

Impact of DFM and RET on Standard-Cell Design Methodology

Paul de Dood

Prolific, Inc., Newark, CA
www.prolificinc.com

Abstract

In this paper, we discuss the impact of Design For Manufacturing (DFM) and Resolution Enhancement Technologies (RET) on the creation and use of Standard-Cell libraries. We demonstrate through examples the various types of design-rules and recommended design styles that DFM imposes and how those requirements can be accommodated on a cell by cell basis to maximize yield while still maintaining minimum area.

1 Introduction

Design rule requirements are imposed to ensure a variety of quality metrics, including area, timing, power and yield. Of course, these requirements generally conflict. For example, the minimum width and spacing design-rules are set to be as small as possible to improve area. Any smaller and the design will suffer from unacceptable yield problems. Any larger, and the area of the design will grow with it – which also decreases the yield and increases the cost. The final values for the rules are generally not a “cliff” where making a rule smaller than a certain number results in zero yield, while exceeding a certain number results in zero defects. Instead, the rules are chosen such that the likelihood of a defect is low (but not zero)[1].

Since the 1970’s, any novice with a compactor could quickly determine the impact of a global design-rule change to the average area of a standard cell. However, it is far more interesting to see on a cell by cell basis how close each geometry can get to its preferred, highest yield configuration. If the vast majority of geometries can be created using the high yield rules with absolutely no penalty, it would be wrong to ignore the DFM requirements. Likewise, unless the yield effects of a single rule would be catastrophic, it would be foolish to increase the area (and cost) globally to unnecessarily increase a design-rule.

RET requirements introduce additional complexity. The use of RET can be grouped into 4 categories: Optical Proximity Correction (OPC), Phase-shifting, Sub-Resolution Assist Features (SRAF) (e.g., scattering bars) and Off-Axis Illumination. The mixture of RET strategies and manufacturing recipes will dictate what the design should look like. These requirements are

represented as a set of design-rules and recommended design styles.

For example, the use of OAI and scattering bars requires either lines to be drawn at an acceptable pitch, or spaced sufficiently far apart such that there is room to place scattering bars – resulting in certain forbidden pitches which cannot be resolved. However, the use of OPC results in the movement of edges such that the final pattern on the wafer is as close as possible to the design intent. Used together, the OPC edge movement can locally affect what the forbidden pitches should be for OAI.

In some cases, the penalty for not using DFM and RET enhancing methods results in large yield penalties. In those cases, the DFM and RET requirements become mandatory design-rules. However, the values are chosen for those design-rules such that yield degradation is an acceptable value. Generally, making the rules more conservative will further improve the yield.

The ProGenesis® software from Prolific enables users to determine the impact of design-rules changes by rapidly and automatically rebuilding an entire library. Also, every design-rule can have both a required and a recommended (or preferred) rule associated with it. By liberally using the recommended design-rules, users can determine how often a particular recommended rule can be applied without increasing area. Finally, all design requirements, including every DFM and RET rule can be individually weighted in importance to create a layout that is truly optimized for area, timing, power and yield.

Given a fixed amount of available space (i.e., slack) in a given layout area, there are potentially multiple yield-enhancing changes that can be made. For example, if there is free space in the design, then wires can be widened, overlaps can be increased, or space can be increased between objects, or some combination thereof. The ProGenesis® compactor uses a novel force-based approach to compaction. All design constraints, called forces, have an associated priority and weight. Each rule can have its independent priority and weight, which is used to determine which preferred rules are most important, and at what values their importance changes. By assigning the appropriate weighting function on a rule-by-rule basis, the optimal yield will be achieved[1].

2 Examples

There are many examples of how designing for manufacturing impacts the standard-cells. We will highlight some of those here. Each example was built using Prolific's ProGenesis® software, which implements all DFM and RET requirements desired as either required or recommended rules.

