Volume 12 Issue 02 Published June 17, 2008 ISSN 1535-864X DOI: 10.1535/itj.1202.05

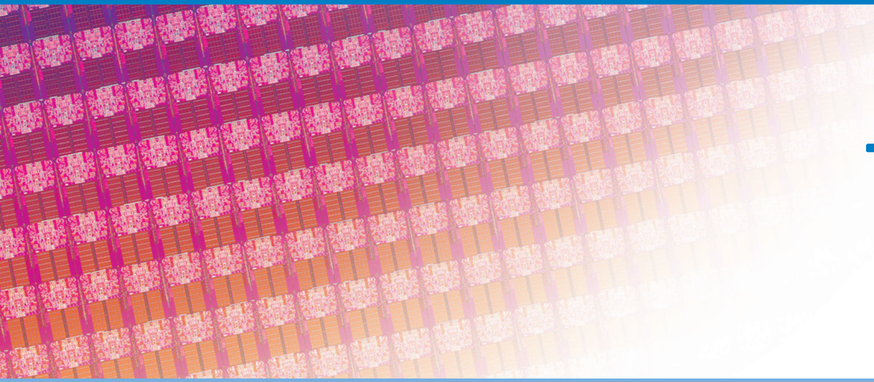

Intel® Technology Journal

Intel's 45nm CMOS Technology

45nm Design for Manufacturing

45nm Design for Manufacturing

Clair Webb, Technical and Manufacturing Group, Intel Corporation

Index words: DFM, 45nm, design rules

ABSTRACT

Co-optimization between design and process is required for a highly manufacturable process technology. This paper discusses this co-optimization and how it meets the challenges for maintaining Moore's Law while delivering new processes and designs capable of fast ramp to high yields. Poly is one of the most critical layers for control of variation, and it needs the most restrictive rules. We show the change in poly rules over the last few processes to illustrate how rules have changed to meet manufacturing requirements. The variation, density, and yields on the 45nm process show the success of this Design for Manufacturing (DFM) methodology.

INTRODUCTION

The difficulties in continuing Moore's Law with the lack of improvement in lithography resolution are well known [1, 2, 3]. Design rules have to change and Design for Manufacturing (DFM) methodology has to continue to improve to enable Moore's Law scaling. In this paper we discuss our approach to DFM through co-optimization across design and process. We define the design rules for a new technology early in the definition process well before the technology development is complete. This early definition of design rules allows the design to start in parallel with the technology development. Early accurate modeling of the design rules and layout is a key to making this process successful. The design rules must meet the requirements of a highly manufacturable process at the beginning of the production ramp for the first product.

Our DFM goals and methodology are different from those of some other manufacturers. The basic rules for drawing transistors, other layers, and DFM requirements are not separate in our definition process. We have a few guidelines such as suggesting that designers use redundant vias where possible, but most of our DFM requirements are included as required rules that all designs must meet for all layout. Some other companies have simple basic layout rules and provide other rules that are guidelines or suggestions for changes to layout or design that would improve manufacturability or reduce variation. Designers make tradeoffs for area and cost to decide if they will implement these guidelines. Some of their DFM changes are available only after analysis of the initial layout. Another difference in our methodology is that we tend toward adding rules to prevent something in layout that might affect design, instead of depending upon modeling of product layout to find problems. We build the requirements for manufacturability and low variation into the basic design rules as hard design rule requirements. Product design starts in parallel with technology development. Design rules must not be changed significantly after design work starts. Our DFM methodology depends upon modeling to define the rules very early in technology development. Therefore, our methodology ensures that products are ready for ramp-up in multiple fabs to high-volume and high-yield manufacturing without changes. This may make our design rules more complicated than those of some of our competitors, but this methodology ensures that all our products are capable of high yield when they tape in their first stepping.

Co-optimization across process and design is required to ensure we understand and balance all requirements. By co-optimizing design and process early in the development cycle, we arrived at a set of design rules that met 45nm process and design requirements. The thoroughness of this early work resulted in these rules being stable through the development cycle, which led to the successful insertion of 45nm technology in high-volume manufacturing, ensuring the continued march of Moore's Law.

The poly layer is the most critical layer for control of variation. Due to this need to control variation, the poly layer was the first layer on which restrictive design rules were used. We need to build the requirements for minimizing variation into the rules. How the rules have changed for the poly layer shows how designs have changed. The analysis of the impact of changes in the poly rule shows how we consider design and process needs in defining design rules.

DESIGN FOR MANUFACTURING GOALS

The first requirement for a new technology is the continuation of Moore's Law. A new technology should have twice the number of transistors. This is increasingly difficult with the lack of improvement in wave length for lithography and in transistor scaling issues. Keeping manufacturing variation in check as we improve transistor density is a scaling challenge as discussed in many papers and conferences [4, 5]. Transistor channel length is a major focus for controlling manufacturing variation. The variation must scale with the poly pitch for the new process.

