Design for Manufacturability with Regular Fabrics in Digital Integrated Circuits

By

Mehdi Gazor

B.S., Electrical Engineering and Computer Science

University of California at Berkeley, 2003

Submitted to the Department of Electrical Engineering and Computer Science in partial fulfillment of the requirements for the degree of

Masters of Science in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

[June 2005]

May 11, 2005

Author ___

Electrical Engineering and Computer Science
May 11, 2005

Certified by __

Duane S. Boning
Professor of Electrical Engineering and Computer Science
Thesis Supervisor

Accepted by _____

Arthur C. Smith
Chairman, Department Committee on Graduate Studies
Electrical Engineering and Computer Science

Design for Manufacturability with Regular Fabrics in Digital Integrated Circuits

By

Mehdi Gazor

Submitted to the Department of Electrical Engineering and Computer Science on May 11, 2005 in Partial Fulfillment of the Requirements for the Degree of Master of Science in Electrical Engineering and Computer Science

Abstract

Integrated circuit design is limited by manufacturability. As devices scale down, sensitivity to process variation increases dramatically, making design for manufacturability a critical concern. Designers must identify the designs that generate the least systematic process variation, e.g., from pattern dependent effects, but must also build circuits that are robust to the remaining process or environmental random variations. This research addresses both ideas, by examining integrated circuit design styles and aspects that can help curb process variation and improve manufacturability and performance in future technology generations.

One suggested method to reduce variation sensitivity in system designs has been the concept of design regularity. Long used in FPGAs, and SRAMs, the concept of repeatable blocks is examined in this work as a method of reducing circuit variation. Layout based variation is examined in three designs with different distinctions of regularity: a Via-Patterned Gate Array (VPGA) FPU, a Berkeley BEE-generated decoder, and a low power FPGA. The circuit level impact on variation is also considered, by examining several circuit architectures. This includes analysis of the novel Limited Switch Dynamic Logic (LSDL) style, which reduces design area and encourages regularity through minimum logic sizing. Robustness to spatial variation and slanted plane effects is examined with a common-centroid based layout methodology for digital integrated circuits. Finally, a methodology is introduced in the form of the Monte Carlo Variation Analysis Engine whereby distributed process variables are fed into repeated simulation runs, output metrics are recorded, and regressions are measured to expose design sensitivities. The results for different layout and circuit design styles identify improvements that may be made to improve robustness to variation. We show that design regularity is a significant factor in mitigating sensitivity to process variation and is worthy of further examination.

Thesis Supervisor: Duane S. Boning
Title: Professor of Electrical Engineering and Computer Science

Acknowledgements

I would not be where I am today without the love and hard work of my parents, so first and foremost I wish to thank them for all I have today. The same can be said for my sister, grandparents, and extended family, all of whom have given me tremendous support and encouragement throughout the years.

The past two years at MIT have been tremendously rewarding for me, both academically and extracurricularly. I have been fortunate enough to learn from the best engineers and scientists in the world and will forever look upon the experience as truly enlightening. I would like to thank my research advisor, Duane Boning, for being the most humble, flexible, and supportive advisor I possibly could have asked for. He entrusted me with the creative freedom early on to explore research ideas, for which I am grateful. These years at MIT would not have been the same without my close friends, Nir Matalon and Kenny Kamrin, or my wonderful office mates over my time, Karen Mercedes González Valentín Gettings, Tyrone Hill, Kwaku Abrokwah, Hayden Taylor, Hong Cai, Shawn Staker, Xiaolin Xie, Karthik Balakrishnan, and Daihyun Lim.

On the research end, I would like to thank the DARPA/MARCO Center for Circuits, Systems, and Solutions (C2S2) for sponsoring my research. I would like to thank Larry Pileggi and Veerbhan Kheterpal of CMU for their early guidance for my research with their VPGA technology. I would like to thank Sani Nassif of IBM Austin Research Lab, for introducing me to LSDL and giving me the opportunity to present my research at ARL. I would like to thank Brian Richards and Bora Nikolic of UC Berkeley, for giving me access to their BEE-generated decoder design for preliminary analysis. I would like to thank Hong Cai right here in the Boning group, for simulating CMP, ECD,

and other profiles with his models for the designs I researched. I would like to thank Frank Honore (MIT), Anantha Chandrakasan (MIT), and Peter Holloway (National Semiconductor), for granting me use of Frank's low-power FPGA design for layout analysis. Finally, I would also like to thank Aaron Gower-Hall and the entire Praesagus team for creating user friendly extraction tools and giving our group the privilege to use them. These models have helped me tremendously for my research.

Table of Contents

Table of Figures

Chapter 1

Introduction and Background

Basic trends in integrated circuit manufacturing are discussed in this chapter, including the growing awareness of the impact of process variation. Models of variation that will be used throughout this work are introduced, as are suggested methods currently employed to help reduce some of the variation sources.

1.1 The Current Status of Integrated Circuits

As of this writing, integrated circuits are being manufactured in 90 nm technology and microprocessors can reach speeds of over 4 GHz. At these feature sizes and performance levels, process variation is a critical design consideration. Unfortunately, the process variation problem will only get much worse, given the predicted scaling that will occur in designs [21]. As technologies continue to scale, designers must consider constraining die size growth while continuing to scale supply and threshold voltage to overcome the predicted limiters in future technologies, namely power delivery and dissipation. Domino logic will continue to lose its performance advantage over static logic, subthreshold current will increase, and designs will become more susceptible to soft errors as supply and threshold voltages scale [12].

Models for 50 nm production of digital circuits have predicted that almost an entire generation of performance gain can be lost due to systematic within-die fluctiations

[23]. Therefore, designers are at a crossroads in time where design tolerances must tighten against increasing process variations to sustain Moore's Law.

1.2 Sources and Models of Process Variation

Process variations have, to a great extent, dictated the style and progression of design in integrated circuits. Designers have attempted more robust designs and CAD engineers have developed more accurate methodologies to ensure that design performance after manufacturing matches that after simulation. Understanding the sources of process variations is critical to developing better design rules for circuits, ensuring accurate tests for robustness, and controlling manufacturing conditions for optimal design yield.

Process variations in integrated circuits can generally be classified into environmental and physical components. Environmental variations, or deviations in operating conditions arising during circuit execution, depend primarily on architectural and operating decisions, such as power grid design and component placement. Variations in power supply, switching activity, temperature, and ambient noise are all examples of environmental variations that can impact a design [20].

Physical variation, however, is also of great concern in integrated circuit design. Figure 1.1 illustrates the classifications of physical variations in design, classifying sources into device and interconnect branches and also into geometry, material, and electrical categories. This figure reminds us that variation is introduced at every stage in the build process of an integrated circuit. For instance, variation may affect layout geometries and induce unexpected coupling capacitances, may continue through manufacturing in the form of excess deposited oxide, lithography deviations and reduced

etch, and then finally through inductance from packaging materials. The possibilities are infinite, as one can imagine. Because the manufacturing process may never be deterministic, a designer's only hope is to curb the known sources of variation through design techniques and methodologies that are known to reduce process variation.

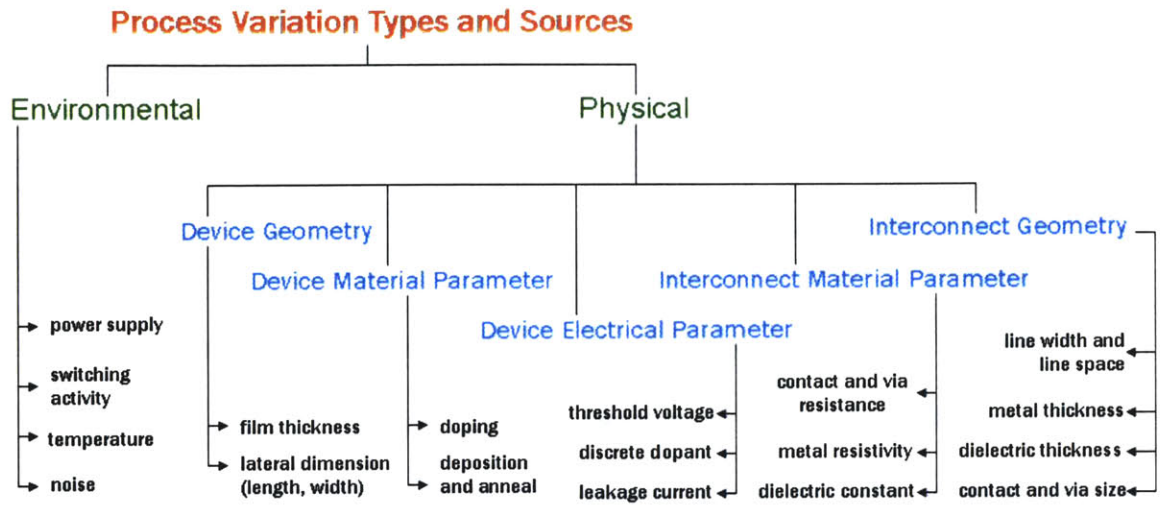

Fig. 1.1. Process variation types and sources [20].

Lumped statistical models, used in this work, are the most basic approach to describing variation. Given any process variable P, the lumped statistical model assumes that P is the sum of some nominal value P_0 and a variation measure ΔP. The variation component is assumed to be a zero mean Gaussian random variable with variance σ_P^2 [20].

$$P = P_0 + \Delta P \tag{1.1}$$

$$\Delta P \sim N(0,\sigma_P^2) \tag{1.2}$$

More accurately, we can characterize the variation components of a process variable into a greater number of more narrowly defined variables. Interdie and intradie

variation are two such ways to classify physical variation sources, as considered in Equation 1.3. Interdie variation is the difference in the value of some parameter across nominally identical die and is typically accounted for in circuit design as a shift in the mean of some parameter value equally across all devices or structures on any one chip. An interdie trend may be oxide thickness, for which variations across an entire wafer may be significant. Intradie variation is the deviation occurring spatially within any one die [20]. Mismatch variation, for example, is something that affects local dies.

$$P = P_0 + P_{interdie} + P_{intradie} + \Delta P \tag{1.3}$$

Because of the small area of the die, intradie variation for any process variable is modeled as a linear function of position, with $\omega_x x$ and $\omega_y y$ as the said components. Interdie variation, however, is more complicated, as many contributions many factor into such a variable. Equations 1.4 and 1.5 detail these components. Interdie variation is treated as a zero mean Gaussian random variable with variance $\sigma_{interdie}^2$ [20].

$$P_{intradie}(x,y) = W(\omega, x, y) = \omega_0 + \omega_x x + \omega_y y \tag{1.4}$$

$$P_{interdie} = P_{fab\text{-}to\text{-}fab} + P_{lot\text{-}to\text{-}lot}(\text{fab}) + P_{wafer\text{-}to\text{-}wafer}(\text{lot}) + P_{die\text{-}to\text{-}die}(\text{wafer}) \tag{1.5}$$

$$P_{interdie} \sim N(0,\sigma_{interdie}^2) \tag{1.6}$$

1.3 Design for Manufacturability

Through years of research and production, a few key rules have developed in the semiconductor industry, which make up the concept of Design for Manufacturability (DFM). This idea encourages providing better process/circuit design prior to manufacturing, to allow for faster yield learning, and stabilizing the manufacturing by

minimizing process variability. DFM can be accomplished through a combination of minimizing process variations, and minimizing device and circuit sensitivities to process variations. This can be accelerated once the process is statistically characterized, exposing inefficiencies [9]. Many of the rules advocated in the DFM method include measures to mitigate pattern dependent variation, such as reduction of geometric primitives, layout regularity, redundancy techniques, pattern density uniformity, and optimal process feature sizing. Other areas of consideration in DFM are design for robustness and design for test. These suggested approaches are detailed below and considered in this research.

Reduction of Geometric Primitives

A geometric primitive refers to the category of shape in a design. The rectangle is the most basic primitive in circuit layout. With each new geometric primitive comes a new set of uncertainty in processing. Primitives process differently, whether in lithography, etch, plating, or through other steps in the process cycle. Complicated Resolution Enhancement Techniques (RET), depicted in Figure 1.2, must be applied to compensate for the distortions that new primitives may introduce. These techniques are costly and often insufficient for arbitrary random layout patterns [5]. Reducing the number of these primitive shapes in a design ensures higher certainty during the manufacturing stages, and is an important component of design for manufacturability.

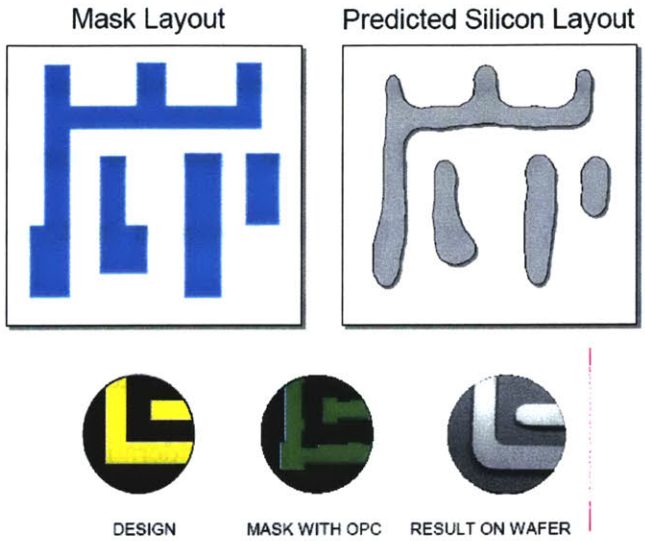

Fig. 1.2. Increasing geometrical primitives increases uncertainty in final design [1].

Layout Regularity

Design performance has been correlated with the methodology and technique used during transistor layout [3]. Figure 1.3 illustrates some of the layout techniques that may affect performance in a design. Number of fingers per transistor, transistor orientation, proximity spacing, local and global effect of polysilicon density, interconnect orientation, choice of metal layer, and coupling capacitance across metal layers are all factors that can be constantly monitored and corrected during the layout process. Consistency is an important factor in this stage of the design process, as is an understanding of the manufacturing impact of layout styles.

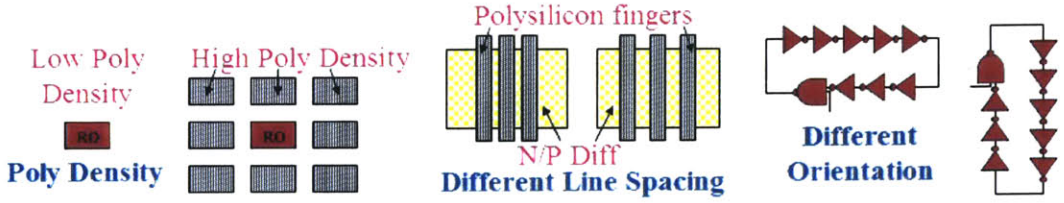

Fig. 1.3. Some layout-induced variation sources [3].

Redundancy to Improve Manufacturability

To improve planarity of polysilicon or metal layers in a design, metal "dummy" fill is inserted throughout each layer. Dummy fill consists of redundant pieces of metal placed in empty areas of a design with automated CAD tools to increase planarity. At the cost of increased capacitance, this fill narrows the layer's range of pattern density and significantly reduces CMP variation. Inserting dummy cells and dummy transistors are also methods for reducing edge effects for transistors, memories, and other blocks in a design. Finally, redundant via placement may improve manufacturability by decreasing the likelihood that a single via causes logical failure in a design due to bad processing. Combined, these techniques are a wise combination to reduce the variation effects of manufacturing and improve overall yield in an IC process.

Pattern Density Uniformity

One of the more understood culprits of systematic process variability is the irregularity in surface topography due to an unevenly distributed pattern density [2]. In Figure 1.4.1, the high pattern density region on the left results in a plating bulge, whereas the low density region to the right creates plating recesses. The electroplating-induced topography variation interacts with additional pattern effects in chemical-mechanical polishing (CMP), as seen in Figure 1.4.2.

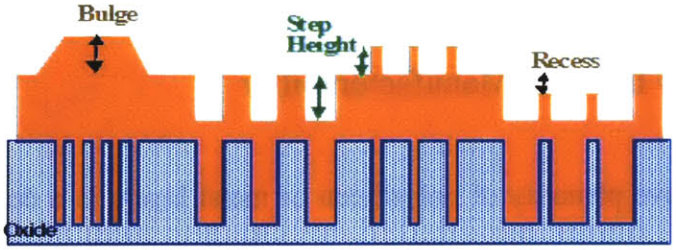

Fig. 1.4.1. Copper plating problems due to irregular pattern density in design [2].

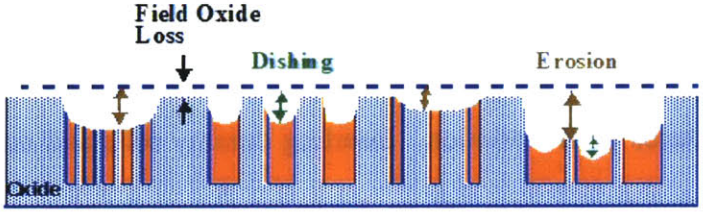

Fig. 1.4.2. CMP dishing/erosion problems due to irregular pattern density in design [2].

Fluctuations in local metal thickness can affect the resistance and capacitance of the region, adding unexpected loads to wires, for example. Clock skew resulting from mismatch in the affected interconnect is one important effect of non-uniform pattern density. Electromigration also becomes an increased concern with sub-optimal non-regular metal thicknesses [2]. In the worst case, logical failure ensues from the improper plating and polish of the design.