One potential yield problem occurs when the transistor gate extension is near diffusion, as shown in Figure 1. Note that in the figures, green is diffusion, red is poly-silicon, contacts are white and blue is metal. In this example, the yield is significantly improved when the gate extension is increased when it is near a diffusion geometry, as shown in Figure 2. This is a sufficiently serious problem that in some processes, the gate extension near diffusion is a required rule, not just a yield-enhancing rule.




Figure 1. Diffusion near gate extension




Figure 2. Increasing gate extension near diffusion

Misalignment of contacts and edge placement errors (EPE) due to lithographic effects can cause metal overlap to not fully enclose a contact. This can result in a catastrophic open-circuit, but smaller amounts of misalignment can also cause timing or failures in the field defects due to the increased resistance and electromigration through the contacts. One way of resolving this potential problem is to increase the metal overlap of contact wherever possible. Generally, when layouts are hand-drawn, the overlap is large in only one dimension and minimal in the orthogonal direction, according to the minimum design-rules, as shown in Figure 3. However, yield can be improved by increasing the metal overlap of contact wherever possible. One supported method of improving yield is pick the optimal contact orientation, as shown in Figure 4. Another method is to increase the metal overlap of contact wherever possible without increasing the area of the cell, as shown in Figure 5. In most cases, the overlap can be increased in all four sides, while in some cases, only 3 out of 4 sides can use the larger overlap rule. In this example, forcing the overlap to be large in all four directions would result in an increase in area while improving only a small percentage of the edges.