The second requirement of a new technology is the ability to ramp very quickly to high volume with multiple designs in multiple fabs. The design and process have to be manufacturable at the beginning of the ramp. Design rules have to be defined early in the process development work to allow product design to be done in parallel with the process development. There must be no major changes to design rules late in process development or during manufacturing ramp. Predictive modeling of the rules must be done well before process development is complete. The design rules may be conservative, to ensure that the design is very robust, but they should not be too conservative, to ensure that we derive the maximum benefit from Moore's Law scaling. The challenge is in defining this optimum robustness.

The third requirement of a new technology is for its yields to be as good as, or better than, the previous-generation technology; and for its learning curve to be as fast as, or faster than, the previous-generation technology. Judgment is required to define the technology to strike the right balance between the difficulty of the technology and the impact of the DFM requirements on the design. Some breakthroughs may be required to meet design requirements, but these cannot be so difficult that they slow the yield learning goal.

CO-OPTIMIZATON

Lithography, optical proximity correction process, and design requirements all need to be understood in defining a new technology. Co-optimization to balance requirements for all of these areas is needed to define a very manufacturable process and design. The different needs of each area must be balanced in the context of different constraints. The design groups want small die size for low cost, but they also want design work to be easy. Flexible design rules are needed to allow designers to optimize the product with minimum effort and area. Lithography engineers also want simple design rules but cannot allow unlimited flexibility. OPC engineers want layout to be very predictable so there are no hotspots that are caused by the use of unexpected combinations of rules. In the definition process, all of these different goals are considered to find an acceptable solution. The ability to consider all aspects of the problem from design to high-volume manufacturing, within an affordable envelope of multi-dimensional constraints, is the key to good DFM solutions.

Co-optimization is done from the beginning with a multi-disciplinary team that includes experts in design, lithography, OPC, and processing. We have design experts in the process development organization who closely interact with the process development engineers, and we also include representatives from the lead product designs and CAD tools' engineers in the definition process. The discussions between these experts start about four years before the technology begins production. There was a significant increase in the co-optimization work for the 45nm technology. More layout studies were done early in the definition process, and there was an increase in the modeling of design rules.

The goal of co-optimization is balancing the risk and difficulty between design and manufacturing. It may appear to designers that our solutions are not balanced, since design is more difficult on each new process. The reality is that the process and patterning choices are often limited. The lack of improvement in resolution of the lithography tools limits patterning options. DFM definition often has to make the choice that will have the least impact upon design. A solution with no impact is not possible, however. The need for low costs may result in changes in design rules for smaller die size or better yields to offset concerns about increased design effort. Judgment about manufacturability will take priority over concerns about increased work for designers.

DESIGN RULE DEFINITION

Modeling is the cornerstone of today's design rule definition process. Building predictive models for the new technology is one of the challenges for a new technology. The rule definition work is front-loaded, with the rules defined before technology development is complete. Figure 1 shows the elements of the design rule definition process.