Fortunately, dummy fill algorithms have the ability to compensate for the non-uniformity of most designs. Though dummy fill can alleviate macro-level fluctuations in pattern density uniformity, it is not a complete solution. Dummy fill introduces significant capacitance to a design, which may be a high cost for high performance circuits. Therefore, a designer must still consider the most regular approach to design and layout with consideration of density in mind.

Optimal Feature Sizing

Local feature size records the smallest width of features in a small area on a chip. Various sizes and interaction distances among local wires, transistors, and other design components create this feature size variation. The resulting non-uniformity may affect pattern density and may result in copper dishing or oxide erosion variation [2], as illustrated in Figure 1.4.2. Optimally, all areas of a design have uniform feature sizes, encouraging consistent processing across the chip. Well characterized feature sizes could remove some of the uncertainty present in the manufacturing process, and are therefore advocated in the DFM methodology.

Performance Robustness

A careful designer will examine the timing slack across his circuit, whether for clock distribution or data arrival delays. More importantly, he or she will design with tolerances for error in mind. These include considerations for process variations, such as transistor mismatch and die-to-die variation, as well as for environment variables, such as power supply deviations. Design for robustness integrates these considerations with well constructed tests for performance under the worst process circumstances. The most robust designs, therefore, are the ones that can operate in light of these adverse conditions.

Design for Testing

Chip design complexity and run costs have forced designers to consider testability of design before fabrication. Design for test is now a critical part of the IC design

process, with much development going into the areas of test structures, scan chains for test, and on-chip measurement techniques [3]. Design for manufacturability intersects with design for test, together ensuring robustness to process errors and a clear understanding of these errors when they do occur.

1.4 Summary and Directions

Process variation has become an increasing concern in integrated circuit manufacturing. The many sources and classifications of variation were described, and a lumped statistical model for quantifying variation in design was introduced. Design for manufacturability (DFM) was presented as a school of thought for curbing integrated circuit process variation during the design stage. Many of the considerations advocated in DFM were introduced and described, with much of the focus on methods to reduce pattern dependent variation.

Going forward, there are two key variation reduction ideas throughout this work. One is that some designs and design styles *generate* less process variation, i.e. more regular layout actually creates less pattern dependent variation. The second idea is that different design styles are more or less robust to the remaining process or environmental random variations, whether systematic or random. This research addresses both ideas, by examining integrated circuit design styles and aspects that can help curb process variation and improve manufacturability and performance in future generations.

The ideas and methods described in this chapter will be used to examine the impact of regularity on reducing pattern dependent variation in digital integrated circuits. In Chapter 2, preliminary comparisons will be done on logical fabrics that advocate the

use of regularity against more traditional non-regular designs. Using the variation models presented in this chapter, this work develops and analyzes an assortment of tests for regular and non-regular circuit architectures in Chapters 3 through 5. A Monte Carlo methodology and sensitivity analysis are later introduced in Chapter 6 as methods to measure and compare design vulnerability to variation. Finally, a novel way to measure and explore the impact of regularity on variation robustness based on layout placement is introduced and evaluated in Chapter 7.

Chapter 2

Manufacturability of Regular Fabrics and Flows

Several of the considerations introduced in the previous chapter for design for manufacturability can be addressed with the notion of regularity in design. This chapter explores the potential of using regular fabrics as a method to reduce pattern dependent variation in designs. Three designs are analyzed and compared. The regular via-patterned gate array (VPGA) is compared against more traditional ASIC and regular FPGA designs to evaluate manufacturability potential. Each of these designs with and without dummy fill is also evaluated to determine the effectiveness of dummy fill in reducing edge effects and pattern dependent variation.

2.1 Motivation for Regularity

As designs become smaller and more intricate, the effects of process variation become more apparent and intolerable. These variations are either random or systematic, the latter of which occur in a reproducible fashion. A single systematic process variable may be the aggregate sum of numerous individual deterministic contributions, as visualized in Figure 2.1. Our goal is to identify these systematic variations and thus adopt layout and design approaches that both minimize these variations in the first place, and are robust to the remaining systematic or random variations.

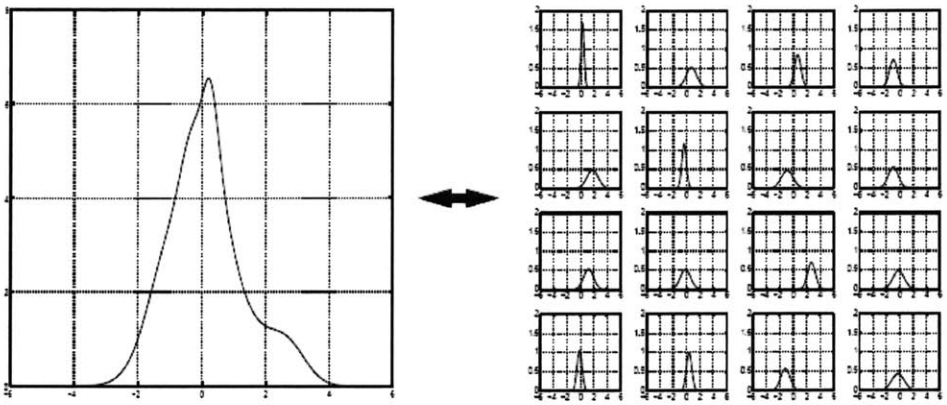

Fig. 2.1. A large number of systematic or deterministic contributions (right) to a parameter can appear in aggregate as a single "random" distribution (left) [7].

To eliminate systematic variation would reduce much of the process variation that tends to occur during fabrication. Though there is no way to completely eliminate systematic variation, one proven way to reduce this type of variation has been through well characterized devices. Once the manufacturability of these small structures is understood and optimized, there is less likely to be systematic variation involved during the fabrication process. It then follows that larger structures built of these smaller units will also exhibit the manufacturability advantages of its smaller units, reducing the intradie variation component visualized in Figure 2.2. This is the core idea behind regularity and regular fabrics [5].

Regular fabrics are a general term for the digital architectures that exhibit these modular, repeatable logical components. Traditional examples include the field programmable gate array (FPGA) which uses repeated configurable logic blocks (CLBs) to create a powerful grid of programmable switches.

Regular structures have the manufacturability advantage over more custom ASICs due to their well-characterized, repeated components. In spite of a smaller die area,

custom designs tend to have lower yields and fewer working dies per wafer [4]. Once smaller pieces are modeled and well understood in terms of the effect of process variation, lithography, and etch, then the regular design that is composed of these smaller logic blocks can be well understood. This advantage has tremendous cost and yield implications, since reliability and robustness can improve over non-regular designs.

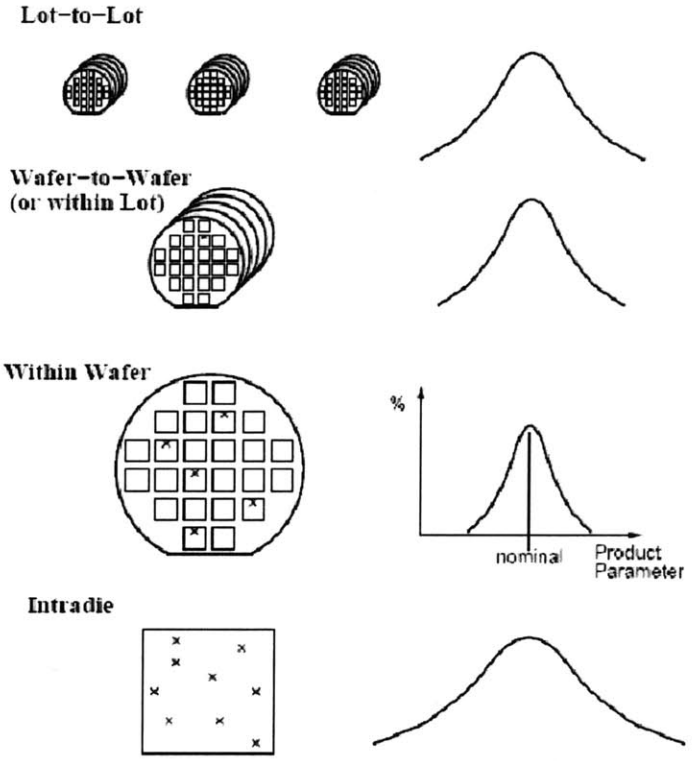

Fig. 2.2. Scope of variation in semiconductor manufacturing [6].

Feature level variation is a large concern in pattern sensitive processes such as lithography, plating, etch, and deposition. Layout printability challenges have become extremely severe due to: 1) high NA (numerical aperture), off-axis illumination schemes (angular, quadrupole, dipole) and small depth of focus; and 2) large mask error enhancement factor (MEEF). Moreover, etch and chemical mechanical polishing (CMP), which depend on intra-layer layout density variation, add to the challenges [5]. The

difficulty of lithography is compounded when the design is composed of an assortment of geometric primitives, orientations, and interaction distances, as suggested in Figure 2.3. It is increasingly difficult to print and pattern complex figures due to larger interaction distances. Resolution enhancement techniques (RETs), which predicatively correct for errors in lithography, are costly and insufficient for arbitrary layout patterns [5].

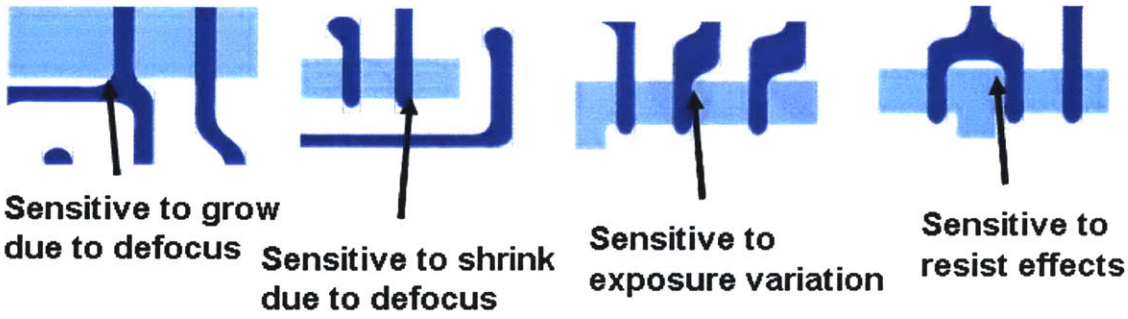

Fig. 2.3. Shortcomings in optical lithography. Source: A. Strojwas & W. Maly, CMU

To reduce the problem complexity and achieve the desired performance and yield objectives, new solutions must explore more deeply the concept of regularity. The ultimate solution would be based on a full chip layout being assembled out of a set of patterns that are guaranteed to print for given lithography, etch and CMP process windows [5]. A regular fabric would ideally have a single logic block optimized for these mentioned properties, reducing the uncertainty in lithography or later stages of fabrication. Regularity at the feature level and device level will be a great step in reducing spatial variation, thereby reducing systematic variation at the intra-die level.

Regular structures are critical in future scaled designs. W. Maly concludes that only by applying in a design highly geometrically regular structures, created out of the limited smallest possible number of unique geometrical patterns, can one hope to contain

the cost of nanometer ICs on the manageable level [11]. A regular architecture that employs many of the advocated regular structures is described and evaluated in the next section.

2.2 Via-Patterned Gate Arrays: A Regular Fabric

The via-patterned gate array (VPGA) is a new digital architecture that takes traits from both custom ASIC and fully programmable architectures, like FPGAs [22]. Developed at Carnegie Mellon University, the basic idea behind these structures is to use repeatable logic units in a grid-like fashion, like FPGAs, to construct a larger operational unit. The first metal layers include the regular logic structures and the higher metal layers are composed of regular interconnect grids. A VPGA is "programmed" one time during fabrication through strategic placement of vias across the interconnect layer, to customize the logical operation of the design. In this respect, via-patterned gate arrays can be considered a regular, yet semi-custom ASIC. Its regular logic structure and the restriction of via placement to only a few metal layers not only decreases mask cost, but also has the potential to increase manufacturability tremendously.

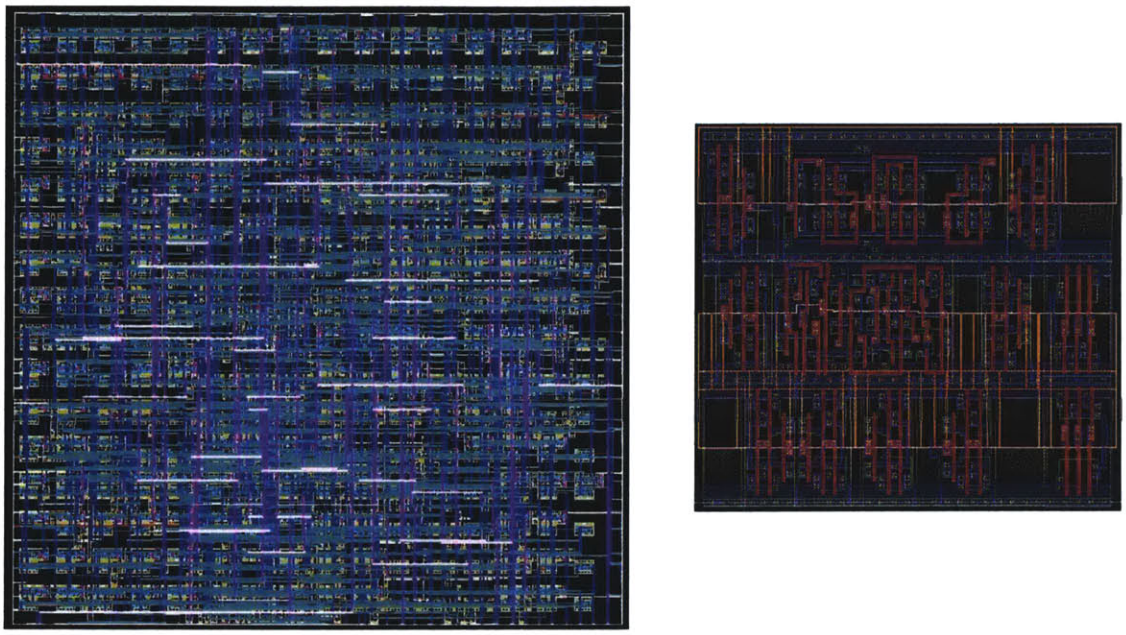

Fig. 2.4. Via-patterned gate array design (left) and single logic block component (right). Source: L. Pileggi.

VPGA cells contain simpler elements such as nands, nots, muxes, and pass gates and are highly optimized to be robust to process variations. Figure 2.4 depicts both a VPGA design and an individual cell component, which consists of simple logic elements. Used together, the cells form a regular fabric that is also tuned for successful fabrication. One major metric of a manufacturable design is the composition of its pattern density. In Figure 2.5, the range of interconnect pattern density across three metal layers can be visualized. In each individual layer, the range is no more than 5%, a number substantially lower than that of traditional ASIC layouts. This uniformity in interconnect density will ensure that many systematic variations during fabrication will be kept close to minimal, improving manufacturability.

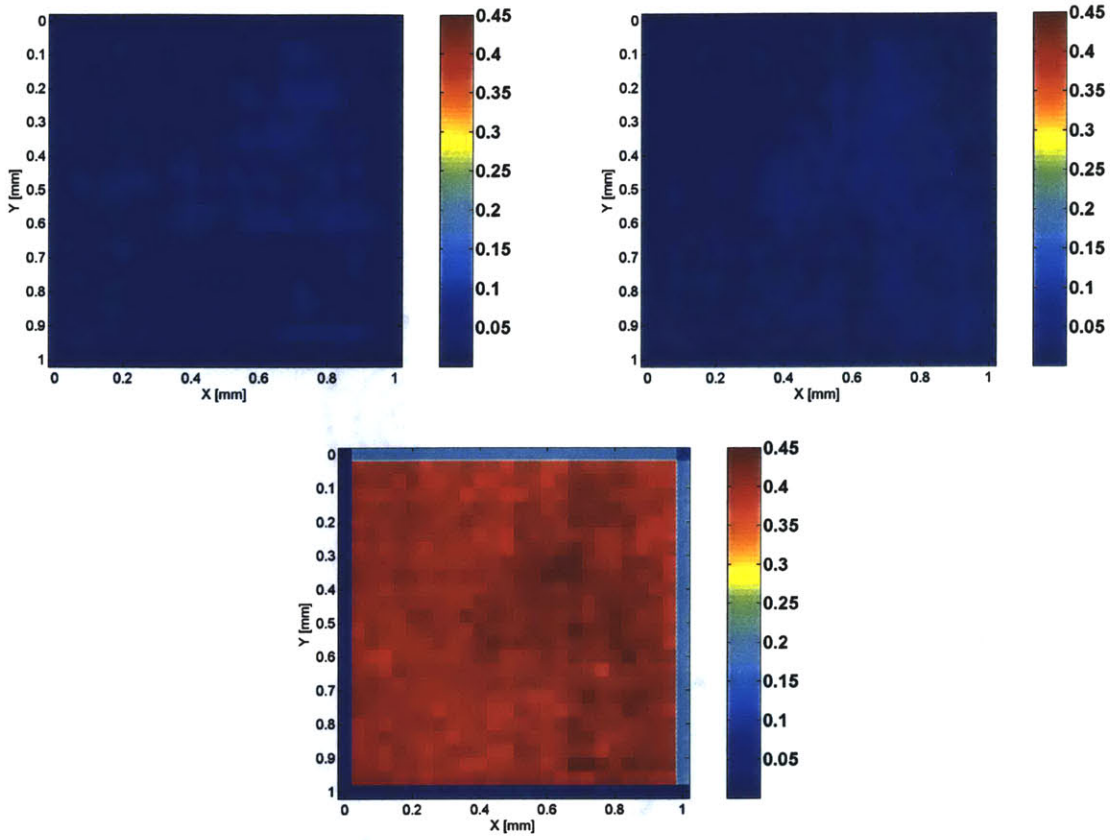

Fig. 2.5. Normalized VPGA Metal Densities. Metal 1 (upper left); Metal 2 (upper right); Metal 3 (lower).

