

RDR or DFM?

Which is best?

Restricted design or design for manufacturing?: the debate at 22nm. By **David Abercrombie**.

Shrinking process dimensions have brought us to the point where it is becoming difficult to manufacture chips with reasonable yield and variability control.

Historically, manufacturers have reduced process variability to allow for successful production at ever smaller nodes. Now, because of lithography and material limitations, production teams are running out of 'knobs' they can turn to remove sources of manufacturing variability.

Reacting to that, foundries are looking to remove variability from the design layout; the concept being that if there is less design variation in what they have to build, they can build it with less manufacturing variation.

The question remains, however, is all process variability bad? More precisely, is all process variation equally bad? This debate has led to two contrasting approaches: restrictive design rules (RDR); and design for manufacturing (DFM).

The pros and cons of both techniques have led to internal debates at foundries and even to 'publicity wars' in regard to what various foundries plan to do in their attempts to capture market share. In reality, the optimal solution will probably be some hybrid combination. Let's take a look at why that is the case.

The RDR approach (figure 1) is to remove layout variation everywhere uniformly, in an attempt to make a very regular layout in all cases. It does this by imposing restrictions that enforce

regularity, consistency and uniformity in design. The concept therefore restricts designers from using 'creative' or unproven features in a layout.

Conversely, the DFM approach (figure 2) is to find and remove the particular variations that cause manufacturing problems. The designer can try any design configuration, but will only find out retroactively that some can't be manufactured.

It's obvious that the RDR approach seems much more restrictive to designers and appears to take away their freedom to be creative. However, when designers are done with an RDR layout,

they're done. But restrictive design tends to use more area, because designers cannot take advantage of irregular layout patterns.

At 22nm and beyond, area is a critical resource that can have a significant impact on the ability of a design team to implement the performance and features needed for marketability. Additionally, restrictive design is a relatively new concept, meaning designers must learn new design techniques and cannot make full use of their past design experience and skill. This shift in design techniques can create an

Fig 1: Traditional 'unlimited' feature design and a restrictive design of the same feature

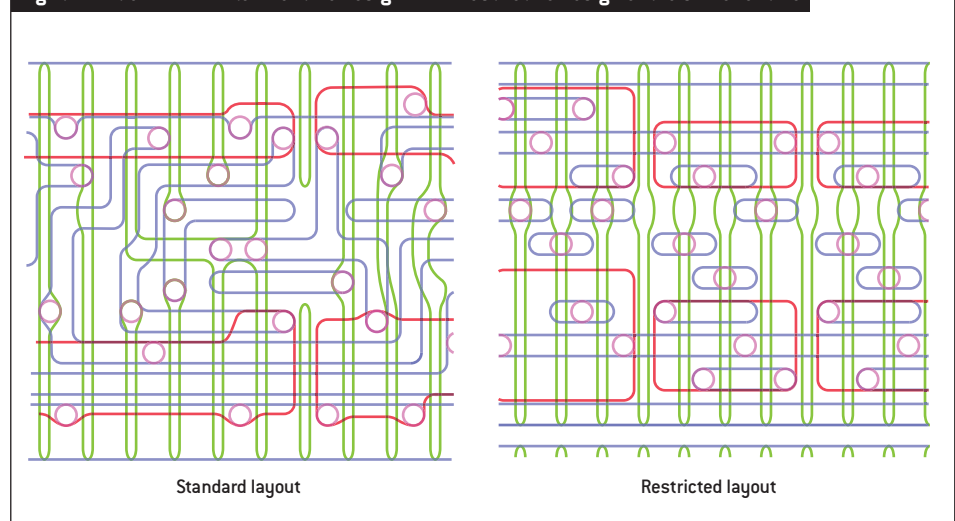



Fig 2: An example of DFM analysis



initial decline in productivity as designers adapt to the new requirements.

DFM, by contrast, allows designers a lot more freedom, but they must often perform a lot of iterative design when features are identified as non-manufacturable. Not only do they have to change these features, they also have to figure out what correction might work best. This iterative nature of DFM provides a lot less predictability in design time and/or effort, even though it allows designers to make full use of their experience.

However, there is another advantage to DFM that might be overlooked in a direct comparison. DFM simulation tools simulate the variability in designs and allow the design teams to use that knowledge in timing simulations. Even though RDR promises less variation overall, the design team can't know what the variation will be. DFM simulations allow designers to predict the specific variation they will get from each transistor.

On the other side of the equation are the foundries. What does a foundry have to do to develop either restrictive constraints or the checking capability to identify DFM hotspots?

It is typically easier for manufacturers to pursue the RDR approach: they only have to characterise the limited set of layout variation they will allow. With the short life cycles of technology nodes, foundries are forced to deliver design constraints prior to having a completed process. So using the RDR approach means all they have to do is make sure they can process those particular features. RDRs provide well bounded goals they have to meet to be able to

manufacture. Additionally, assuming they are successful at some point in time in manufacturing those features, they know they will always be successful in manufacturing those features throughout the lifetime of that production node.

In the DFM approach, by contrast, it is impossible for the foundry to characterise ahead of time anything the designer might do in a free form environment. Instead, they attempt to capture the process behaviour in a DFM simulation (such as lithography and planarity). A DFM simulation essentially allows the designer to simulate the process during design and verification, identify which features won't print or manufacture, then work to modify them until the design is manufacturable. The challenge is twofold: it's difficult to get those simulation models built prior



ABERCROMBIE:
 "... IS ALL PROCESS
 VARIABILITY BAD?
 MORE PRECISELY,
 IS ALL PROCESS

VARIATION EQUALLY BAD?"
 THIS DEBATE HAS LED TO TWO
 CONTRASTING APPROACHES:
 RESTRICTIVE DESIGN RULES AND
 DESIGN FOR MANUFACTURING."

to the completion of the process; and, because the process changes over time, the models must be updated and revised continuously.

Although requiring RDRs enables a foundry to establish a long term, predictable scenario, it doesn't take advantage of the fact that, over time, the process matures to the point where it's capable of doing more than it originally could on the same node. The restricted layout of RDRs doesn't allow the foundry or designers to take advantage of that maturation. However, as the process matures, DFM models are revised to match and the next designers will gain even more freedom in layouts, because the process is now capable of handling more.

Some foundries are betting that the safety of the RDR – guaranteeing their ability to produce the best performing chips in the shortest time – will outweigh the cost to the freedom of the designer. Others are attempting to apply as little design restriction as they possibly can in the hope that DFM 'safety nets' and process maturation will allow them to overcome early yield challenges, while gaining more business through design freedom.

A hybrid solution, in which the foundries implement some amount of layout restriction, but leave some creative freedom to the designer while applying the safety nets of DFM to capture the corner cases, will assure a sufficient 'pain to pleasure' threshold to both the designers and the foundries to optimise their business relationships.

As an eda vendor partnering with both these parties to reach success at 22nm and beyond, it's not rational to assume that you're only going to support one or the other of these techniques. It's also not going to be sufficient to support both, but in isolation. The tools will need to deal with both forms of constraints simultaneously, regardless of whether they're implementation or verification tools. Ultimately, the tools need to facilitate the adherence to the upfront restrictions, while easily identifying the hotspot issues within that design context.

Author profile:

David Abercrombie is programme manager for Advanced Physical Verification Methodology with Mentor Graphics

ne For further information on any subject visit: www.newelectronics.co.uk