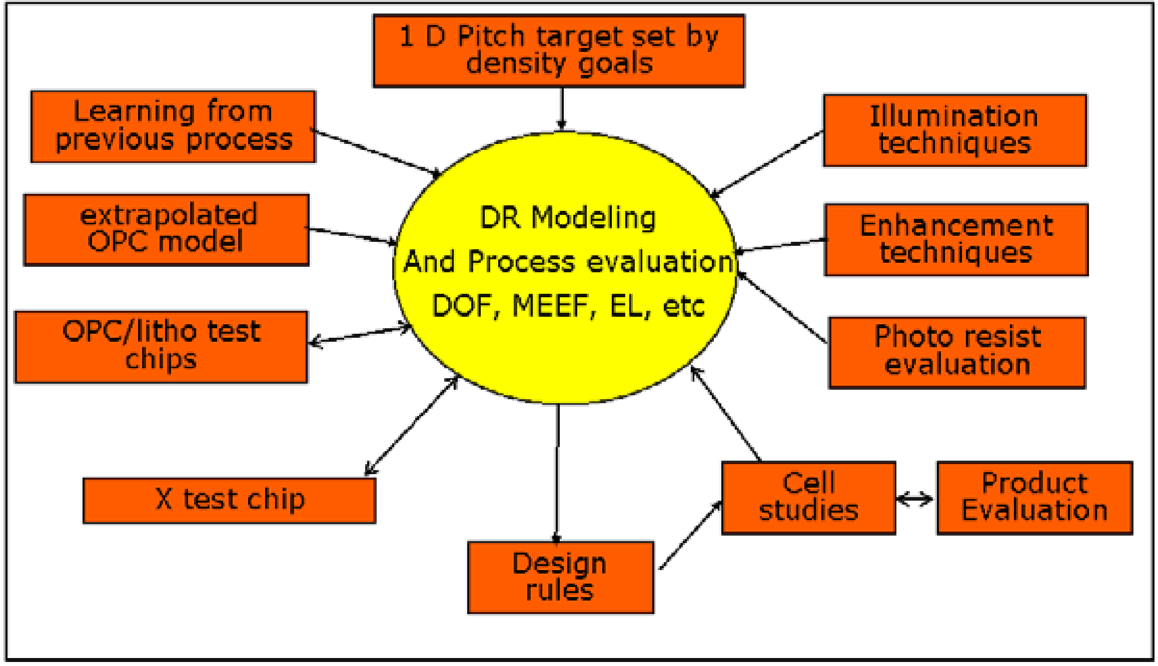

Figure 1: Design rule definition

The transistor density scaling goal drives the 1D pitch requirement for the key layers like poly and Metal-1 (M1). This 1D pitch must scale by 0.7X per process generation. 2D rules such as line end-to-end space or minimum line length can be more difficult to scale than the 1D rules. Isolated and wide lines may have scaling problems. The techniques needed to meet the 1D scaling requirements may make scaling the 2D rules more difficult. A change in illumination technology to get good minimum line width and minimum space may not allow the same scaling of rules for wider lines, or it may make scaling of rules different for X and Y directions.

Design rules are not changed after the beginning of manufacturing ramp. Learning about difficult process issues feeds forward into the design rule definition of the new process. Some design guidelines on previous technology may become hard rules on the new technology. Some structures that caused significant process problems and/or required significant process changes may be eliminated by new rules. Device or interconnect models may be simplified by the elimination of parametric variables caused by simplified design rules.

Modeling of design rules starts with extrapolation of OPC models from the previous technology. New lithography tools, illumination techniques, enhancement techniques, and resists are evaluated to determine the best method for 1D scaling, and to understand the changes needed in other rules. Learning about the capabilities of the new tools is a continuous process during rule definition. Typical and worst-case layout topologies are analyzed to evaluate process issues like Mask Error Enhancement Factor (MEEF) and depth of focus. Test reticules are created to calibrate the models.

Cell studies are done using the rules generated from the modeling studies. Data are extracted from designs on older technologies to understand the requirements for critical layouts and how rule changes might affect design. Standard logic cells, register files, SRAM bits, and metal routing are all included in the layout studies. As rules mature, product groups are included in the evaluation. Evaluation of the patterning capabilities and the impact on layout is a continuous and iterative process until rules are final. Important layout topologies identified by design are analyzed by using the models, and they are included in new test reticules. Design rules that limit meeting the transistor density goals are evaluated with the models to understand if rules need to be changed or improved.

Figure 2 shows a typical simulation from a study of a design rule. This simulation studied the process margin for line end-to-end space as a function of line length. A high MEEF is an indication of poor process control. Some line lengths have insufficient control of the end-of-line space.

This can create line-to-line shorts. One option to fix this would be to change the illumination as shown by the different lines on the graph. Another option would be to create a design rule that does not allow line lengths of 0.2 to 0.3u. If a rule change is proposed, cell layouts are done with the new rules to understand the impact on transistor design. A study may be done on data from designs on previous processes to determine if the design rule being changed is commonly used. If the rule change is shown to cause a significant change in area, we would consider other pattering solutions.