Another useful metric of manufacturability, demonstrated in Figure 2.6 for Metal 1 of the VPGA layout, is uniformity in feature sizes across a layout layer. One way to characterize this is to put feature sizes in a number of "bins," each bin holding a defined range of feature sizes. These bins can be extracted for each discretized region on a layout. By examining the binning plots of Figure 2.6, the relative percentages of feature sizes that fall within a certain bin can be seen. The narrower the range of the feature sizes are across a design, the less vulnerable such a design is to many manufacturing uncertainties, including copper dishing and oxide erosion, which can lead to plating and CMP variation [2]. For the VPGA evaluated here, over 90% of all

interconnect features in the design are grouped into Bin 1, or features smaller than 0.35 μm in width. Along with the uniform pattern density across the metal interconnects, the process binning is optimal for manufacturability.

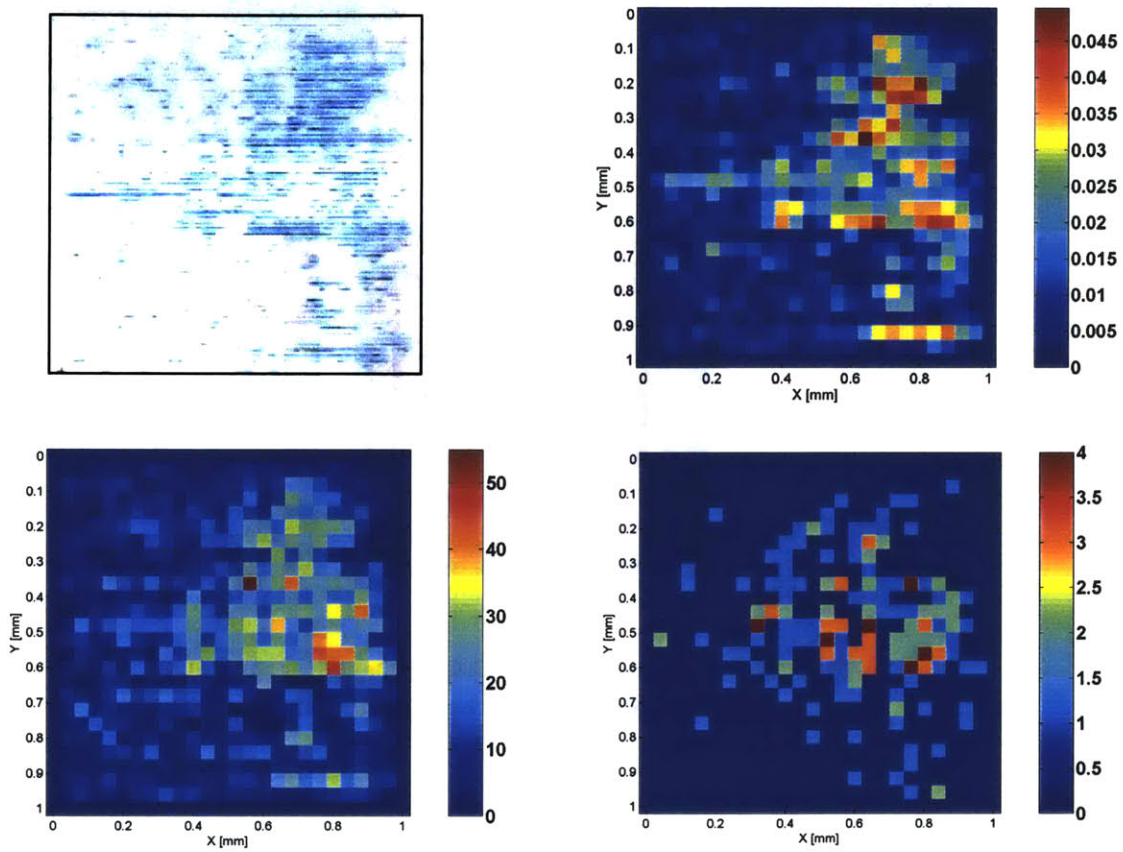

Fig. 2.6. Process binning for VPGA FPU Metal 1 interconnect. Metal 1 interconnect layout (upper left); Metal 1 interconnect pattern density (upper right); Bin 1 [0-0.35 μm] (lower left); Bin 2 [0.35-0.50 μm] (lower right).

Moreover, when actual process models are run on the VPGA design, the results are dramatic. Based on models for copper CMP developed elsewhere in other work, we can predict how each metal layer will manufacture, for an example plating and CMP process [19]. Metal 4, which is one of the grid layers for the VPGA, was used to demonstrate the manufacturability of regular, repeated structures. Copper dishing, oxide

erosion, plated thickness, and post-CMP copper thickness models were applied to the VPGA FPU layout to determine the degree of variation in processing. The results, shown in Figure 2.7, highlight the uniformity of the processing involved. Ignoring the edge effects caused by the layout edge, the degree of variation is low relative to more irregular designs: final post-CMP copper plating thickness varies less than 3%, compared with up to 12% or more in some cases for custom ASIC designs.

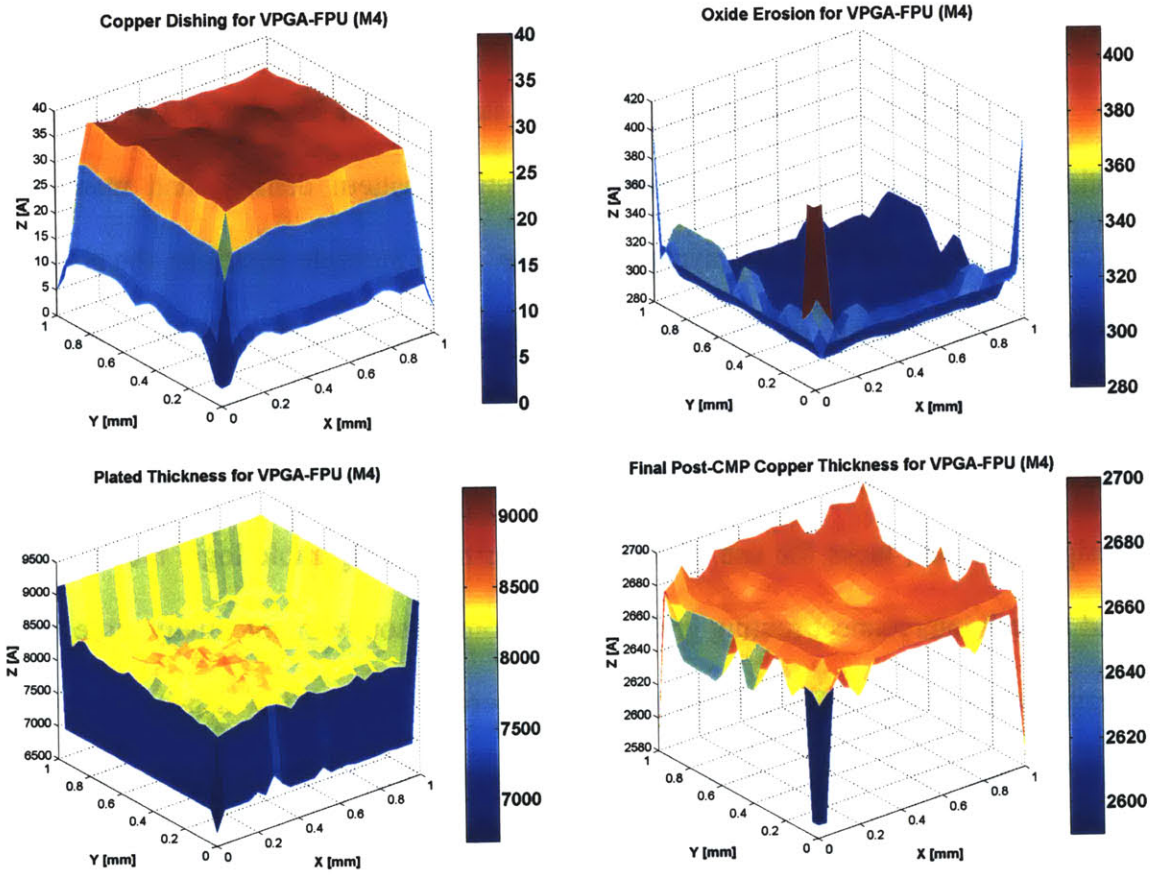

Fig. 2.7. Modeled copper thickness (upper left), oxide erosion (upper right), plated thickness (lower left), and post-CMP copper thickness (lower right) for VPGA FPU's metal 4 interconnect.

2.3 Berkeley Emulation Engine (BEE): A Semi-Regular Flow

The Berkeley Emulation Engine (BEE) project at the UC Berkeley Wireless Research Center is a system of flows and parallel FPGA processors for emulation [24]. Its goal is to speed the chip design to hardware verification process to a single day. Relatively well documented, the project uses an intelligent, semi-regular flow to optimize and layout an ASIC design. It also utilizes an efficient dummy fill algorithm to reduce process variation effects, and is thus an interesting candidate for design flow examination.

A semi-regular ASIC design flow uses standard cell libraries and algorithms to optimize interaction distance, interconnect routing, pattern density, and many other measures to create as regular a layout as possible. Over multiple iterations, the placement tools in the flow decide on the optimal placement of standard cells. With an intelligent flow, these cells are then routed to achieve reproducible, optimal results. Both the VPGA and a BEE-generated ASIC are synthesized designs with standard cell routing. The BEE design, however, lacks the use of well characterized configurable logic blocks (CLBs) and a mesh interconnect routing grid, and thus comparison with potentially more regular designs such as the VPGA are of interest. A generated 4092-bit low density parity check (LDPC) decoder [14] was chosen for comparison to the VPGA design. The metal 3 interconnect layer was extracted from the LDPC decoder and analyzed in the same fashion as its VPGA counterpart. The proprietary Praesagus extraction tool was used to extract a pattern density map from the design, as depicted in Figure 2.8. The design with and without dummy is evaluated for comparison.

From a manufacturability standpoint, a few things stand out with the BEE-generated decoder. The first is the impact of the dummy fill algorithms in the CAD flow.

As mentioned earlier in this work, dummy fill refers to the act of inserting redundant blocks in each layer of a design to create a more planar surface after patterning and CMP. In this design, for example, dummy fill reduces oxide erosion variation on the die from 25% to 10% and post-CMP die copper thickness variation from 8% to 3%, as depicted in Figures 2.9 and 2.10. However, good dummy fill algorithms are increasingly complex and time and resource intensive to apply, limiting their widespread use.

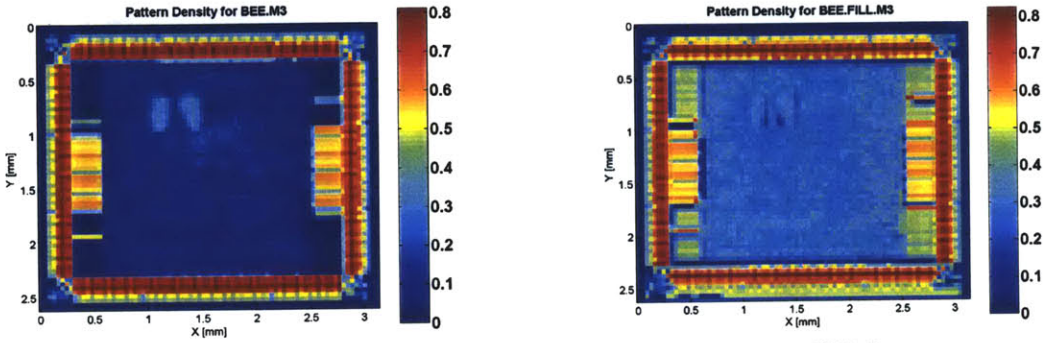

Fig. 2.8. Pattern density (with and without dummy fill) for BEE-generated layout.

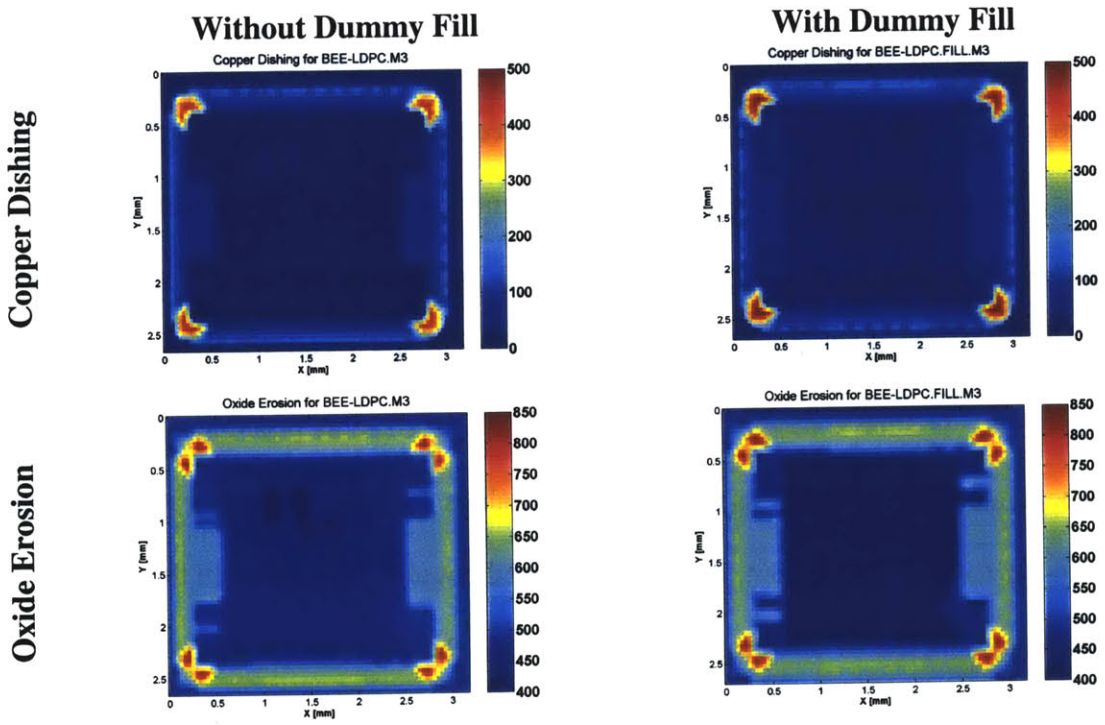

Fig. 2.9. Copper dishing and oxide erosion (with and without dummy fill) for BEE-generated layout.

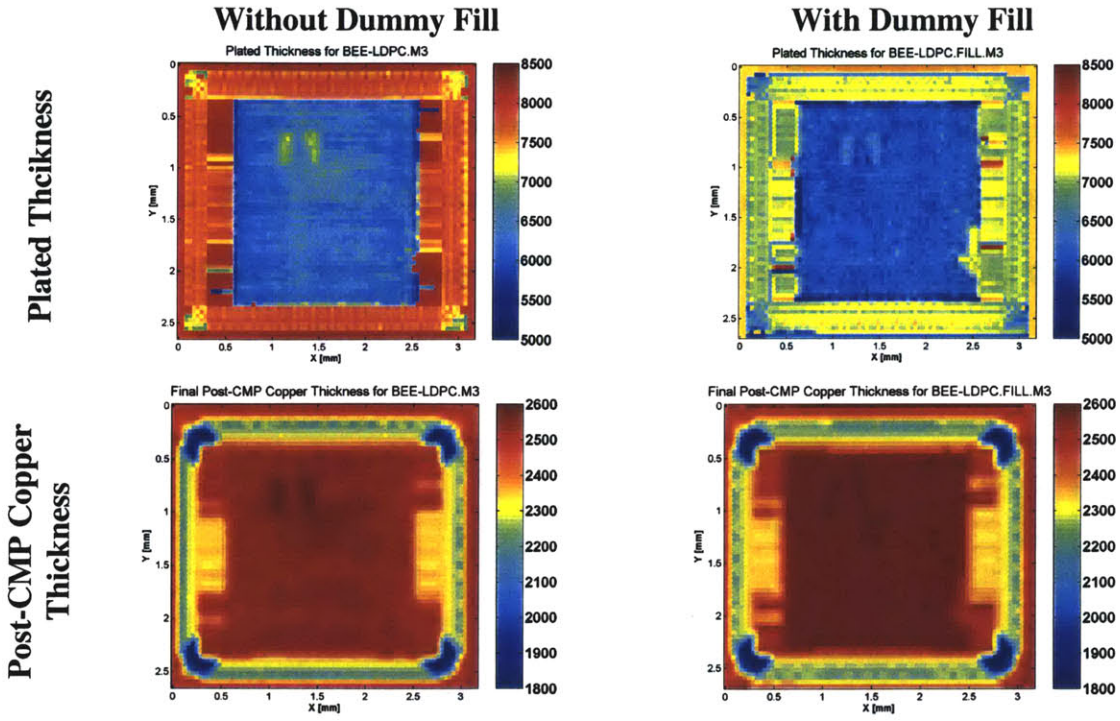

Fig. 2.10. Plated thickness and post-CMP copper thickness (with and without dummy fill) for BEE-generated layout.

