**Chip design leaps forward**
By Purnima Gauthron, Sequence Design Inc.

Today, miniaturisation is one of the most relentless pressures in the electronics industry. Basically, consumers want ever smaller products. At the same time, of course, these same consumers want each successive generation of products to deliver ever higher performance, while consuming less power. And - naturally - they want them available faster and more cheaply.

If this sounds to you like a challenge not too far removed from squaring the circle, you'd be right. It's a tough nut to crack. And the onus of responsibility for doing so lies with the silicon vendors - the companies which design and manufacture the chips which lie at the heart of every almost every electronic device currently available.

In their efforts to provide viable solutions, silicon vendors have, over the years, made many advances. One of these is the steady reduction of chip geometry - the scale of physical design layout of a chip. Today, several companies are already working at geometries of 0.18 micron, and some are producing products designed at an even smaller scale. It is developments such as this which have enabled the emergence of concepts such as the SOC (System On a Chip). This, as it's name suggests, is a chip which integrates all necessary system elements - microprocessor core, memory, logic, communications etc - onto a single piece of silicon. Obviously enough, this approach is more space efficient than the use of discrete chips, and there can be performance and power advantages, too.

But while ever smaller geometries promise much in theory, they also - in practice - present a major design challenge. And so significant is this challenge that the current generation of design tools is, increasingly, inadequate for optimising chip performance. Today, in fact, it is no exaggeration to say that, at 0.18 micron, as much as half of a chip's performance is left on the table as a direct result of the implementation of current design orthodoxy. Worse, unless the industry accepts a radical reappraisal of this orthodoxy, there is little prospect of improvement.

But, before you reach for the scotch in despair, read on. Because there is good news: a new approach is available, and it represents a real step forward. It is an approach which has turned conventional design wisdom on its head to provide fresh insight and a new - and immensely powerful - design tool. It is a tool which will underpin the next generation of SOC designs.

What is this tool? To appreciate its significance, we need to be clear on one of the most important issues facing chip designers. This is timing closure - basically the process of achieving target chip performance after final physical layout. In today's immensely complex and high-speed designs, this is an extremely difficult task. However, the semiconductor and EDA (Electronic Design Automation) community has invested enormous amounts in developing toolsets to address the issue, and it can be done reasonably effectively.

At least, it can be done reasonably effectively on geometries above 0.18 micron. The problem is that, as geometries shrink to even smaller UDSM (ultra-deep sub-micron) levels, conventional wisdom is pretty clear - current physical design tools cannot cope. Why should this be? Well, consider the physics: the key to timing closure is to control the critical signal paths in a chip. This delay derives mostly from the capacitative loading of the interconnect on the logic which drives it. And herein lies the problem. All major physical design flows today start basically with the known logical structure of the design and try to "estimate" what this loading will turn out to be. They then add gate and buffer drive until they're sure they've overdriven everything which might cause trouble. In other words, the conventional model of chip design is essentially logic-centric.

The difficulty is that the "estimate" extensions - initially using wireload models, and more recently using placement-based wirelength estimates, don't work properly at advanced geometries. Below 180 nanometers, most wire capacitance is determined by neighbouring wires, and increasingly by the actual signals on those wires determined in the last stages of detailed routing. As a result, estimate-based methods invariably rely on severely pessimistic loading assumptions - otherwise known as over-design - to eliminate all possible timing violations.

How much over-design occurs? At 180 nm, a reasonable figure for estimate-based timing-closure strategies is that 30 to 50 percent of a chip's performance is left on the table - and more as geometries shrink. In competitive markets, this is a prohibitive sacrifice. At Sequence we call this corrupt by "construction."

So what's the solution? Simple: abandon those estimates. Nanometer timing closure requires interconnect-driven, not logic-driven methods. The fundamental problem is that uncertainty in wire capacitance leads to crippling over-design. The first point in the design flow where the wire geometry can be known without uncertainty is after detailed routing is complete. Timing closure problems can (and must!) be solved for high-performance design after - not during - synthesis and place-and-route. Interconnect-driven methods will consistently produce faster chips than estimate-based mechanisms.

But how can we change interconnect without new routing iterations? The key is that timing closure really affects only a fraction of the chip, typically a few percent. But *which* few percent? Since estimation methods can't quantify loading, they can't tell which wires are really the problem, and therefore must rework so much of the chip that the problem becomes large. But with exact interconnect and delay knowledge, careful optimization can deliver both fast chips and rapid design closure, without rerouting.

This leads to the radical corollary that all this rewriting of standard physical EDA might really be unnecessary. Not that synthesis and place-and-route are obsolete; gates still need to be created, placed and connected. But the core goal of timing closure will be addressed by post-layout, interconnect-driven timing optimisers, which will surgically restructure wire topology and drive, without requiring place-and-route changes – not by more synthesisers and more routers.

To enable designers to implement this paradigm shift in design philosophy, Sequence has introduced Copernicus, the industry's first interconnect-driven timing closure solution. Copernicus has full visibility into the actual layout topology. A powerful, post-layout optimisation tool, Copernicus uses a unique patented methodology to identify timing "hot spots" and then proceeds to surgically correct them in one direct pass, all the while maintaining the integrity of the routed signal nets. How does Copernicus accomplish this? At the heart of the tool lies the ExactTopology engine which orchestrates Copernicus's sub-engines: true 3-D extraction technology, advanced delay calculation and user-constraint driven static timing analysis. These work in incremental concert to continually ensure accurate interconnect modelling on all signal nets, associated capacitive loading and consequent timing delays. The result: elimination of setup timing violations and design rule violations in one smart pass through the tool, more aggressive chip performance and reduction in design size. (see figure).

A major side benefit to this approach is to have design teams continue to use their choice of synthesis, placement and routing flows and "dropping-in" Copernicus to close timing. This allows synthesis to concentrate on its strength: deriving efficient logic structures, rather than on its weakness: estimating post-layout coupling-capacitive loading effects. It also allows routing tools to work on producing better routability and die area, thereby extending the life of what these tools were originally meant to do.

At 0.25um and below, design teams are having a hard time meeting performance requirements due to the "guardbanding" of present tools. The Sequence timing closure approach allows designers to protect their current investment in synthesis, place and route. At the same time, design teams can avoid the trap of *estimating* routing effects. By operating after routing, this approach can see the exact effects of the 3D topology and make local, precise changes to unclog and fix real timing bottlenecks. This new approach inherently takes advantage of achievable silicon performance and produces faster designs, with far fewer design iterations.

Ref:  timing_closure_PG_V1.doc