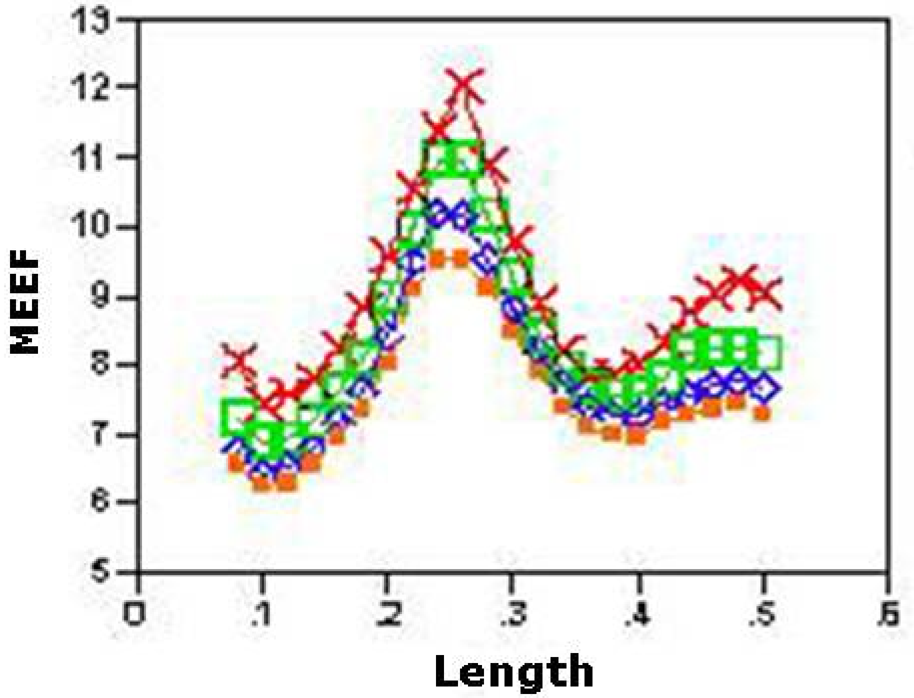

Figure 2: Length vs. MEEF

The main process development vehicle is the X-chip test vehicle that includes large SRAM designs, process and design test structures, and process-sensitive circuits. There are earlier mask sets that include some of the rules and features of the new process, but the X-chip is the first mask set that has large circuit blocks with all of the rules. The processing of the test chip is used to validate the rules, not define the rules. There may be a few changes in rules depending on what we learn from the test chip. The number of rules that are added or changed have to be limited, because product designs will have started before the test chip is processed. As the development process continues, the justification for design rule changes becomes more difficult. By the time production starts, a design rule problem has to be fixed by process unless it is impossible to fix that way.

Figure 3 shows our trend in the number of rules over time for the 45nm process. The timeline is relative to tape-out of first design on the technology. In addition to added new rules, there were some minor changes for better and worse rule values. The number of rules increased during the design and layout of the test chip as the modeling work continued and OPC flows were developed. There was a small increase in the number of rules after processing of the test chip. These changes were not all due to patterning issues. New understanding of the transistor of metal processing can create the need for new rules. By the time masks were created for the first product, the rules were stable. There were only a few changes in design rules after the first design started. The 45nm rule stability was good, but there were more changes than desired. On the 32nm process we had fewer changes during the test chip design and after the start of the first product. Process development continued during debug of the first design, but the rules did not change during this time. Process improvement continues through the life of the process to reduce defect density, and cost, but this is done without changing design rules.

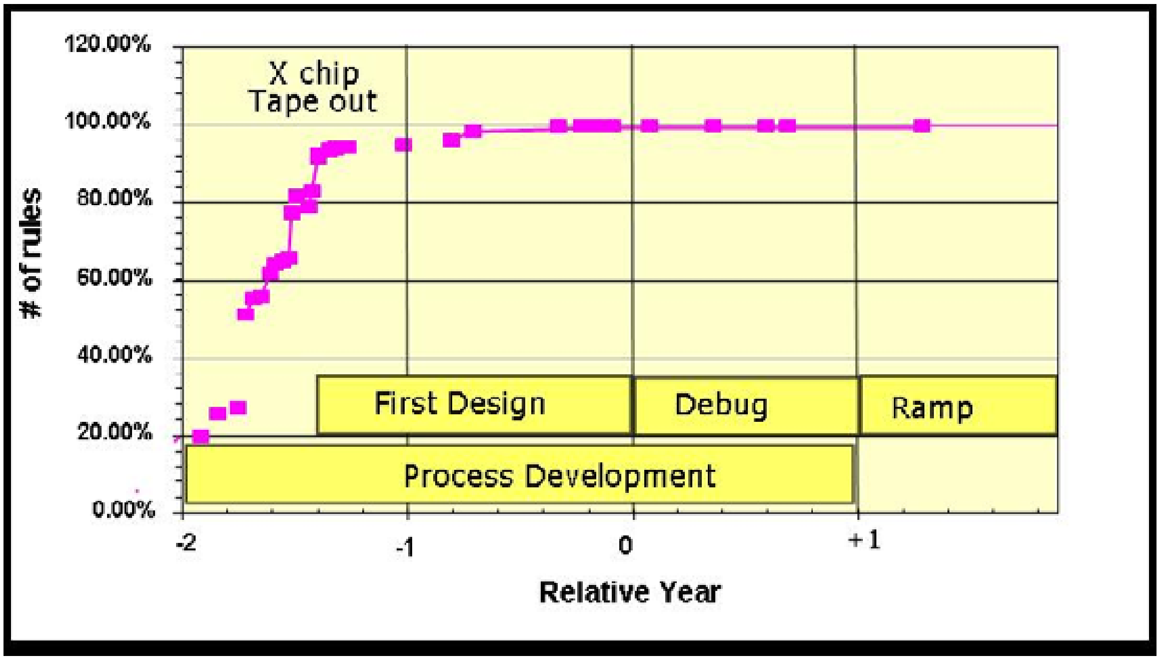

Figure 3: Changes in number of poly rules

EVOLUTION OF POLY RULES

The poly layer is one of the most difficult layers to pattern and process. Control of poly CD is one the most critical requirements in the process, due to its affect on transistor performance and variation. Poly CD control must scale for the new technology to keep the percentage variation of the channel length constant. Contacted gate pitch is a big factor in SRAM cell area and logic transistor density. These critical requirements make poly the first layer to need new pattering solutions and design rules. In the following sections, we show how poly design rules have changed over the last few processes and how DFM methodology has evolved and has been used in the definition work.