The second item of interest is the post-manufacturing topography of the pads and corners of the chip, seen in Figure 2.11. The steep drop in landscape surrounding the circuitry is not corrected by the dummy fill and contributes the greatest source of variation in the design. Although the decoder was designed within its technology's DRC specifications, this example illustrates that there is much room for manufacturability optimization. In such a design (1 mm^2), as planar a surface as possible will ensure the most accurate CMP and process results. The penalties for poor processing may affect the pads in the form of throughput, adding unexpected coupling and delay in this high-speed chip design. Increased attention to dummy fill in the edge and pad boundary regions would lessen these post-CMP edge non-uniformities.

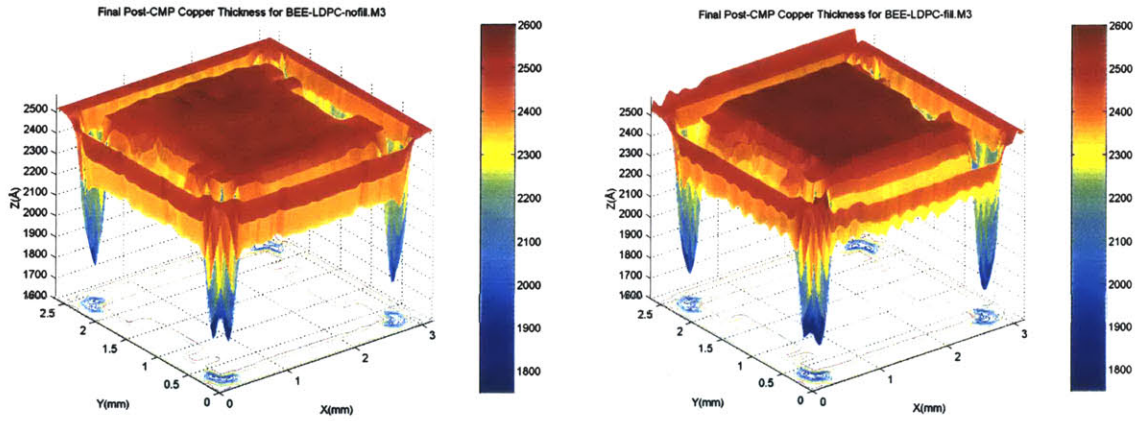

Fig. 2.11. 3D visualization of BEE-generated decoder metal 3 layer, post-CMP. Without dummy fill (left) and with (right).

The semi-regular flow of the BEE must utilize a time-intensive dummy fill algorithm process to create a layout that is comparable in manufacturability to the VPGA FPU examined. Despite the semi-regular nature of the flow and its use of standard cell libraries for routing, the post-CMP profile of the BEE-generated layout reveals that it is not as manufacturable, with respect to this class of pattern dependent process variations, as the VPGA examined previously. The regular architecture of the VPGA has intrinsic manufacturability advantages, including a grid interconnect structure and repeated blocks of standard cells. Nevertheless, it appears that a semi-regular flow, like that of the BEE, may be able to achieve comparable results with an advanced dummy fill methodology.

2.4 MIT Low Power FPGA

Up to this point, we have analyzed a novel regular fabric and a flow-generated ASIC. A last model for comparison is an established and well understood regular fabric in the form of an FPGA created by Honore and Chandrakasan [8], originally designed to apply circuit-level power reduction techniques for use in power aware systems. From the

VPGA examined before, we saw the effects of regularity in reducing system level pattern density variation. With the Berkeley decoder, we witnessed the advantage of using dummy fill for reducing process effects. In this MIT FPGA, both elements are combined for reducing manufacturing variation significantly.

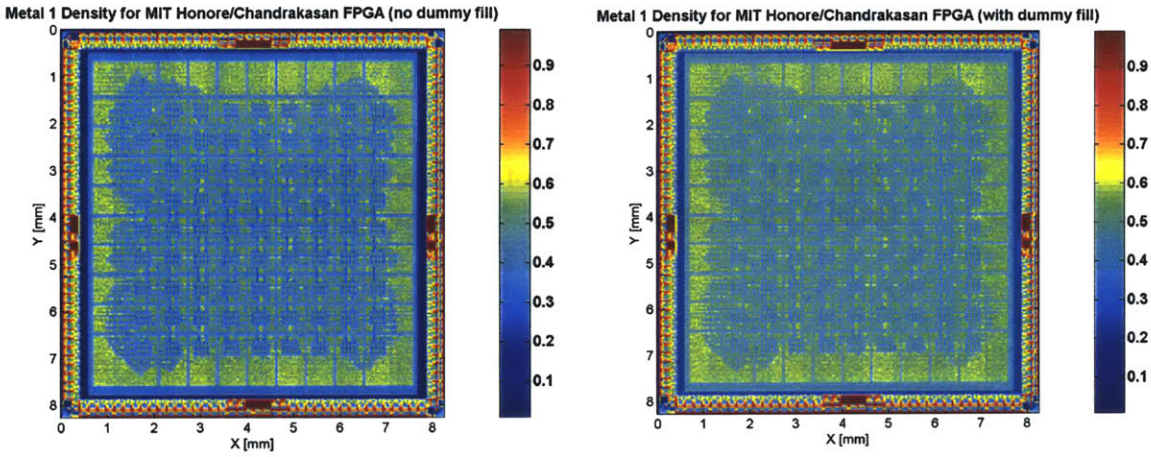

Fig. 2.12. Pattern density (with and without dummy fill) for MIT FPGA.

With this design, like the Berkeley decoder, the most dramatic effects of dummy fill are at the edges. From Figure 2.12, we notice that without dummy fill the edge of the FPGA die increases in pattern density from 0% to 38% to 60%, a poor gradient. This is in comparison to the 30% to 50% to 60% gradient when dummy fill is used. Though throughput may not be as great an issue as with the Berkeley decoder, the subject of edge non-uniformity is nevertheless an important one that demands further review.

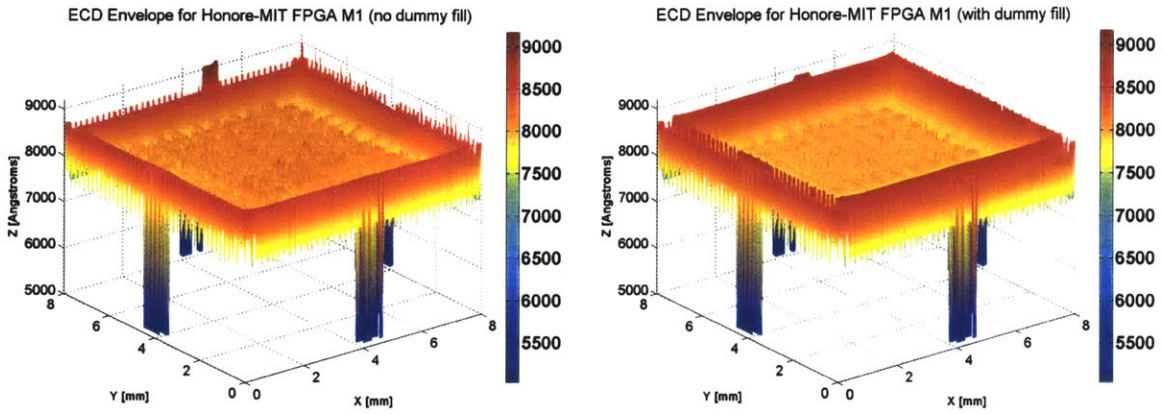

Fig. 2.13. Electrochemical deposition (ECD) Copper damascene envelope profile for MIT Honore/Chandrakasan Low Power FPGA. Without dummy fill (left) and with (right).

The dummy fill also creates a slimmer ECD envelope on the chip, and less ECD variation on the edges, as seen in Figure 2.13. After the copper CMP, the results may be dramatic, as depicted in Figure 2.14. When dummy fill is used in this regular layout, the predicted CMP envelope variation is reduced from 45% to near 5%. In terms of reducing manufacturing variation, this data suggests that the FPGA architecture contributes much less to the cause than the actual dummy fill placement. As this and previous results have indicated, design for manufacturability must address a solution beyond just the placement of regular blocks. Even an FPGA, which has traditionally been considered a highly manufacturable structure, may run into manufacturing problems related to these types of layout pattern systematic process variations.

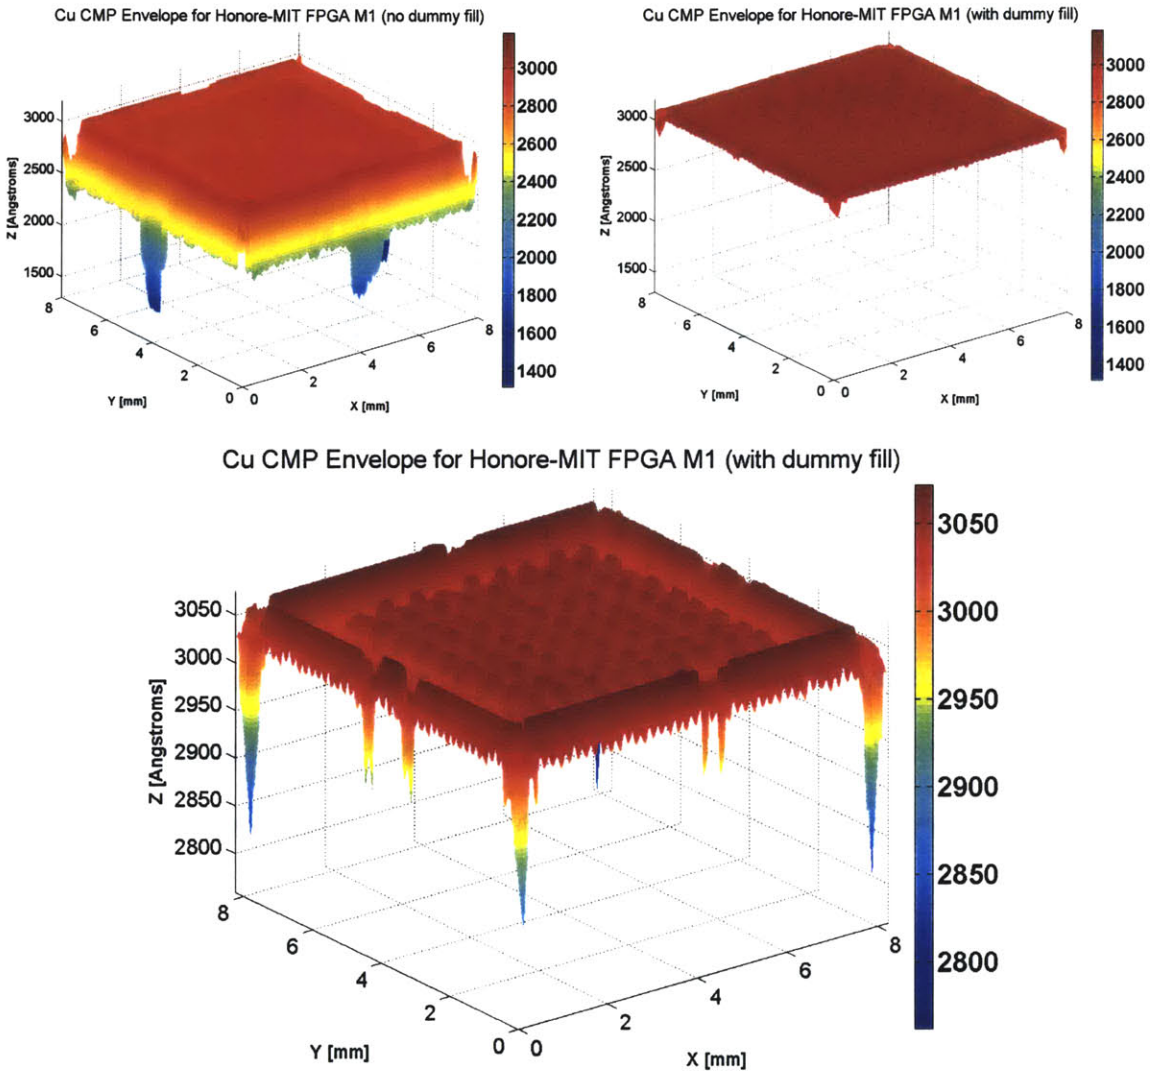

Fig. 2.14. 3D Visualization of MIT Honore/Chandrakasan low power FPGA, post-CMP. Without dummy fill (upper left) and with (upper right). Zoom of profile with dummy fill (bottom).

2.5 Results Summary

Regular fabrics in circuit designs appear promising from a manufacturability perspective. Because it is made up of well-characterized, predictable, repeated components, a regular design is less vulnerable to process variations during manufacturing. Pattern-based

dependency analyses of a via-patterned gate array (VPGA) interconnect layer confirmed the highly uniform pattern density of the structure and relative robustness to dishing and oxide erosion. Comparisons were made to a more traditional ASIC and FPGA, each with and without dummy fill, to evaluate the improvements in pattern dependency variation using such compensative layout techniques. Dummy fill does significantly planarize layers, as suggested by pattern density and post-CMP plating data, but potentially at a high capacitive cost that an intrinsically regular design, such as a VPGA, does not pay. Edge effects, and the impact of dummy fill in reducing them, were also examined.

Chapter 3

Exploring Regular Circuit Architectures

Shifting one step down from the system level to the circuit level, our next goal is to determine the potential impact of a regular circuit family in reducing systematic performance and process variation. Ultimately, isolating the robustness characteristics of each circuit family may deliver some understanding towards the impact of architecture choice on system level manufacturability. This chapter explores one such regular circuit style, that of limited switch dynamic logic, versus more traditional dynamic and static styles. The qualities that contribute to a regular circuit family are discussed and an adder design is introduced as a benchmark for further analysis in this research.

3.1 Limited Switch Dynamic Logic (LSDL) Architecture

Dynamic logic has traditionally been a faster performing alternative to static logic, though it trades power and robustness for its performance advantage [15]. Dynamic logic's clocked nature makes it more difficult to manage and introduces another dimension of variation that must be controlled for correct operation. More advanced forms of domino logic, like the dual-rail version, have been developed to solve many of the robustness problems, but power still remains a critical factor.

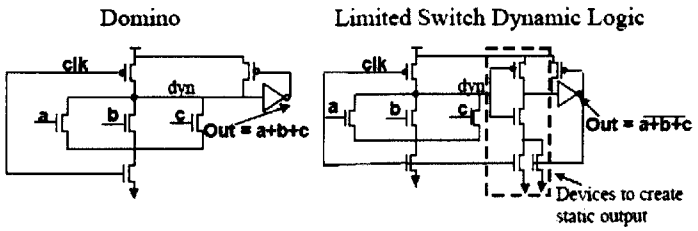

Fig. 3.1. Domino vs. LSDL schematics [10].

By introducing latches into the logic at every stage of a domino design, we can generate what IBM has termed limited switch dynamic logic, or LSDL [25]. Typical schematics for domino and LSDL logic are depicted in Figure 3.1. There are numerous advantages to such a technique. Introducing an inverting latch into the logic essentially separates the logic stage from the gain stage. As a result, logic stages can be kept close to minimum size, and only the gain stages need to be sized, as suggested in Figure 3.2. The result is a much smaller design, which has promising cost implications. In this smaller design, cells more closely resemble each other, not only in size, but also in content. All cells include the same pattern of latching transistors and a similar pattern of (close to) minimum sized domino transistors. From a performance perspective, the LSDL clock can resemble more of a pulse than a standard domino clock, given the latching capabilities of the circuit. Duty cycle can be reduced well below 50%, allowing for higher speed designs. Moreover, the latched nature of the architecture ensures less activity through switching at internal nodes in the design, which can help reduce overall power consumption.

The concept of embedding logic functionality into latches is not a novel one. It was extensively used in the design of the EV4 DEC Alpha microprocessor and many other high performance designs [15]. True Single Phase Clock (TSPC) is a popular

method of integrating logic with latches. LSDL, though structurally similar to previous techniques, varies in its approach to clocking and circuit sizing, exploiting the regularity of this circuit topology in a way not considered before.

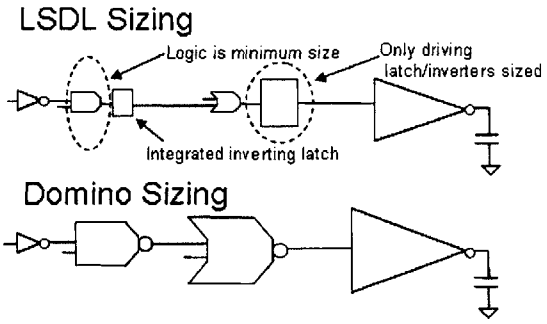

Fig. 3.2. LSDL sizing vs. domino sizing.

Though this circuit style has promising performance and power advantages over traditional static and domino designs, another interesting aspect is the regularity of such a style. There is logical regularity in the fact that there is separation of gain and logic elements and less internal switching, but there is also considerable layout regularity. Since logic is kept small, the majority of area in such a design would be taken up by the LSDL latches, which themselves are simple, well-characterizable, repeated structures.

These principles were examined and an approach was established to examine the effect of the regularity of LSDL on not only performance, but also robustness to process variation and manufacturability. A comparison was made between domino, static, and LSDL adders for the purpose of understanding the significance of this regularity better.