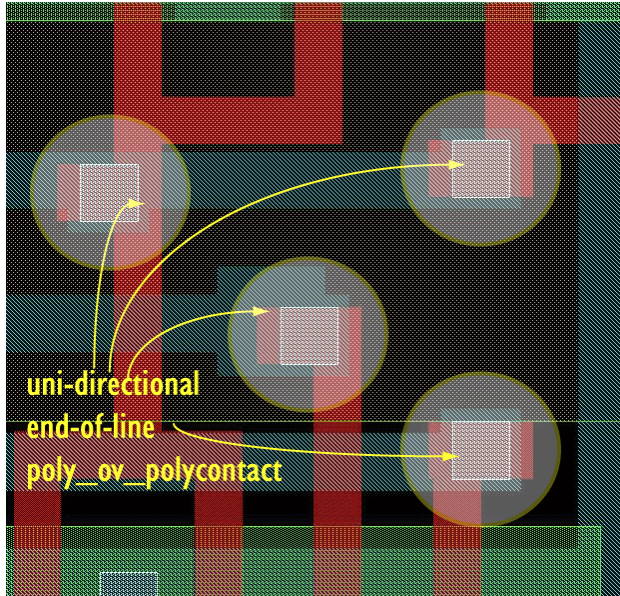


Figure 3. Minimum metal over contact overlap




Figure 4. Choosing optimum end-of-line depending on context

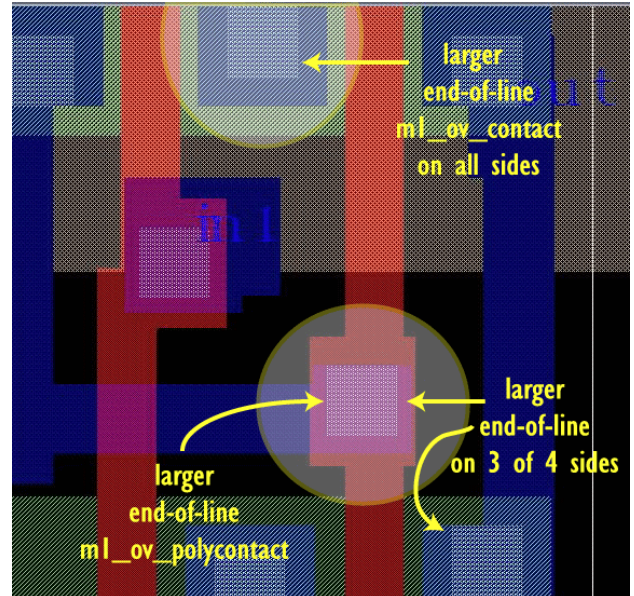


Figure 5. Increased metal over contact overlap

In general, the ability to set a geometry to its preferred width or spacing depends only on neighboring geometries. However, often the various preferred rules can conflict with each other. For example, suppose that you have a set of series gates as in the nand4 in Figure 6. Because the width of the cell is dictated by the contacted p-transistors, there is a small amount of slack available in the n-transistor region. One method is to simply give all of this slack to a single gate-to-gate spacing, which would make the spacing between those pairs of gates equal to: $min + slack$, where min is the minimum spacing and $slack$ is the additional slack available. The remaining gate-to-gate spacings would therefore remain at min . However, the critical yield issues often happen when geometries are exactly at the minimum spacing rule, and therefore only one of the three gate-to-gate spacings has improved. A better solution is to distribute the extra slack to eliminate as many minimally spaced geometries as possible, as shown in Figure 6. In this case, the gate-to-gate spacing for each of the three spacings is set to $min + slack/3$, and therefore none of the three gate-to-gate spacings is at the critical minimum spacing.



Figure 6. Minimum gate to gate spacing




Figure 7. Distributed gate to gate spacing

The use of Off-Axis Illumination results in the resolution being enhanced for certain pitches and degraded for other pitches. For the degraded pitches, scattering bars are inserted at the optimum pitch. However, because of the spacing between scattering bars and design geometries, there are pitches that are degraded but are too small to allow scattering bars to be inserted. These pitches are therefore disallowed and become forbidden pitches. To resolve these situations, the spacing between the geometries must either be decreased to an acceptable pitch, or increased such that scattering bars can be inserted. For example, Figure 8 shows two gates drawn at a forbidden pitch. By applying the forbidden pitch rule, the compactor was able to resolve this spacing by bringing the gates closer together into one of the acceptable ranges, as shown by Figure 9.




Figure 8. Gate to gate spacing at a forbidden pitch

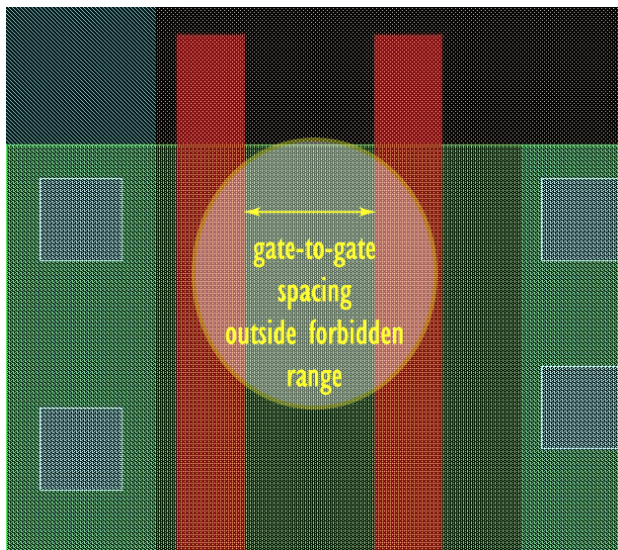


Figure 9. Gate to gate spacing at a legal pitch

The use of strong phase-shifting requires regions to be colored as either 0 or 180 degrees, where neighboring geometries must always have opposite phase. Depending on the layout, it may not be possible to assign a phase to each of the phase regions without assigning the same phase to two neighboring phase regions. This phase conflict must be resolved by either moving the phase-shifted geometries further apart, or by increasing their width such that they no longer need to be phase-shifted to reach the desired resolution. In the example of Figure 10, the entire poly layer is being phase-shifted – not just the gate layer. In order to solve a phase conflict, the width of the center poly wire is increased so that it no longer requires phase-shifting, which breaks the phase conflict, as shown in Figure 11.