130nm Process

Our 130nm process had simple rules. There was limited early modeling of layout. Layout had random combinations of poly widths, spaces, and device orientation. Products had some issues with poly corner rounding that affected small devices. This issue was helped by DFM guidelines for layout of small devices, but the guidelines came after initial design of lead products. There was limited involvement from product design engineers in the early design-rule definition. Rule definition was primarily simple scaling of rules from previous technology generations.

90nm Process

The 90nm process included more restrictions for poly layout, and the number of poly rules increased by 47%. All devices had to be in the same orientation except for the memory bits. The difference in printing of shapes parallel and orthogonal to the scan direction was one reason for this change. Poly-over-field-routing was allowed in the X or Y orientation. The memory bit shown in Figure 4 is a unique topology with transistors in both orientations that could be modeled and characterized to account for any difference due to device orientation. Poly corner rounding and modeling were analyzed and modeled early in the definition process.

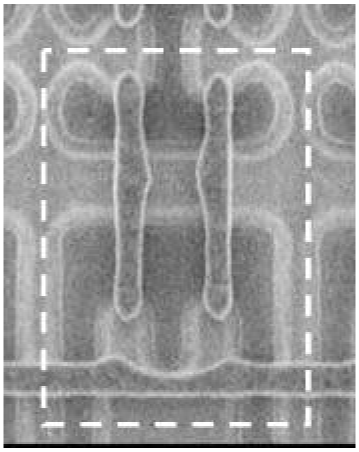

Figure 4: 90nm SRAM bit layout

The main design impact of the one-device orientation is that a block of layout cannot be placed in two orientations, rotated 90°. I/O buffers on the top and right edges for the die must have unique layouts, even if circuits are identical. This added some layout effort, but had no affect on die size. Since there was no die size impact, the rule change was a better solution for random layout than trying to model differences due to device orientation, as was done for the memory bit.

65nm Process

The number of poly rules increased by 65% for the 65nm process, and rules were more complicated. Rules changed to allow the use of phase shift masks. All devices including the memory bit had devices in one orientation as shown in Figure 5. This layout is almost ideal for patterning with simple rectangular shapes on all layers.

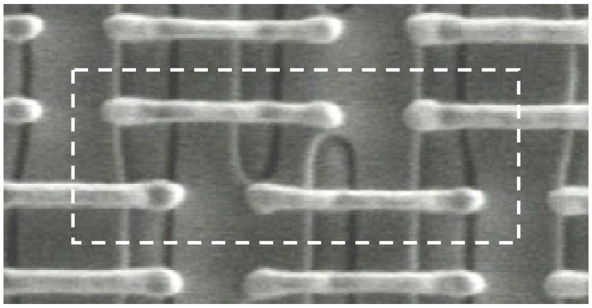

Figure 5: 65nm SRAM

There were differences in poly layout rules based upon pitch and poly space and orientation. The complex rules had little effect on transistor density. Many of the new rules affected special cases that did not occur often or were easy to fix. For example, layout of a minimum width device has some new rules for end-cap, but design did not use a lot of minimum width devices. Where they were used, the end cap rules were usually easy to meet due to area being limited by other rules.

Cell layout studies and modeling of layout increased for this generation for worst-case structures. The layout was random and it was difficult to determine all worst-case layouts.

45nm Process

Meeting the transistor density and process requirements of the 45nm technology required significant changes in design rules. The high-K metal gate transistors on this technology are the biggest change in transistors in 40 years [6]. 193nm patterning tools were needed to minimize cost and risk. Poly had to be printed using 193nm dry tools and still had to meet our need for 0.7X scaling of the pitch. Variation could not increase. The early modeling work increased significantly from the 65nm process. This included earlier involvement of the OPC experts in the design-rule definition and evaluation phases. Our goal was a more comprehensive evaluation of rules and layout topologies through modeling.

There were some changes in poly layout on the 90nm and 65nm process, but the layout has remained very random. Design could have used any poly pitch >= minimum pitch, and any channel length >=minimum was allowed. The different channel lengths could be randomly mixed. Transistors were in one orientation, but poly routing could be in either orientation. Corner rounding of poly close to devices could impact transistor performance.

Early in the definition work we asked if poly patterning in logic could be similar to the poly patterning in the SRAM bit introduced in the 65nm generation. One of the big concerns for making poly layout more regular was limiting the channel layout choices available to design. Design had always had few restrictions on the channel lengths. We wondered if this freedom was necessary. Figure 6 shows the channel lengths used in one 65nm design. Most of the devices at 0.10u and 0.11u are in the SRAM bits. The few devices at longer channel lengths were primarily in analog and I/O circuits. 99% of the devices in random logic had minimum channel length or minimum + .01u. The main reason for using longer channel length in logic was to reduce device leakage. Channel lengths can be limited in logic as long as there are options for low leakage devices. Device leakage is strongly dependent upon Le, so a very small change in Le can reduce leakage by 3X. Higher voltage can also be used to reduce leakage. There were some circuits where a longer Le device had to be replaced by two or more minimum channel length devices. Based on analysis of data like this and other layout studies, the design was partitioned into three groups: logic, analog, and SRAM. Analog and SRAM are treated as special cases.