3.2 Domino and Static Circuits

A traditional domino or static circuit is not regular in sizing like an LSDL circuit may be, but is much more widely used. Static CMOS circuits have the primary advantages of

logical robustness (i.e. low sensitivity to noise), good performance, and low power consumption with no static power dissipation [15]. The static circuit is therefore the classic comparison to use as a baseline when measuring robustness and reliability.

Advantages of dynamic logic include fewer transistors (thus fewer load capacitances to drive at each stage), increased speed performance, and no static power consumption. Disadvantages include increased overall power consumption, and an assortment of clock issues, including slew, clock feedthrough, and timing problems. Dynamic domino logic was used in this comparison for its similarity to LSDL, to quantify the extra advantages of the LSDL design compared to static logic.

All circuit styles were sized using the logical effort technique [15]. In a chain of circuits, this method uses the fact that complex gates must work harder, or place more effort, to produce similar responses. As indicated before, only the LSDL design takes advantage of this principle, by integrating simple inverting latches, which themselves have low logical effort. Overall, since a chain sizing focuses on these inverters, overall design size is reduced. The static and domino designs do not possess this intrinsic regularity component.

3.3 Adder Design

Adders are generally well understood designs, and quite regular structures. They are designed using blocks of logic in a repeated fashion and can be scaled to a larger size using more of these blocks. For these reasons, a 16-bit carry look ahead adder was a good structure to compare different circuit architectures against one another. By examining the strengths and weaknesses of each of these architectures to process

variations, it is possible to understand the impact of design style on not only performance but manufacturability.

The 16-bit size for the adder was chosen to begin with, to be expanded to 32-bit if necessary. The simulation time per run had to be kept reasonable to complete the thousands of Monte Carlo runs planned for each adder in the limited time available. The carry-look ahead style was chosen for its speed and common use. The logarithmic nature of its computation ensures maximum speed and reflects modern adder designs. Higher radices in logarithmic adders reduce the number of stages, but make each gate more complicated. Higher radices are better for low fanout where intrinsic delay dominates. Sparse trees reduce the number of gates and wires, but increase the fanout on the internal nodes. Finally, a Kogge-Stone tree style was selected because these trees yield a minimum number of stages for a given radix [17]. With these considerations in mind, a radix-2, sparseness-2, Kogge-Stone CLA adder was selected for design. For n bits with radix r, this CLA architecture presents a $log_r(n)$ delay.

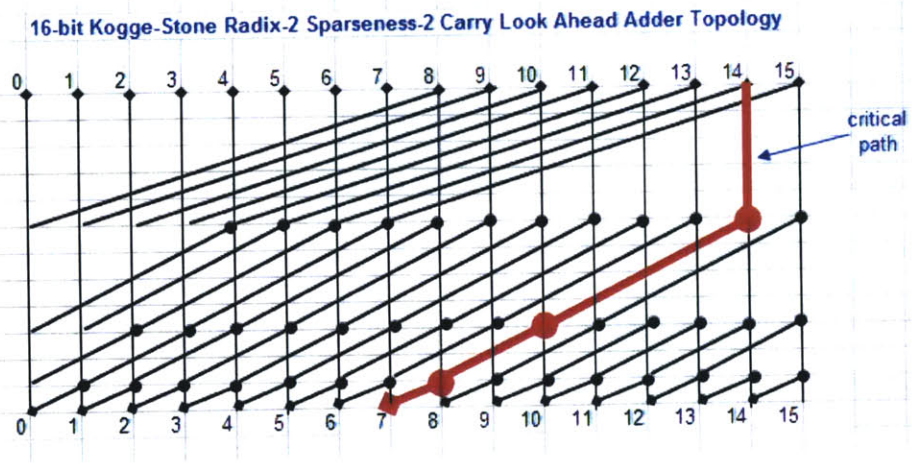

Fig. 3.3. Topology for 16-bit CLA adder design. Critical path is measured at SUM14.

Layout of the three adders in question was performed manually, optimized for density. The domino design has the characteristic of having mostly n-channel devices. The only p-channel devices are pre-charge transistors and auxiliary transistors, such as keepers or within inverters. Since all logic is sized throughout a domino sizing chain, this creates an increasingly large gap in the area of n-channel devices versus p-channel devices with a domino layout cell. A loss in pattern density results within larger cells.

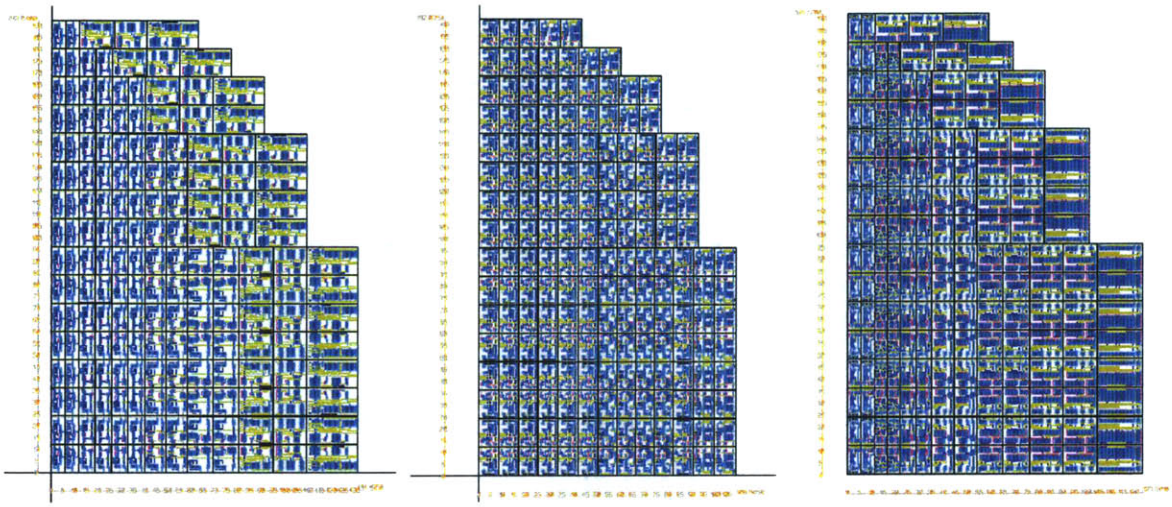

16-bit Domino CLA Adder 16-bit LSDL CLA Adder 16-bit Static CLA Adder

Fig. 3.4. Adder layouts, by circuit architecture.

The LSDL design, though related to the domino design, can be optimized much more for density. The first major difference between the architectures is the minimum sizing of the logic transistors. The result is smaller, more modular transistors that can be used to fill small gaps in the layout. In a comparable domino design, these small patches would remain unused and dummy fill would need to be utilized to fill such gaps. Furthermore, integrated latches add more p-channel transistors to the mix and help to even out the balance of n-channel and p-channel transistors, thus eliminating the area discrepancy between these transistors within the adder cells.

The minimum logic sizing of LSDL has a very significant effect towards the end of the adder sizing chain. In Figure 3.4, notice how much larger the cells toward the end (right edge) of the LSDL design compared to the domino design. It should be obvious that these LSDL cells are much smaller than their domino counterparts. This is primarily due to the fact that the LSDL logic in these cells is close to minimum size, whereas the same logic in the domino design consumes approximately 20 times more gate width area. The LSDL design is thus denser than the domino design overall. But more significantly, the architecture allows for regularity in layout patterns given the modular nature of the smaller transistors throughout. As a result, the pattern density range of this architecture is much tighter than the other adder architectures examined.

The static design architecture has the unique characteristic here of a relatively equal number of p-channel and n-channel transistors. Because all transistors are sized up in the sizing chain, and the quantity of p-channel versus n-channel transistors remains relatively constant per cell, the static design is relatively high-density. However, larger transistor blocks in the static adder decrease modularity, as they did in the domino adder. Small gaps within the cells cannot be filled with smaller devices in this case.

The final areas were determined by each adder's width, since the bitslice height was kept constant across all adders. The domino adder sized 192μm x 131μm, the LSDL adder sized 109μm x 192μm, and the static adder sized 192μm x 124μm. Compared to the LSDL adder, the domino adder was 20% larger and the static adder was 14% larger.

3.4 Results Summary

The idea of regularity was explored at the circuit level, with a case study for limited switch dynamic logic (LSDL) versus more traditional dynamic and static design styles. The sizing approach to LSDL circuits is of interest, as it allows for minimum sized logic and results in smaller designs with higher densities. A 16-bit carry-look ahead adder design was implemented in three circuit architectures: LSDL, static, and domino logic. The density advantages of the LSDL are suggested from the adder sizes, as the LSDL adder consumes 20% less area than both the domino and static adders.

Chapter 4

Robustness Analysis through Historical MOSIS Models

A preliminary performance analysis of the adders previously developed is conducted in this chapter, via historical models obtained from the MOSIS foundry. Each adder design is simulated with nearly sixty run-time SPICE models, each of which encapsulates parametrized process data obtained from an actual MOSIS manufacturing run and contains intrinsic variations. Adders were measured for maximum speeds, average power, and leakage power to gauge the performance of each adder style while exposed to historical manufacturing variations.

4.1 Historical Parametric MOSIS Models

It is often difficult to develop theoretical models for process variation performance, though empirical models are plentiful. One such resource for empirical models is the MOSIS integrated circuit fabrication service for commercial, research, and education institutions. After each run in the MOSIS foundry, data from the wafers is taken to create parametric SPICE models for the run [13]. For each historical run, process data is compiled and a model of best fit is created to simulate the run. These models can be

downloaded from the MOSIS website and used with SPICE decks to estimate the run-to-run performance of the design under measured process variation.

MOSIS models integrate the many variations present in the fabrication process into a net chip-to-chip variation set, capturing all of the correlations within the parameter set for each run. For our studies, we use this intrinsic quality to test for the maximum speed of the various adders, measure average power at these maximum speeds, and measure leakage power all across the numerous process runs. These studies thus focus primarily on the impact of chip-to-chip variations (referred to as $P_{interdie}$ in Equation 1.5). Later chapters will consider systematic layout within chip variations, as well as random process variations.

4.2 MOSIS-Based Robustness Analysis

To test for speed, scripts were used to spawn SPICE decks with clock speeds in the expected range for maximum adder speed. The metric for successful operation was the proper propagation of the logic signal along the critical path of the design. This required the signal to reach a desired voltage threshold and hold state for a certain hold time. Above its maximum speed, each adder did not operate reliably over 90% of the time for the given the process run.

In average power tests using the historical models, adders were set to operate using a constant speed set at the average maximum speed measured for that adder. Power was calculated internally within HSPICE with a .MEASURE POWER AVG statement over the length of the simulation period. Historical MOSIS run models were used to gauge the variation in power due to process at these speed thresholds.

Finally, leakage tests were performed by setting all data signals to 0V and running each adder versus all historical models. Current through the power supply was measured and integrated over the duration of the simulation to calculate total leakage power, as defined in Equation 4.1.

$$P_{\text{leakage}} = V_{\text{dd}} \bullet \int I_{\text{supply}} dt, \text{ for inputs grounded} \tag{4.1}$$

The results of the maximum speed, average power at maximum speed, and leakage power tests can be seen in Figures 4.1 through 4.3. Mean maximum frequencies

Adder	Mean maximum speed	Std.Dev./Mean of maximum speed	Average Power at avg. maximum speed [mW]	Leakage Power [nW]
Domino	1.30 GHz	4.57%	27.9	9.1
LSDL	2.30 GHz	5.69%	23.7	11.8
Static	1.25 GHz (data)	4.76%	10.3	11.8

Table 4.1. Data results for historical MOSIS runs.

were 1.25 GHz for the 16-bit static carry look ahead adder, 1.30 GHz for the domino adder, and 2.30 GHz for the LSDL adder. These numbers, along with other performance data from the historical runs, can be seen in Table 4.1. As expected, the static adder has the lowest maximum speed, and has relatively low deviation near this speed. However, it consumes by far the least average power of the adders tested.

Between the clocked designs of LSDL and domino, the performance difference is quite dramatic. The mean maximum speed of the LSDL adder was close to 77% higher than that of the domino adder. The greater deviation of the LSDL may be attributed to this sizeable difference in speeds being reached.

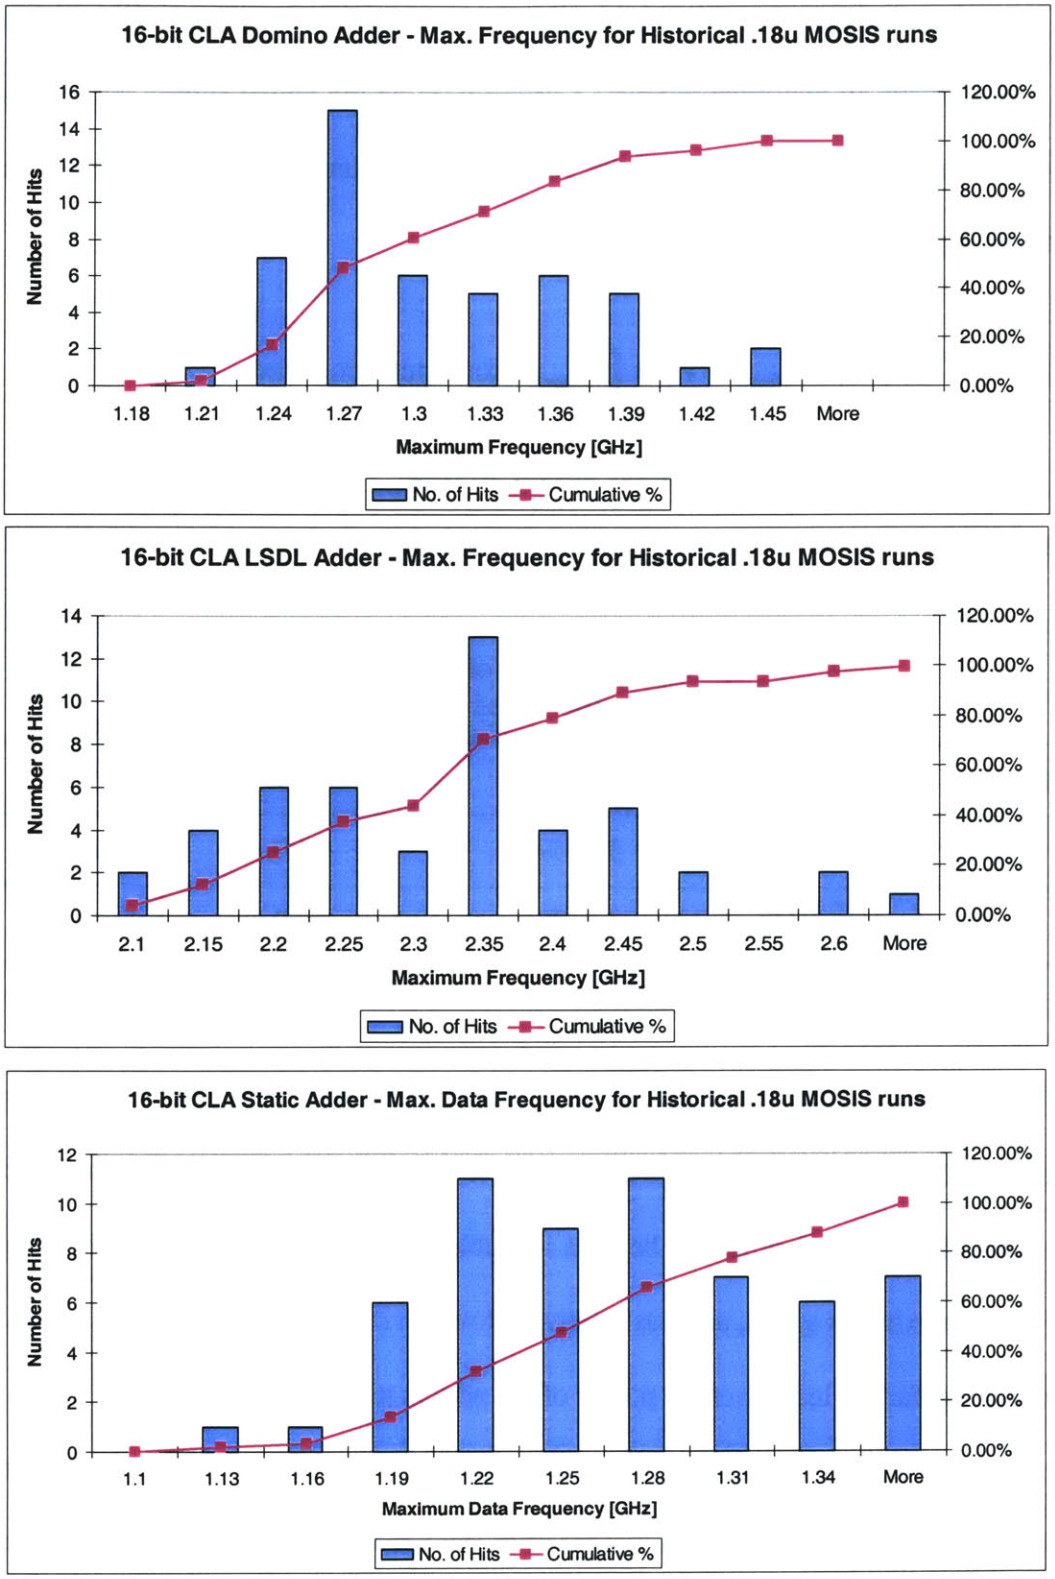

Fig. 4.1. Maximum adder speeds and distributions under historical MOSIS process models.