Figure 10. Poly PSM phase conflict




Figure 11. Increasing poly width to fix phase conflict

It is common for the spacing rules for a given layer to be a function of width. However, at the 90nm node and below, the number of different width-dependent spacing rules is increasing and the widths at which they are triggered are getting smaller. At the 130nm node, typically only power-rails were wide enough to trigger a width-dependent spacing rule. However, at 90nm, even the metal overlap of contact or an equivalent width wire can have a different spacing requirement than a minimum width wire. The more common case is shown in Figure 12, where the spacing from the diffusion contact to the power rail is too close because of the width of the power rail. This is resolved in Figure 13, increasing the space between the contact and power rail.

Note that in Figure 13, the ProGenesis® compactor was able to maintain the 3 diffusion contacts by moving the entire structure up, rather than simply etching back the metal, which would have resulted in the loss of the bottom-most diffusion contact. This demonstrates the advantage of using a compactor to resolve these complex rules to find the global solution, rather than using a simple script or manually cleaning up the errors which may result in a valid, but sub-optimal solution.




Figure 12. Metal to wide metal spacing




Figure 13. Increased spacing to wide metal

Another method of improving the RET friendliness of a layout is to reduce the number of vertices. The complexity of the OPC depends on the number of vertices, as do the resulting data files. While often jogs are necessary to minimize the area of the layout, some vertices can be removed without an area penalty. For example, the drain region in the leftmost transistor of Figure 14 has a small notch due to the mismatch between the diffusion contact size and the transistor diffusion extension. This can be easily filled without area penalty, as shown in Figure 15.

This example also demonstrates the addition of redundant contacts, which is another DFM enhancement. Because contacts are often a source of yield problems, having redundant contacts can improve the yield and certainly reduce the resistance. In Figure 15, the source contact has been doubled to reduce resistance and increase yield, without an area penalty.

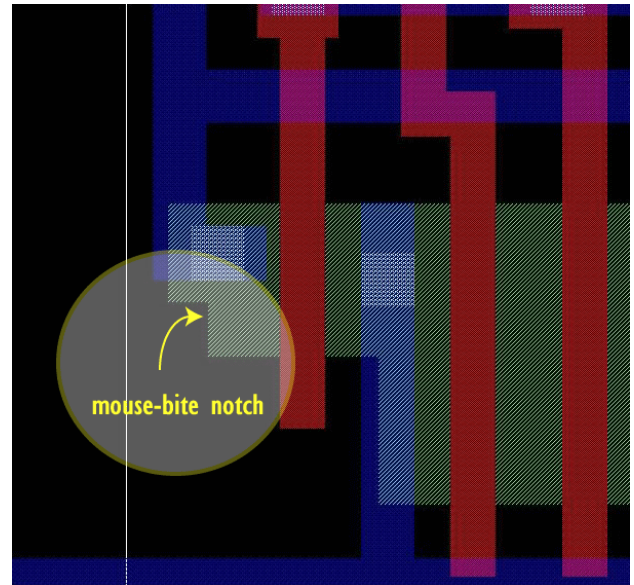


Figure 14. Minimum capacitance transistor drain




Figure 15. Reducing number of vertices (and resistance) on transistor drain

3 Conclusion

This paper has demonstrated through examples the impact of design for manufacturing and resolution enhancement technologies on standard cell design. Even though there are many additional requirements on the layout design, we have shown that the additional requirements can easily be handled by a DFM aware automated layout creation system. Furthermore, by experimenting with altering not only the required design-rules, but also the recommended rules, the smallest and best possible yielding layout can be created and the standard cell library can truly claim to be Designed For Manufacturing.

References

- [1] J. A. Torres, D. Chow, P. de Dood and D. J. Albers, "RET Compliant Cell Generation for sub-130nm Processes", *SPIE*, 2002.