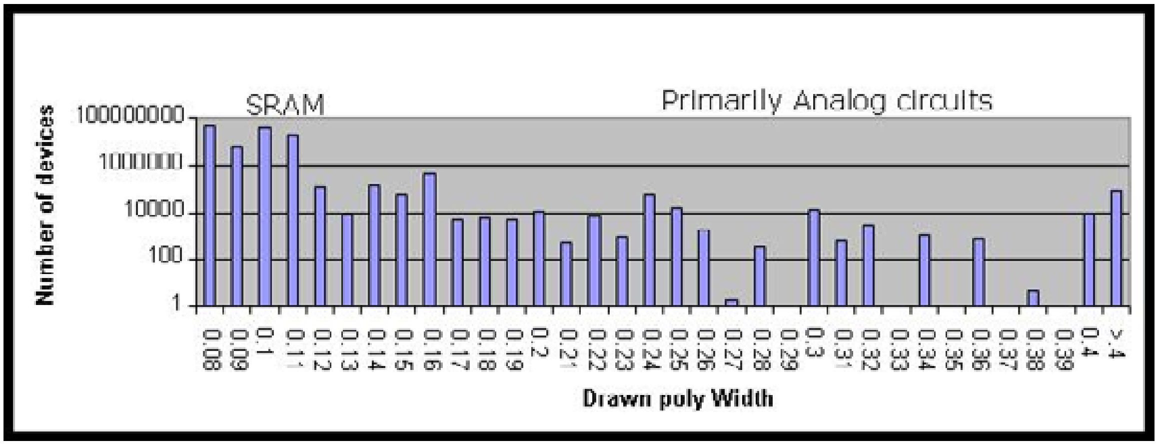

Figure 6: Product channel lengths on 65nm design

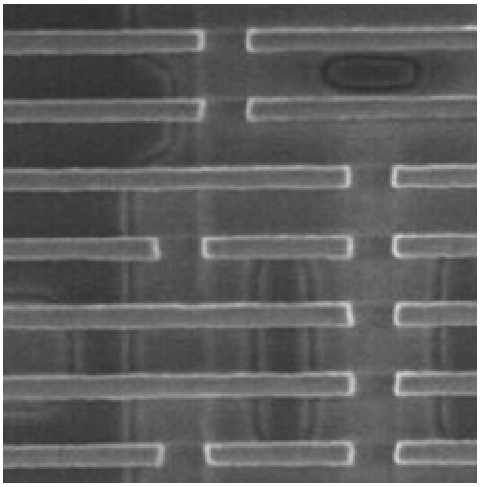

Logic

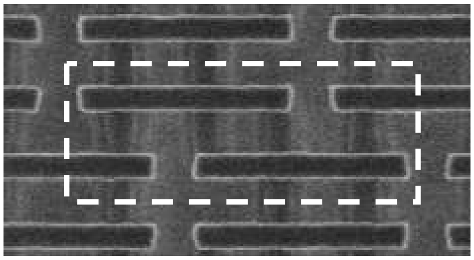

SRAM

Figure 7: 45nm SRAM and logic poly layout

The 45nm process has trench contacted-based local routing. This eliminated the need to use poly routing orthogonal to the transistor gate and the wide poly used for poly contacts. Layout studies were done to understand if one poly pitch was possible and how this affected rules for other layers. The local routing and limitations on transistor channel length had allowed logic poly layout to be one pitch and one direction as shown in Figure 7. The number of layout rules for logic layout was reduced by 37%. This reduction is not large, as the simple layout might indicate, because the poly rules include rules for poly spacing to other layers, and there are several rules for end-caps and poly end-to-end.

OTHER LAYERS

The regular layout for poly also simplified patterning of contacts and M1. The contact pitch is the same as the gate pitch. M1 parallel to the gates has the same minimum pitch as poly.

Very restrictive rules similar to poly rules were not needed to meet process or patterning requirements for other layers, but the rules for some other layers are complex. The metal pitch choices do not need to be as restrictive as the poly layer pitches, since metal variation requirement is not as tight. The metal variation has less impact upon path delay than the poly CD variation. Metal layer design rules have some types of design rule restrictions used on poly layers on previous technologies. Printing of isolated metal lines was one of the issues on 45nm technology. Rules were added to restrict the use of isolated lines. Product layout uses several metal widths and a range of spaces. It is difficult to limit the width and pitch choices for metal, due to the need for wide lines and spaces to optimize RC delays, capacitance, and power delivery. CAD tools must change to support the changes in metal rules. One of the learnings from the 45nm technology was that the CAD tool work needs to start earlier.