Moreover, the power difference at the maximum speeds is dramatic, considering that the LSDL adder is being tested at a 77% higher clock speed. Average power of the LSDL circuit is nearly 18% lower at its maximum speed than that of the domino adder. As is clear from Figure 4.2, however, the deviation in power for the LSDL design is smaller than in both the static and domino designs, suggesting good robustness in this respect.

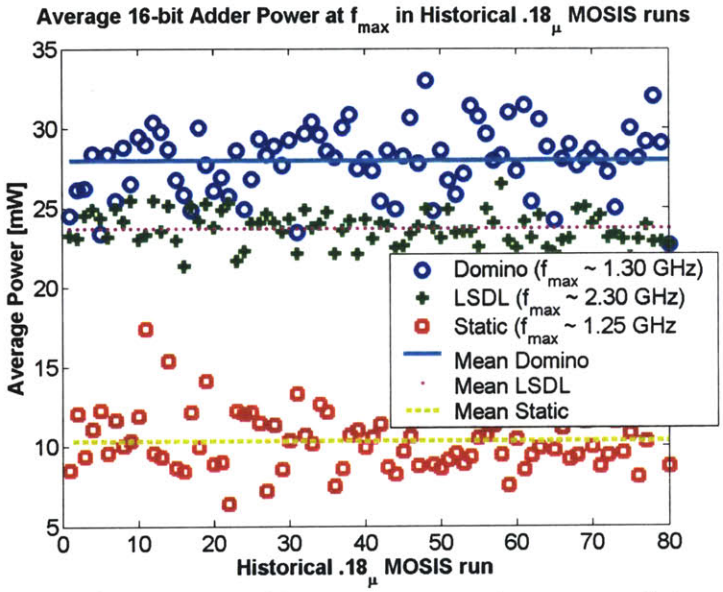

Fig. 4.2. Average adder power at maximum speeds in historical MOSIS process models.

Tests for leakage power did not reveal a sizeable difference among the architectures examined. Many of the techniques used in modern digital IC design to mitigate leakage power, such as high-V_t stacked foot transistors, were not used here. What this analysis suggests is that no design tested here was tremendously more susceptible to leakage current through process variations than another. Both the LSDL and static styles exhibited almost identical leakage power values, as illustrated in Figure 4.3.

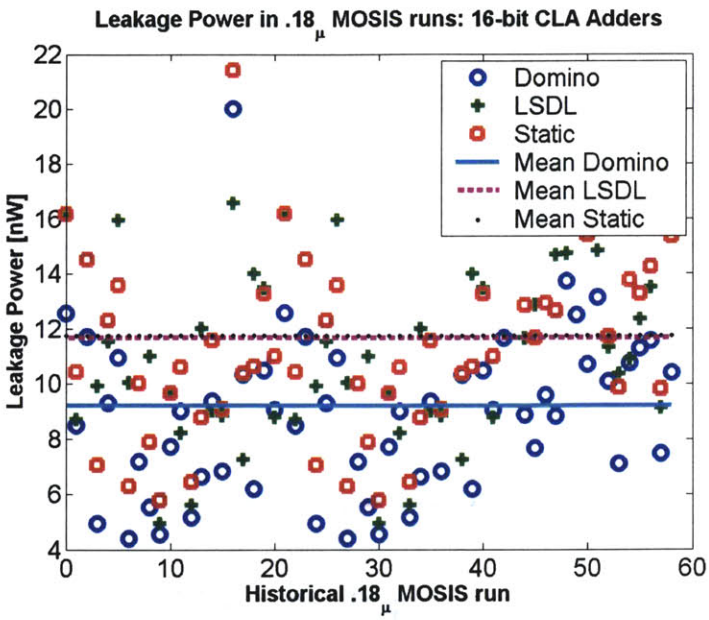

Fig. 4.3. Leakage power in historical MOSIS process models.

4.3 Results Summary

The three adder designs under test were measured for average maximum speed, average power, and leakage power using the fifty-eight SPICE models from MOSIS for their 180 nanometer runs. The goal was to test the designs for robustness with historical parametrized manufacturing variation models. The 16-bit carry look ahead LSDL adder was by far the fastest design, reaching a average maximum speed of 2.30 GHz, followed by average maximum speeds of 1.30 GHz and 1.25 GHz for the domino and static adders. The static adder used the lowest average power at its maximum speed, at 10.3 mW, followed by 23.7 mW and 27.9 mW for the LSDL and domino adders. Leakage power was roughly the same among all the adders. The data highlights the significant speed versus power tradeoff in the LSDL versus static designs.

Chapter 5

Analysis of Circuit Architectures for Process Variation-Induced Errors

Many of the end effects of process variation in circuit performance are well known, such as clock skew and the extra slew caused by CMP-related capacitive interconnect load variation. We next examine how various circuit architectures behave under non-ideal input conditions for the clock, data signal, or individual process variables, for example. We also introduce a particular form of systematic within die variation, $P_{intradie}$ as described as Equation 1.4, in the form of a slanted plane model. This is a local linear process variation gradient across a chip, and we evaluate the impact of such a gradient on each design. Finally, we look at pattern dependent layout variation for the adders, to quantify the potential effects of each design layout on process variation. In this case, a more detailed process model is used to evaluate systematic within-die variation as a function of position, $P_{intradie}(x,y)$. The work in this chapter attempts to better understand the underlying strengths and weaknesses of the various circuit families to formulate ways of increasing circuit robustness through design.

Adder designs throughout this chapter were tested at 77% of maximum recorded operating speeds with no parasitics or variation introduced. This was done to ensure a comfortable operating region for each design, such that 130% of the tested speed fell

under the maximum speed. The tested speeds were 2.0 GHz, 1.42 GHz, and 1.25 GHz for the LSDL, domino, and static designs.

5.1 Clock Jitter

Skew and jitter are two of the major problems in clock distribution throughout a circuit. Clock jitter refers to the temporal variation of the clock period at a given point on the chip [15], depicted in Figure 5.1. There are three classified sources of jitter: clock signal generation, environmental variation, and coupling capacitance.

During clock-signal generation, noise can couple into the voltage-controlled oscillator that powers the clock, introducing temporal variation to the clock signal. Another common source of jitter is power-supply variation, which is dependent heavily on switching activity. The final major source of jitter is coupling capacitance between the clock wire and adjacent data lines [15].

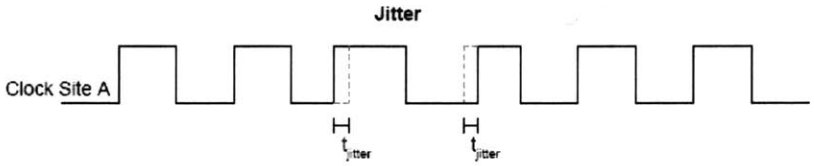

Fig. 5.1. Clock jitter [16].

In comparing the performance of the 16-bit carry look ahead domino adder versus its LSDL counterpart, careful consideration was taken to ensure viable jitter analysis. The amount of jitter introduced was proportional to the clock period and introduced as a varying clock duty cycle. Duty cycle for each clock was varied from 30% to 50% in small intervals to emulate a jitter effect. This translates to clock high durations varying

from 210 ps to 350 ps for the domino adder and from 150 ps to 250 ps for the LSDL adder, all in 10 ps intervals. The effects of the tests are demonstrated in Figure 5.2.

Here, the numerous jittery input clocks are plotted together with the corresponding outputs for the adder critical path (SUM14). Jitter seems to affect performance in the LSDL adder less than in the domino adder. Because the LSDL was designed to operate on short clock pulses, it is better suited to withstand the temporal clock variations in pulse width. The domino adder does not have this capability and fails when the clock pulse becomes too narrow, as noted by the thick baseline at 0 V. Extreme jitter causes latches throughout the LSDL adder to lose charge, as exhibited by the decreasing values of the high-value on the output. Based on this testing trend, more jitter than this will eventually cause logical failure. However, it is apparent that LSDL is much better built to handle jitter in the signal.

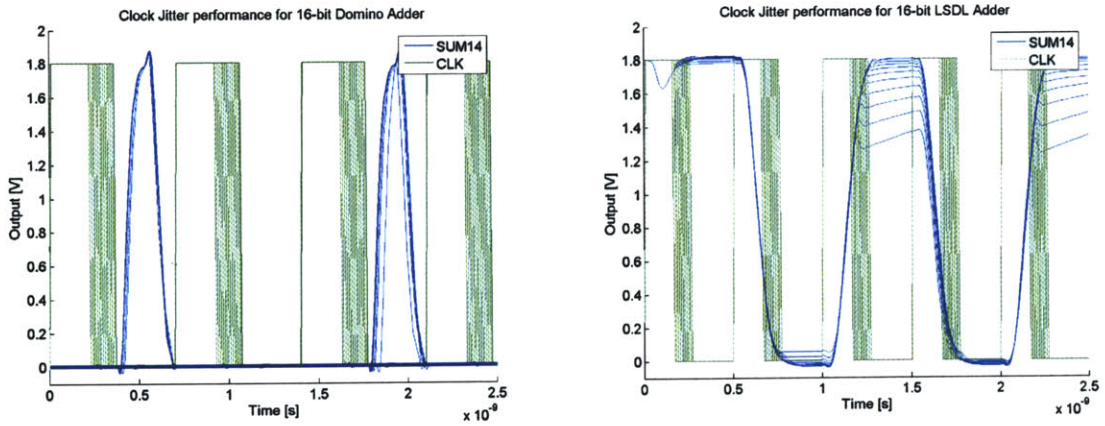

Fig. 5.2. Clock jitter comparisons for domino (left) and LSDL (right) adders.

5.2 Clock Skew

The spatial variation in arrival time of a clock transition on an integrated circuit is commonly referred to as clock skew [15], depicted in Figure 5.3. The most commonly

attributed sources of clock skew include manufacturing device and process variations, intra-chip interconnect variation, and environmental variations.

The manufacturing stage introduces many of the sources of skew. Non-uniformity in pattern density can result in uneven plating topography and may result in excessive post-CMP dishing or erosion due to induced non-uniformity in the polishing rate across the interconnect. Such unevenness introduces localized resistance and capacitance variations across interconnect, which may add or detract unexpectedly to or from the clock load. Gate oxide thickness variations affect threshold voltage and gain, and dopant variations may cause otherwise identical devices to perform differently, for example. Some sources of process variation are random (e.g. random dopant variations), while others are systematic (e.g. pattern density effects). Design for manufacturability seeks to minimize the creation of systematic variation and to make the circuit as robust as possible to both systematic and random variations.

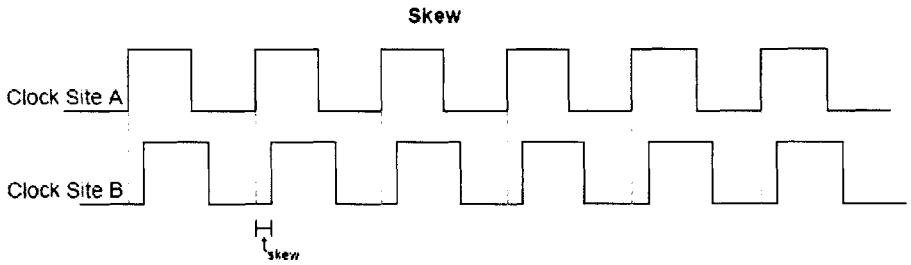

Fig. 5.3. Clock skew [16].

Environmental variation is the last major category of skew sources. Especially suspect is temperature variation across chip. The temperature itself is of concern, but of greater concern is the distribution, or uniformity, of the temperature across chip.

For clock skew analysis, clock delays between -10% and +10% of the clock period were introduced at the final sum stage of each adder, emulating an improperly

balanced clock tree. This translates to -50 ps to +50 ps clock delays for the LSDL adder and -70 ps to +70 ps delays for the domino adder at the sum stage, in 10 ps intervals. Comparing the clocked domino and LSDL adders in their robustness to clock skew, we again see that the high performance of the LSDL design may come at the expense of extra sensitivity to clock timing. For sum stage clock skews of -50 ps to -30 ps, the LSDL adder has difficulty changing output state, suggesting more vulnerability to negative clock skew than its domino counterpart. Though the latches in the LSDL design create better response to jitter, they also introduce numerous timing constraints for achieving correct operation. Clock skew has the ability to push these constraints and cause errors, as seen in Figure 5.4. The domino design continues to respond strongly despite the shifting clock signal.

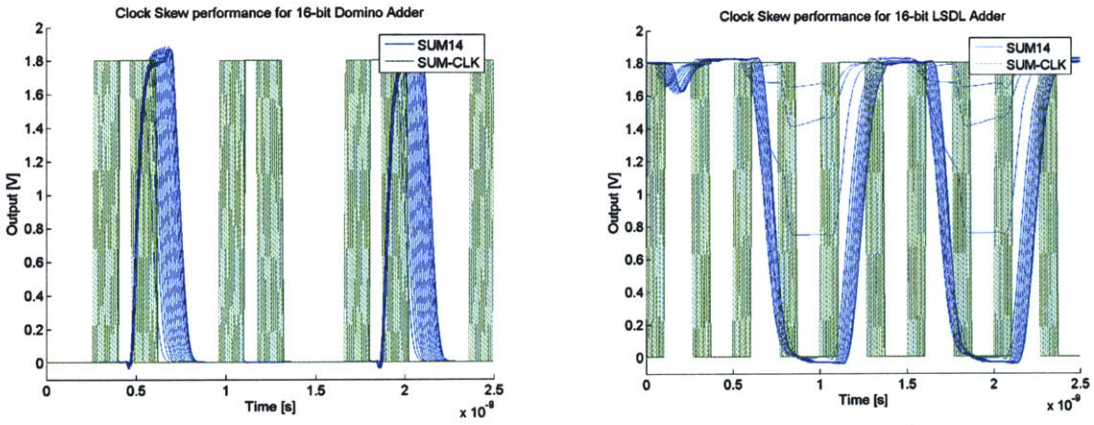

Fig. 5.4. Clock skew comparisons for domino (left) and LSDL (right) adders.

5.3 Clock Slew

Slew rate of a signal refers to the rate of change in the rise or fall of that signal and is widely considered an operational problem in digital design, as suggested in Equation 5.1.

A large magnitude for a slew rate, as depicted in Figure 5.5, is undesirable, as it can cause numerous operational problems in a circuit.

$$\text{Slew rate: } \frac{d}{dt}(V_{high \rightarrow low}) \text{ or } \frac{d}{dt}(V_{low \rightarrow high}) \tag{5.1}$$

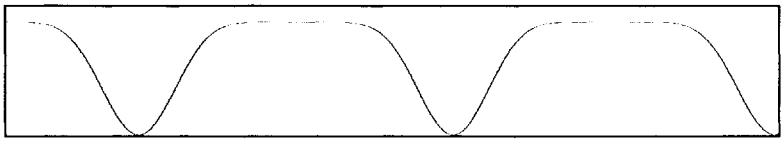

Fig. 5.5. Non-desirable clock slew.

Clock-dependent architectures are especially vulnerable to slew, as excessive delays in changing clock state often lead to transmission delays and logical errors. Clocked latches are also vulnerable. Edge-based latches, for example, are designed to work with sharp clock edges. When fed a lazy clock signal, these latches may leak charge, as the edge of a signal can become harder to detect with slewing. In pipelined designs, which depend on multiple stages of signals (like an adder), the probability for logical failure increases at the output. Slow output transitions decrease the gain in a circuit and harm the circuit's ability to drive a large load. In a multiple stage design — such as an adder or multiplier — this is a critical element, since stages are accurately sized to be responsible for their loads.

LSDL is an integrated edge-based latch, intrinsically pipelined, clock-based logic. Due to its architecture, we can expect this circuit style to be vulnerable to clock slew. The domino design is less latch intensive than the LSDL design and probably less vulnerable to such a problem. Figure 5.6 visualizes the responsiveness to clock slew in both the LSDL and domino cases. In this test, paired clock rise and fall times from 0 ps

to 160 ps were introduced for each adder, with the input clocks and corresponding outputs plotted. Clock slew does affect the performance of LSDL and domino as predicted. The slow response of the clock seems to delay the entire response of the adder, though this response is more severe in the LSDL case. The gain in stages of both designs was not significantly affected, likely due to the inverter-based driving buffers inserted between stages.

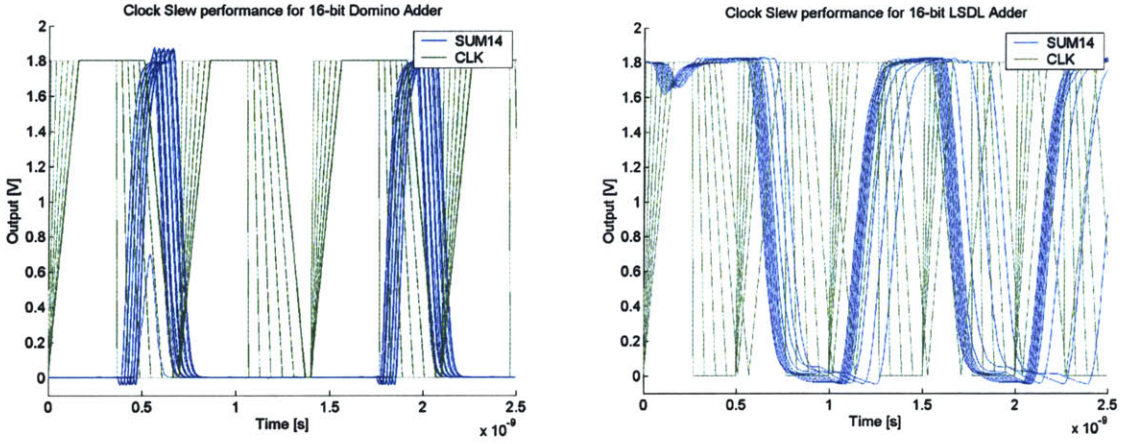

Fig. 5.6. Clock slew comparisons for domino (left) and LSDL (right) adders.

