advantage of short cuts, building on the tool's strengths and working round the weaknesses, all the things you already know how to do with your existing tools. But you also have to build libraries, sort out the interfacing problems with other tools, and devise a sensible path for legacy code. Even if the tool works out of the box there is still a load of other things you need to do. It is logical to go on using the existing tool for as long as you can -- and engineers are nothing if not logical.

The argument is that the move from schematic to code, from RTL to high level languages and now from high level languages to ESL tools -- all these took place only because the users had no choice but to make the change to a new tool set if they were to continue to get working chips into production.

45 nanometre is the process point where the move to ESL will become essential, since without the ability to see and modify at the high level of abstraction that ESL provides, it will be impossible to make the changes needed to improve all the other steps along the flow to manufacturing. And there are points along the chain where the other tools are breaking already. Verification is becoming even more complex and time consuming, even with formal methods and techniques like assertions. Budgeting for power and then controlling power are significant problems. Place and route is obviously linked to power

decisions on the one hand, and on the other hand, to the issues of design for manufacture (DFM).

Of course Mentor, as you would expect, has the tools that will make ESL a reality any time now.

While Mentor (along with the other big boys) is working on trying to fix the entire tool chain, there are, as always, other companies beavering away at point tools that slot into the tool chain to provide help where it is needed. For **Tela Innovations**, the answer to the problem of DFM lies in grid topologies. Just a quick recap on the DFM issue: smaller process geometries, 45 nm or lower, require complex lithography stages and some clever tricks to make the structures in the silicon look like the patterns the design tools created. Nice neat rectangles in the designs come out as circular blobs on the silicon unless there is a clever trickery between the final design and the mask making. The trickery is effectively double- guessing what the lithography does to the geometric structures and working backwards to distort the shapes of the elements so that the lithography produces the ones you want. For example, drawing a four-pointed star may produce a square. This new stage in the design flow requires detailed knowledge of the process, is tremendously compute intensive, is currently limited to a specific process running in a specific manufacturing facility, and adds yet more time to the time between beginning the project and seeing first silicon.

Tela's solution for avoiding this problem starts with a defined topology with a strict one-dimensional on-grid layout, fixed pitches and widths, and all contacts and vias on the grid. The Tela authoring tool fits into the standard tool flow and, once set up with the appropriate process data, sucks in the cell level net list and spits out a Tela net list in a standard format, such as GDSII or LEF. Tela makes a number of claims for this approach. They say it results in 10-15% smaller die area and that simulation shows a 2.5X reduction in leakage. (A test chip was in fabrication in early April.) The simple topology, with features defined only in one dimension, makes it very suitable for double patterning, a lithography technique developed to get around some of the issues of using light-based lithography. It also reduces the physical variation that in turn produces electrical variation in the finished chip. Oh, and it effectively removes the fat volume of design rules from the designers' desks.

If you look at an SoC as a collection of IP linked together, then there is a need for interconnect. Using buses quickly runs into issues with signal propagation and timing. Point to point interconnect with multiplexers is greedy, using large areas of silicon real estate, and has timing issues, particularly connecting

blocks with unconnected clocks. In fact, timing is a significant and growing issue. At least two companies are addressing the timing issues, and both are university spin-outs. Silistix's technology is CHAIN and comes from work at the University of Manchester in England. They describe it as a delay-insensitive chip area interconnect and support it with CHAINworks tools that slot into the conventional tool chain. CHAIN is asynchronous, and the company claims that using the technology provides predictability in project development and reduces development time by up to 40%. This dramatic number comes from three sources: faster development of the interconnect, removing timing closure issues and, as the chip is simpler, reducing verification time. If that were not enough, using CHAIN will reduce overall chip power consumption (by up to 30%), provide increased performance (by up to 50%), and reduce manufacturing costs by up to 20%. (I don't think it has the cure for the common cold, nor does it make tea at the same time. Yet.)

Like Silistix, Elastix is a European university spin-off with offices in Silicon Valley. For Elastix, the university is the Universitat Politècnica de Catalunya (UPC) in Barcelona, Spain. Elastix says that variability in logic circuits means that system clocks have to run at worst case speeds, which equates to slow and inefficient chips. Instead, their approach of "elastic clocks" turns variability into an opportunity. The designer can choose to make the chip run at significantly higher performance or significantly lower power. Elastic clocks track variability and run close to the actual speed of the logic. When the logic runs slow (because it is at a corner of the die or the die is slow, the voltage is low, the temperature is high, or the logic values are such that they will run slow), then the clock runs slow, and vice-versa. An SoC can have blocks with rigid timing while others, and the interconnect, use elastic clocks. Elastix hasn't yet made a public release of its tools, although they are out in use with early adopters, but says the tool will fit between placement and clock tree synthesis in a classic design flow. Again, using elastic clocks will dramatically reduce the issues at timing closure.

Synopsys has discovered multi-core and parallelism and is working hard to get its products to take advantage of multi-core processors in multi-processor compute farms. Their feeling is that even if there are only minor improvements in the time taken to carry out, for example, verification, that will at least reverse the trend for longer and longer processing runs. And the company hopes for more than minor improvements in some cases. One product that is already seeing significant performance improvement is HSPICE 2008.03 for analog and mixed signal simulation. Running on one-core performance is already 3X the previous version, and it is greater than 6X on a four-core processor. As well as tuning existing products, Synopsys is promising that there

will be new products optimised for multi-core, multiple-processor execution from the ground up.

But parallelisation doesn't always mean multi-core processors, and Synopsys is also looking at ways to speed up the traditional linear path for aspects of the flow. The first fruit of this effort is the new Proteus Pipeline Technology, which takes the flow at tapeout and starts running the mask data preparation software soon after mask synthesis has begun, instead of waiting for it to finish.

So there we have some of the highlights from the wonderful world of EDA. Nothing earth-shattering today, but who knows what DAC will bring? If Wally Rhines is right, the next couple of years will see the long talked-up change from HDL to ESL, with the battleground changing from Verilog vs. VHDL to System Verilog vs. SystemC, and with vendors queuing up to provide the latest and greatest. If Wally isn't right, how are you going to design the next generation of chips?

Dick Selwood, IC Design and Verification Journal

June 10, 2008