45NM MANUFACTURABILITY

Standard Cell Area Scaling

The first goal for any new process is to maintain the transistor density scaling trend. Figure 8 shows the standard cell area scaling trend. This analysis used a large standard cell library from our microprocessor designs and is weighted by typical cell usage. Transistor density scaling has followed the 0.5X density improvement per process generation. The 45nm process meets this trend despite the change to more complicated and restrictive design rules.

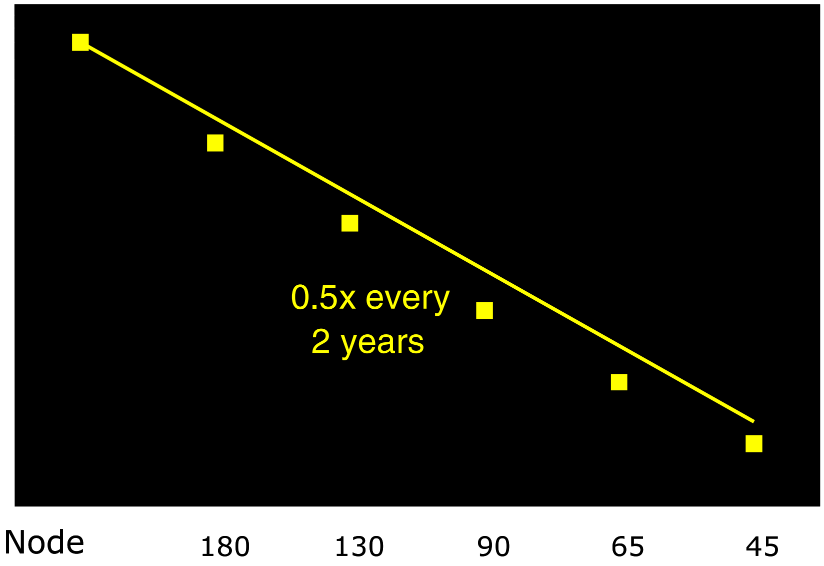

Figure 8: Standard cell density

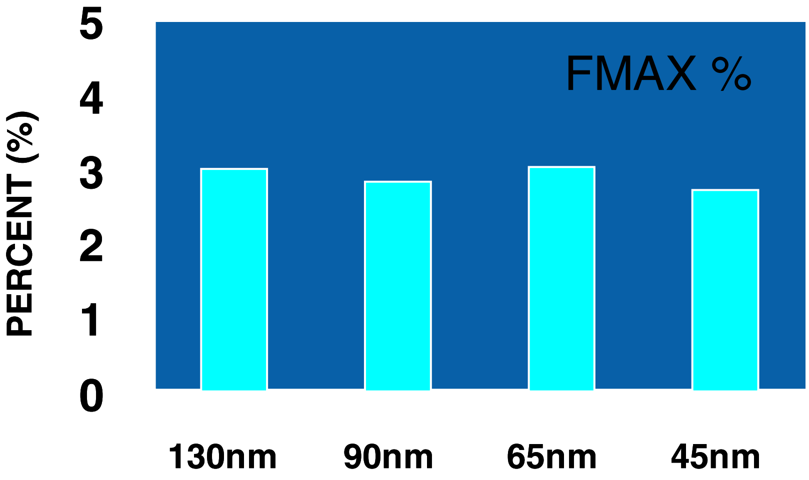

Figure 9: Within wafer variation of oscillator frequency for the 130nm through 45nm technology generations

Variation

One of the requirements for poly patterning is scaling of variation. Figure 9 shows the within wafer variation of oscillator frequency for the 130nm through 45nm generations. The variation in frequency has remained at less than 3%. This frequency variation includes the affects of poly CD, VT, and other sources of device variation. Variation has scaled with the decreasing poly pitch, despite the changes to new transistors, changes in poly patterning, and device sizes.

Yield Learning

The final measure of success of DFM is process yield. Our yield learning trend history is shown in Figure 10. The yield learning rate on 45nm technology is as fast as previous processes and is trending toward lower defect density than the 65nm process. There is no way to measure how the DFM rules contributed to this learning. Design rules, process definition, process tools, and many other things affect the yield learning rate and final high-volume manufacturing yields. All of these things are necessary to meet the manufacturing goals.