5.4 Data Slew

All periodic signals are vulnerable to slewing, including data signals. Though it is often not as dire a condition as clock slew, it is nevertheless worthy of examination in robustness studies. With a slow responding data signal, the logic state may not be evaluated in time for both dynamic and static designs. From Figure 5.7, we can detect three trends in performance to data slew. To test data slew, these adders were exposed to a perfect clock, with data signals of varying slew rate, with rise and fall times of 0 ps to 160 ps. Dynamic logic, when hold time and arrival time requirements are met, is

relatively robust to data slew. This type of design only evaluates data during the high clock state, and can withstand data slewing, given proper timing considerations during design.

In the static adder, we see the ripple effect of data slewing at the output. Since no clock is present to shut off current, slow arriving data ripples through logic to directly affect the arrival of the output. Despite the uncertainty in arrival time, static logic is still quite robust to errors caused by data slewing. Data slewing in static logic, unlike dynamic logic, may contribute to glitching. As the data signals settle and the logic evaluates these settled signals, glitches may occur. However, these glitches seen were minor and did not contribute to logical failure. In dynamic logic, as we see in the case of LSDL, mistimed data arrival or slewing can lead to catastrophic failures. Dynamic logic has the potential to discharge state given an erroneous signal, without being able to charge up until the next clock cycle.

One disadvantage of LSDL is its vulnerability to clock and data slew, moreso than regular domino logic. Integrated latches also introduce another area of vulnerability to misbehaving signals. Process variations, whether by affecting the capacitance and resistance of wires or by introducing a temperature gradient across chip, have the potential of inducing these slewing effects.

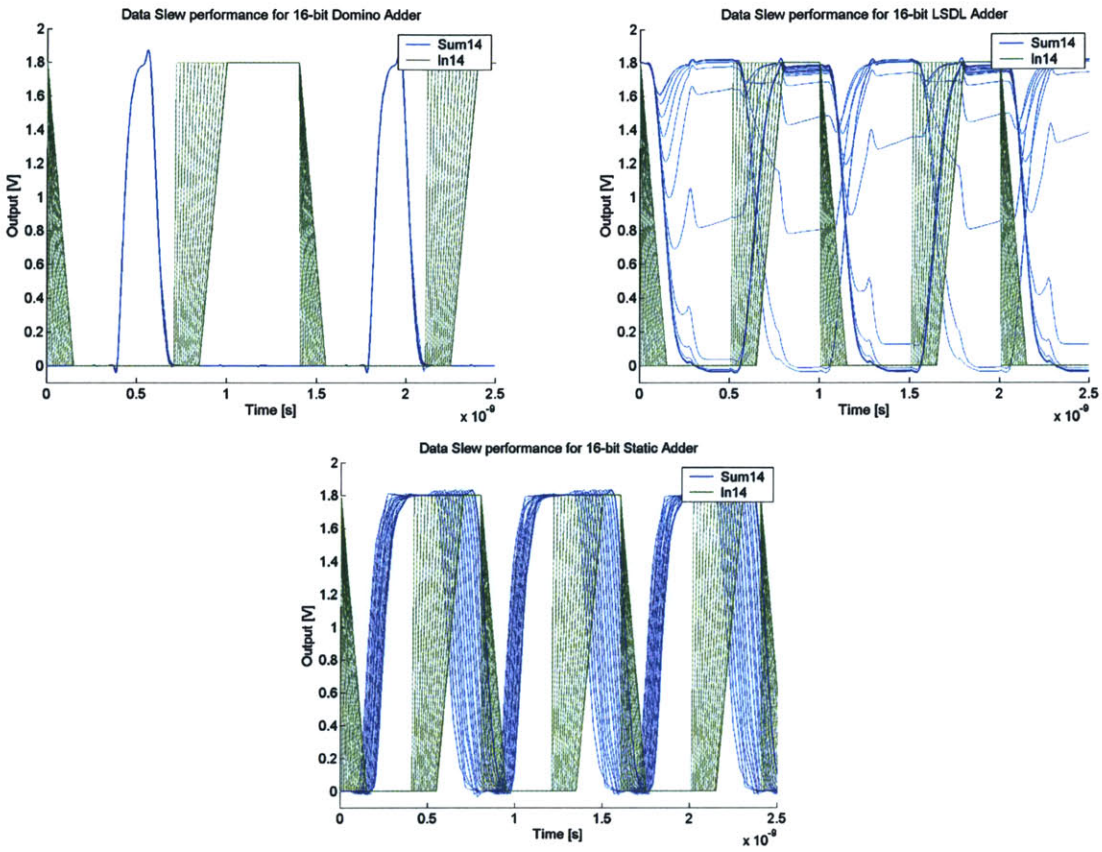

Fig. 5.7. Data slew comparisons for domino (upper left), LSDL (upper right), and static (bottom) adders.

5.5 Slanted Plane Effect

A certain degree of spatial variation is introduced across the wafer by large scale tool and process asymmetries during a process run. At chip-scale, we can experience a wide range of different spatial patterns, resulting from a projection of a complicated wafer-scale variation across different chips on the wafer. Though a linear trend in these patterns is rare across a whole wafer, it is much more likely on a centimeter or chip-scale. The different local process variation gradients that may develop across a single chip are the focus of this section. These directional trends have been termed "slanted plane" effects in processes and are illustrated in Figure 5.8. For simplicity, the orientation of the wafer

plane can be reduced to a two-dimensional model for analysis. Process variation can be introduced in a gradient fashion in such a model, to simulate the effect of such wafer orientations on circuit performance. In particular, we consider a variation source as described by Equation 1.4, repeated below, where ω_0, ω_x, and ω_y can in general be random variables. In our analysis here, we consider eight planes as shown in Figure 5.8.

$$P_{intradie}(x,y) = W(\omega, x, y) = \omega_0 + \omega_x x + \omega_y y \tag{1.4}$$

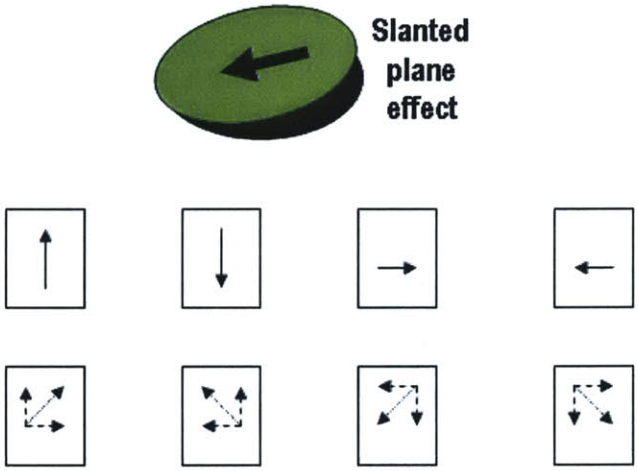

Fig. 5.8. Slanted plane chip-scale effect and directional trends examined.

In the analysis, circuit layouts were examined to determine a spatial grid for applying the gradient variation. In Figure 5.9, one of the adder designs is divided into 11 grid regions between the cells, depicted as the smaller box structures in the adder. All regions were simulated with values for threshold voltage and temperature deviations, from 0 to 10% of the nominal values, based on their location, to emulate a simple process variation gradient. Temperature and oxide thickness, which in turn affects threshold

voltage, as determined in Equation 5.2, were determined to be the most applicable spatial random variables to examine based on processing techniques.

$$V_t = 2\varphi_b + \frac{\sqrt{2\varepsilon_{si} q N_A \varphi_b}}{C_{ox}} + V_{fb}, \text{ where } C_{ox} = \frac{\varepsilon_{ox}}{t_{ox}} \text{ and } \varphi_b = \frac{kT}{q}\ln(\frac{N_A}{N_i}) \qquad (5.2)$$

The magnitude of the cell variation depends on the directional trend being examined and the grid region in which the individual cell is located. The spatial variation analysis is summarized in Figure 5.10 on three different 16-bit adder architectures, where adder performances resulting from a 0 to 10% variation trend with strengths into eight different directions are simulated.

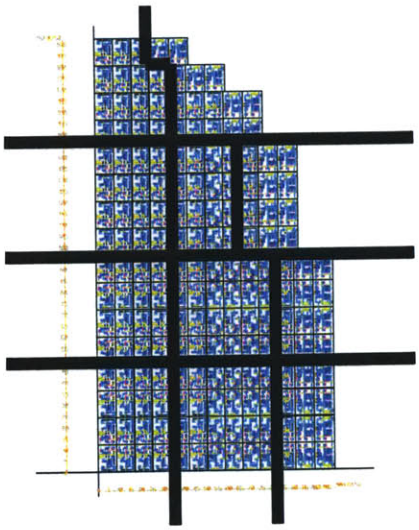

Fig. 5.9. Grid division of design for slanted plane performance simulation.

From this analysis, it appears that for even such a huge gradient in these process variables (10% across a single 400 μm chip), there is little discrepancy in overall performance. Moreover, there appears to be little difference among the directions of the variation trend. From Figure 5.11, we see similar signals at the outputs despite the

direction trend. Therefore, it appears that the slanted plane effect, for the variables considered, may have a negligible effect on the performance of this design. However, it has to be said that these designs, measuring approximately 150 microns by 300 microns, are very tiny compared to standard ASIC and VLSI designs, which may be more exposed to the slanted plane effect.

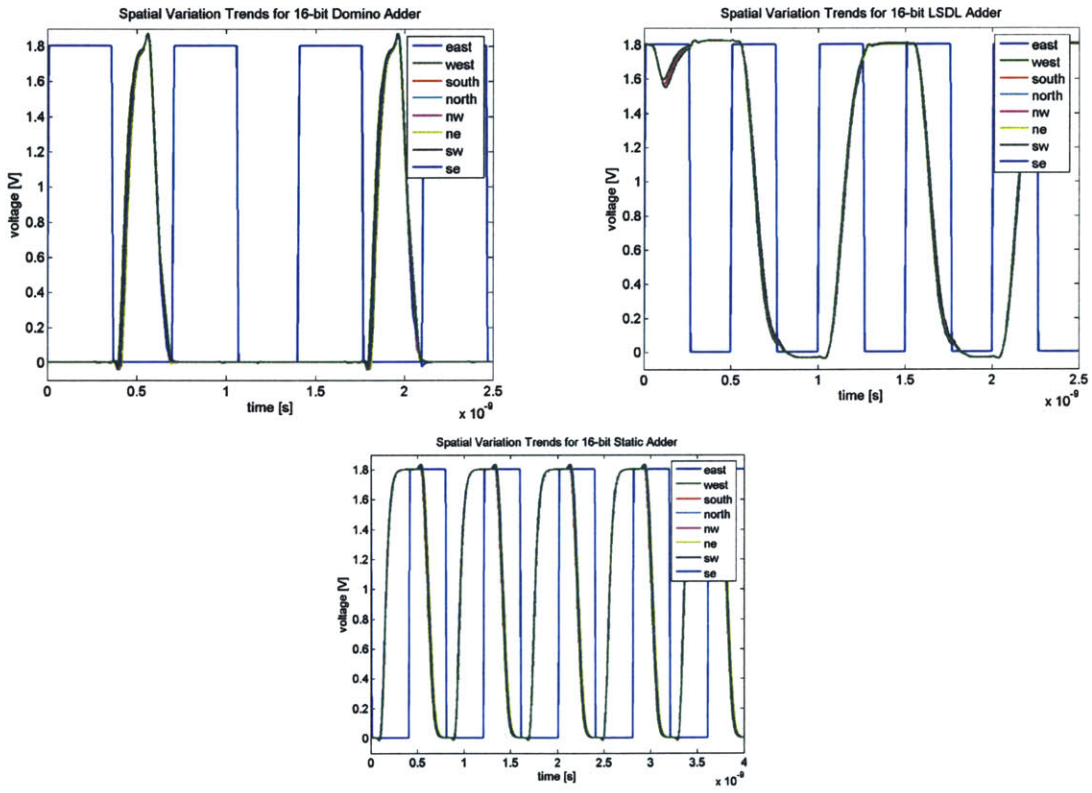

Fig. 5.10. Slanted plane effect analysis for eight directional trends (strength into direction).

5.6 Isolated Process Variable Analysis

Understanding the impact of individual process variables on the performance of the adders is a critical element in painting a complete picture of the sensitivities to process variation. These tests not only serve to expose the vulnerabilities of different designs, but will also serve to verify the later results attained through Monte Carlo regression analysis.

Each design was tested at 77% of its maximum speed, with the individual process variable examined modified to approximately 110% of the nominal value. This corresponds to a simple form of variation, $P = P_0 + \Delta P$, where ΔP is 10% of P, as described in Equation 1.3. Only one variable is modified per run for each adder, along with a nominal run where no variables are modified. Standard deviations of the distributions were neglected in these cases, since performance trends rather than actual performance numbers were desired. The variables included n-active region critical dimension, polysilicon critical dimension, temperature, oxide thickness, voltage, and threshold voltage. The results are depicted in Figure 5.11.

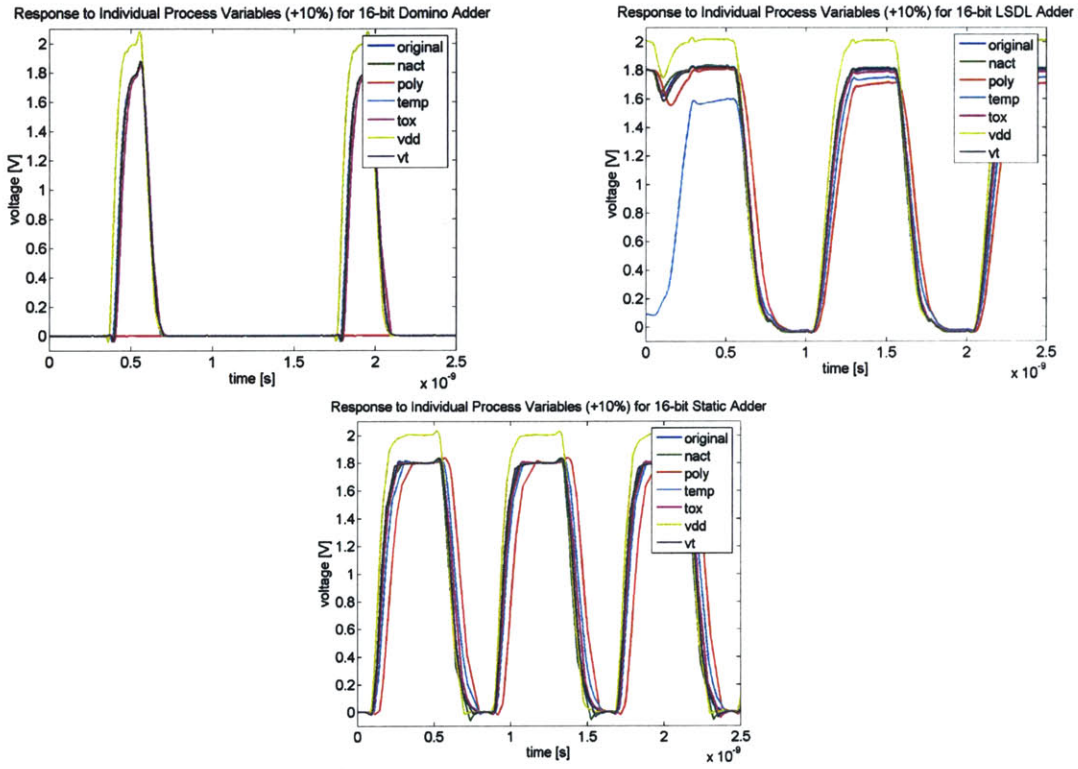

Fig. 5.11. Impact of isolated process variable (110% of nominal value) on adder performance. Domino adder (upper left), LSDL adder (upper right), static adder (bottom).

We notice in the domino case, a 10% increase in polysilicon critical dimension would result in a 100% failure rate, crippling the adder. The LSDL adder adjusts to individual process variables through a decreased high output value and increased fall time. The LSDL adder may be affected by decreased drive strength, which would cause the output to be lower than expected. The static adder generally appears the most robust to the individual variables. This adder is mainly affected through increased output signal delay and not as much through output magnitude.

5.7 Layout-based Variation Trends

A tremendous amount of knowledge can be extrapolated from a design layout. Tools exist to extract such vital information as pattern density, effective feature width/length, minimum feature width/length, maximum feature width/length, and process binning. Each of these tests speaks a great deal about the layout-induced variation that may occur in the design, though pattern density extraction is most widely used of these tests.

As described in greater detail in previous chapters, pattern density affects the topographical uniformity of a post-plated chip. Uniform pattern density is ideal for processing. Therefore, the range of pattern density across a design is more significant than the magnitude of the pattern density.