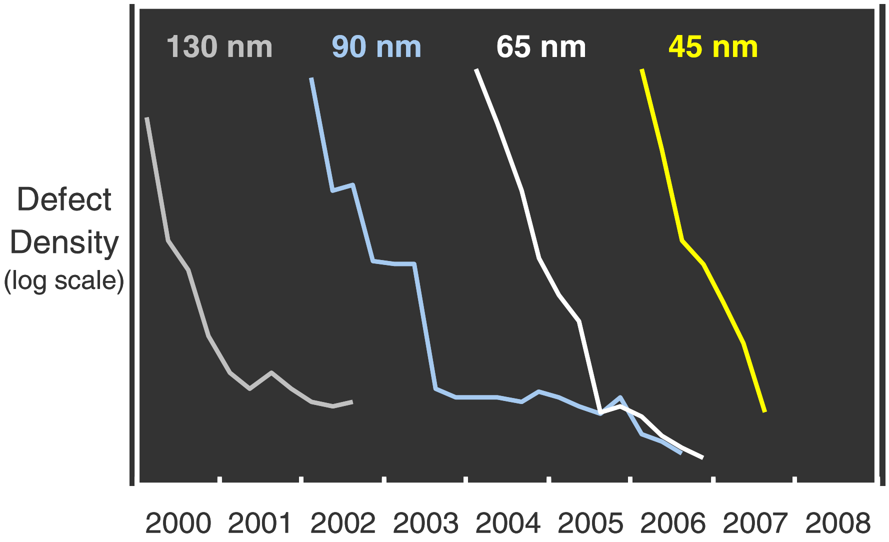

Figure 10: Yield trend

SUMMARY

Our 45nm technology maintains Moore's Law on a technology with a new transistor technology and dry 193nm lithography. Transistor density is on the 2x trend line and variation scaled. Poly rules were changed to allow simple one-pitch poly patterning with no impact to transistor density or product performance. Extensive modeling of the rules provided design rules that were stable before first product tape-out and robust enough to allow a fast high-volume technology ramp.

REFERENCES

[1] G.E. Moore, "Lithography and the Future of Moore's Law." In *Proceedings SPIE, vol. 2440*, 1995, pp. 2–17.

[2] Borodovsky, Yan, "Marching to the Beat of Moore's Law." In *Proceedings SPIE*, *vol. 6153, Advances in Resist Technology and Processing XXIII,* March 2006.

[3] Alfred K. Wong, "Microlithography: Trends, Challenges, Solutions, and their Impact on Design." *IEEE Micro*, *vol. 23, no. 2*, March/April 2003, pp. 12–21.

[4] Shekhar Borkar, "Parameter Variations and Impact on Circuits and Microarchitecture." *40th Design Automation Conference, 338,* (*DAC'03*), 2003.

[5] Kuhn, Kelin J., "Reducing Variation in Advanced Logic Technologies: Approaches to Process and Design for Manufacturability of Nanoscale Cmos." *IEDM* 2007, pp. 471–474.

[6] Mistry, Kaizad A, "45nm Logic Technology with High-k+Metal Gate Transistors, Strained Silicon, 9 Cu Interconnect Layers, 193nm Dry Patterning, and 100% Pb-free Packaging." *IEDM* 2007, pp. 247–250.

ACKNOWLEDGMENTS

Barbara Chappell, Sam Sivakumar, Charles Wallace, Mark Bork, and Robert Bigwood.

AUTHOR'S BIOGRAPHY

Clair Webb is an Intel Fellow. He has been with Intel for 31 years. He has worked on DRAM, SRAM, and cache designs. For the last few process generations he has worked on design rule definition and design issues for new process technologies. His e-mail is clair.webb at intel.com.

BunnyPeople, Celeron, Celeron Inside, Centrino, Centrino logo, Core Inside, FlashFile, i960, InstantIP, Intel, Intel logo, Intel386, Intel486, Intel740, IntelDX2, IntelDX4, IntelSX2, Intel Core, Intel Inside, Intel Inside logo, Intel. Leap ahead., Intel. Leap ahead. logo, Intel NetBurst, Intel NetMerge, Intel NetStructure, Intel SingleDriver, Intel SpeedStep, Intel StrataFlash, Intel Viiv, Intel vPro, Intel XScale, IPLink, Itanium, Itanium Inside, MCS, MMX, Oplus, OverDrive, PDCharm, Pentium, Pentium Inside, skoool, Sound Mark, The Journey Inside, VTune, Xeon, and Xeon Inside are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Intel's trademarks may be used publicly with permission only from Intel. Fair use of Intel's trademarks in advertising and promotion of Intel products requires proper acknowledgement.

*Other names and brands may be claimed as the property of others.

Microsoft, Windows, and the Windows logo are trademarks, or registered trademarks of Microsoft Corporation in the United States and/or other countries.

Bluetooth is a trademark owned by its proprietor and used by Intel Corporation under license.

Intel Corporation uses the Palm OS® Ready mark under license from Palm, Inc.

LEED - Leadership in Energy & Environmental Design (LEED®)

Copyright © 2008 Intel Corporation. All rights reserved.

This publication was downloaded from
http://www.intel.com.

Additional legal notices at:
http://www.intel.com/sites/corporate/tradmarx.htm.

For further information visit:

developer.intel.com/technology/itj/index.htm