By examining the pattern density of the CLA adders built and described in Chapter 3, we can develop further insight into the impact of circuit architecture on manufacturability. The pattern density across all designs can be examined more closely in Figure 5.12. Polysilicon pattern density in the static adder hovers near 13%, compared with 7% for domino and 14% for LSDL. In both the domino and static adders, we see a

wide range in pattern density of nearly 12% and 11%, respectively. In comparison, the LSDL adder exhibits a 5% range in pattern density. This uniformity in density of the LSDL adder is quite significant, for it helps ensure the evenness of many processing steps, such as plating, etch, and CMP.

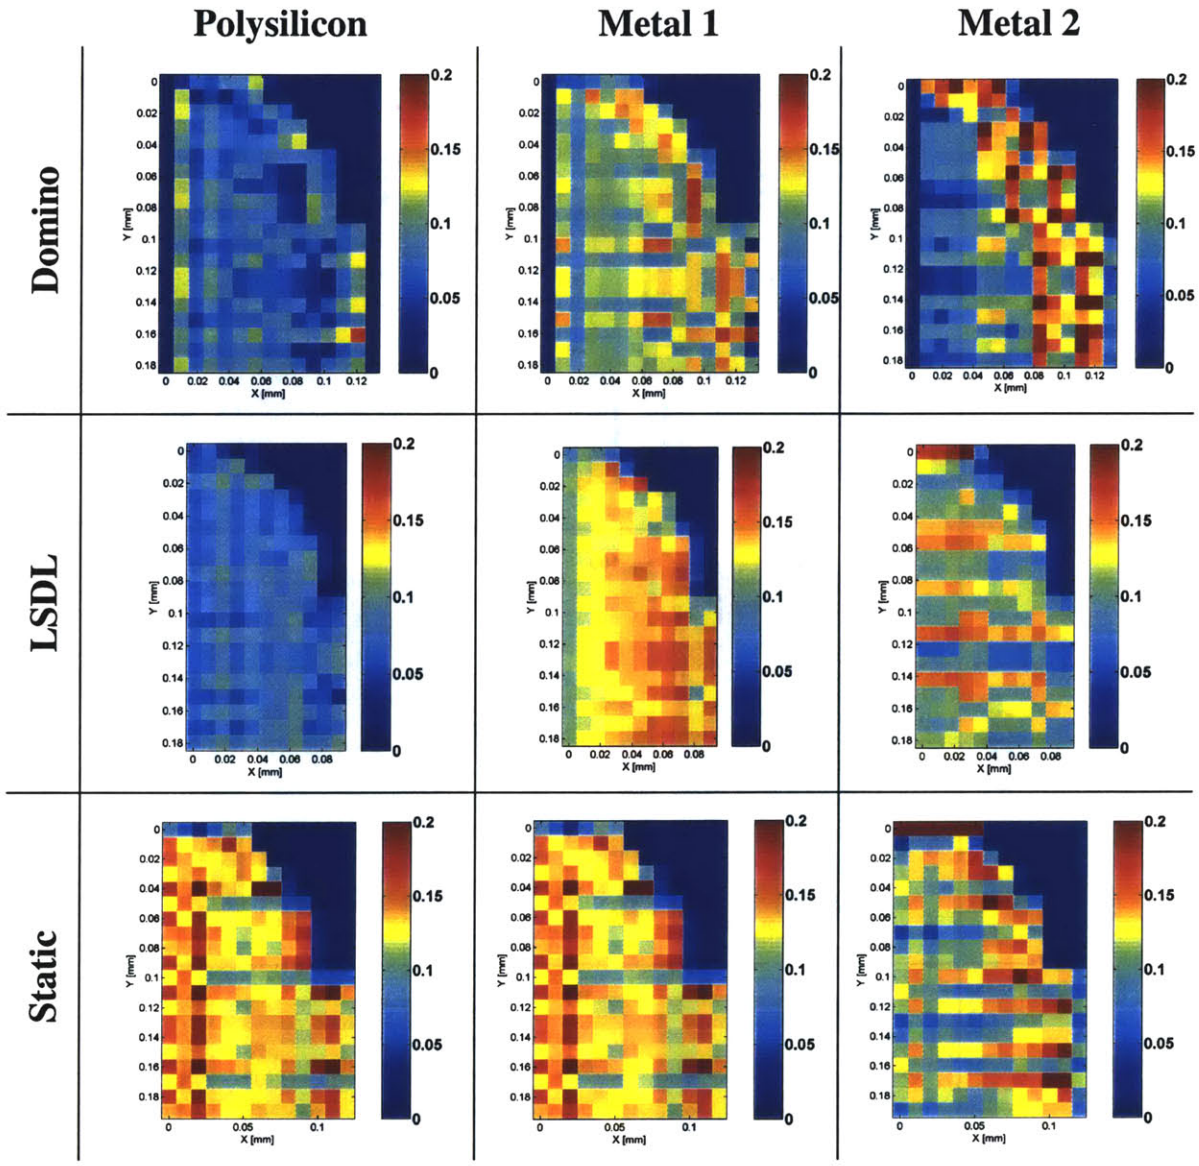

Fig. 5.12. Adder pattern density analysis.

Feature size binning is another useful layout-based metric to examine, visualized in Figure 5.13 for the adders. Dishing in CMP depends on line width, and thus a wide range of line width would indicate large variation of dishing. When most lines are in one

size bin, for example, this indicates a small range of widths in a chip, thus small dishing variation. Generally, the range of variability often depends on the range of pattern sizes. Thus, there is smaller range of variability (of dishing, thickness, etc.) for smaller range of pattern sizes. From Figure 5.13, we can see that LSDL has the smallest range of pattern sizes. Therefore, LSDL enables more uniformity in feature sizes than static and domino.

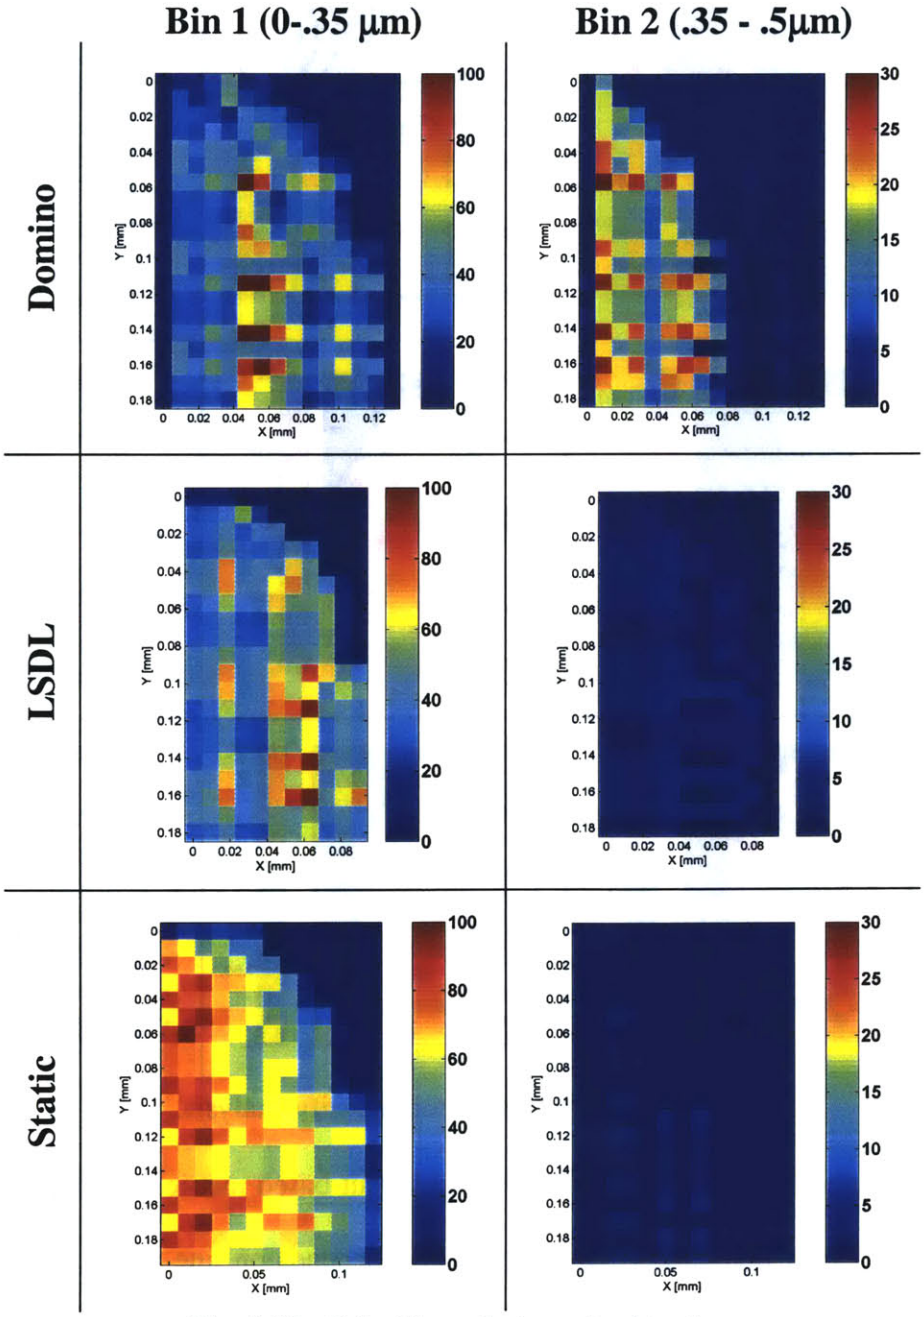

Fig. 5.13. Polysilicon feature size binning.

5.8 Results Summary

In comparing static, domino, and LSDL adders in an assortment of tests, the LSDL design stood up much better to clock jitter, but showed vulnerability to potential clock skew and poor clock and data slew rates. Care must be taken to ensure proper timing constraints are met in the high performance LSDL design. Overall, the static design stood up the best to single process variables and the domino design behaved the worst. A method to model the slanted plane effect, which suggests a local linear process variation gradient across a chip, was devised and implemented for the adders, but exhibited little performance effect for the variables tested. Pattern density analysis revealed much less density variation in the LSDL adder and more optimal feature size binning, compared to the other adders. This suggests that the LSDL design may induce less pattern-dependent manufacturing variations.

Chapter 6

Monte Carlo-Based Variation Analysis Engine

Since there are thousands of potential sources of variation in thousands of strengths, an accurate model of systematic and random process variation is difficult to create. Previous chapters have considered pattern dependent intrachip variation, as well as interchip variations. Random variations remain to be considered. Here, integrated Monte Carlo methods within HSPICE are used to compare the characteristics of several adder architectures simulated with varying degrees of both random interdie and random intradie process variation. This chapter describes the methodology built to conduct these Monte Carlo analyses. The breakdown of interdie versus intradie tests is described. Sensitivity analysis and correlation statistics are performed on the Monte Carlo data, to then extract key vulnerabilities of each adder style.

6.1 Methodology

With any deterministic black box system, like a digital integrated circuit, introducing a certain set of inputs to the system will result in reproducible outputs to the system. This property, though subtle, enables a powerful technique for empirically analyzing variation in the design under test. By running through a diverse set of inputs, as a Monte Carlo simulation may do, the designer may map the wide range in performance of his or her

design in response to distributions of those inputs. Ultimately, with enough data, this input to output relation will reveal information about the sensitivities of the design to the input variables introduced. This chapter discusses the methodology for ultimately enabling such sensitivity analysis.

This system takes inputs of process variables, their corresponding probability distributions, the designation as an interdie, intradie, or a total variation test, and modifies the SPICE deck under test accordingly. The output of the system is the run-time data from the hundreds of Monte Carlo runs. This model assumes that the underlying simulation system correctly captures the functional dependency of the output on both the SPICE deck and the process models and distributions.

HSPICE was chosen as the simulation tool of choice, given the hundreds of process model variables it uses for simulation as well as its built in Monte Carlo functionality. One may select as many of the HSPICE model variables as desired to behave as random variables to the designed system. For every model variable selected as a random variable to the system, a probability distribution must accompany it. The introduced random variable distributions are all assumed to be Gaussian and uncorrelated with one another. The probability distributions are important as to avoid overexposure to certain model variables in the system and skew sensitivities in the final results. For this methodology, distributions were extracted from MOSIS models and used for the test random variables introduced.

Before the methodology can be implemented, the type of system variation test must be selected among mismatch (intradie), interchip (interdie), and overall variation, the combination of both previous types. The mismatch variation model takes in Δw and

Δl as Gaussian random variables as well, adding these values to every transistor's gate width and length, with new values per instance. The interchip variation model selects for simulation only those model process variables that affect a design from the macroscopic scale. This model therefore does not introduce the Δw or Δl random variables per transistor. The overall variation model introduces both macroscopic and microscopic random variables to the design, combining the mismatch and die-to-die (inter-chip) variation models.

The final component is the SPICE deck for the design under test. This netlist will be modified based on the other inputs to establish the appropriate simulation. After the random variables, distributions, and variation model have been chosen, the output measurements desired on the output signal must be specified with the .MEASURE statements in the SPICE deck. Once complete, an automated script runs through the SPICE deck of the design to parse and modify additions to enable the Monte Carlo functionality and introduce the random variables. A pseudo-random seed is chosen and the Monte Carlo simulation is launched.

Upon completion of the simulation, a parser script scans the piped output of the simulation for the .MEASURE keywords used. This script records the input random variables, their values, the output measurements, and their values into an ASCII table for further numerical analysis.

The analysis stage is the final stage in this Monte Carlo methodology. Once the data table has been built, various Matlab scripts are used to perform multiple regressions between the inputs and outputs. Based on the regression model in Equation 6.1, we perform a first order polynomial fit to calculate the slope of the output metric versus the

input variable. We then proceed to calculate the relative input to output sensitivity metric.

$$y_j = a_0 + a_1 x_1 + a_2 x_2 + a_3 x_3 + \ldots + a_n x_n \tag{6.1}$$

$$S_{x_{ij}} = \frac{\partial y_j}{\partial x_i} \cdot \frac{\sigma_{x_i}}{\sigma_{y_j}} \tag{6.2}$$

$$S_{x_{ij}} = \frac{\partial y_j}{\partial x_i} \frac{\overline{x_i}}{\overline{y_j}} \tag{6.3}$$

Although many ways of defining sensitivity exist, this analysis uses the method outlined in Equation 6.2, rather than the more common method detailed in Equation 6.3. In the latter, each input variable is assumed to be as important as the next, even if it contributes predominantly to the final output metric. In the former method, which is employed in this analysis, each input variable is weighted based on its contribution to the variation in the final output metric. This method is generally a good technique for analysis of global sensitivity, as would be done for a Monte Carlo simulation [18].

By calculating these normalized sensitivities, the designer may be able to identify key weaknesses in his or her design and rebuild these areas for robustness improvement. The entire Monte Carlo analysis process is depicted in Figure 6.1.

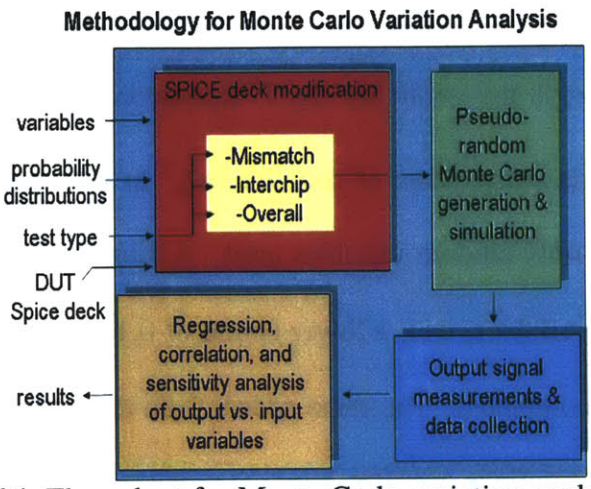

Fig. 6.1. Flow chart for Monte Carlo variation analysis methodology.

The following sections highlight the use of the developed Monte Carlo methodology on the designed adders under test. By comparing the performance and robustness of the static, domino, and LSDL adders to a barrage of process variations, regressing on the outputs, and calculating design sensitivities, the goal is to analyze the strengths and weaknesses of each adder quantitatively. These adders are tested at the same speeds as in the previous chapter, at roughly 77% of their maximum operating speeds when no parasitics or variations are introduced. The following subsections discuss these tests and analyze the results.

6.2 Mismatch Variation

Mismatch has generally been a problem with analog integrated circuits, where accurate current or voltage biasing in pair circuits is required. For these tests, we generally define mismatch as the marginal process errors encountered when processing transistor gate widths and lengths, independent of the transistor. This harmful effect was simulated by adding the independently instantiated Gaussian random variables Δw and Δl

to all of the individual transistor gate widths and lengths. This is described in Equations 6.4 and 6.5, where w_i is the width for the i^{th} transistor, w_{0_i} is the nominal (designed) width, and each Δw_i is independent of Δw_j, for $i \neq j$. In this analysis, these gate widths and lengths were the only random variables used. A three-sigma mark of 20% was used in introducing these variables, with a lower bound of 0.18 μm for gate length, to ensure simulations actually completed. The same adders from the previous tests were used, and an overlay of all of the outputs over two-hundred runs can be seen in Figure 6.2.

$$w_i = w_{0_i} + \Delta w_i \tag{6.4}$$

$$\Delta w_i \sim N(0, \sigma_{w_i}^2) \tag{6.5}$$

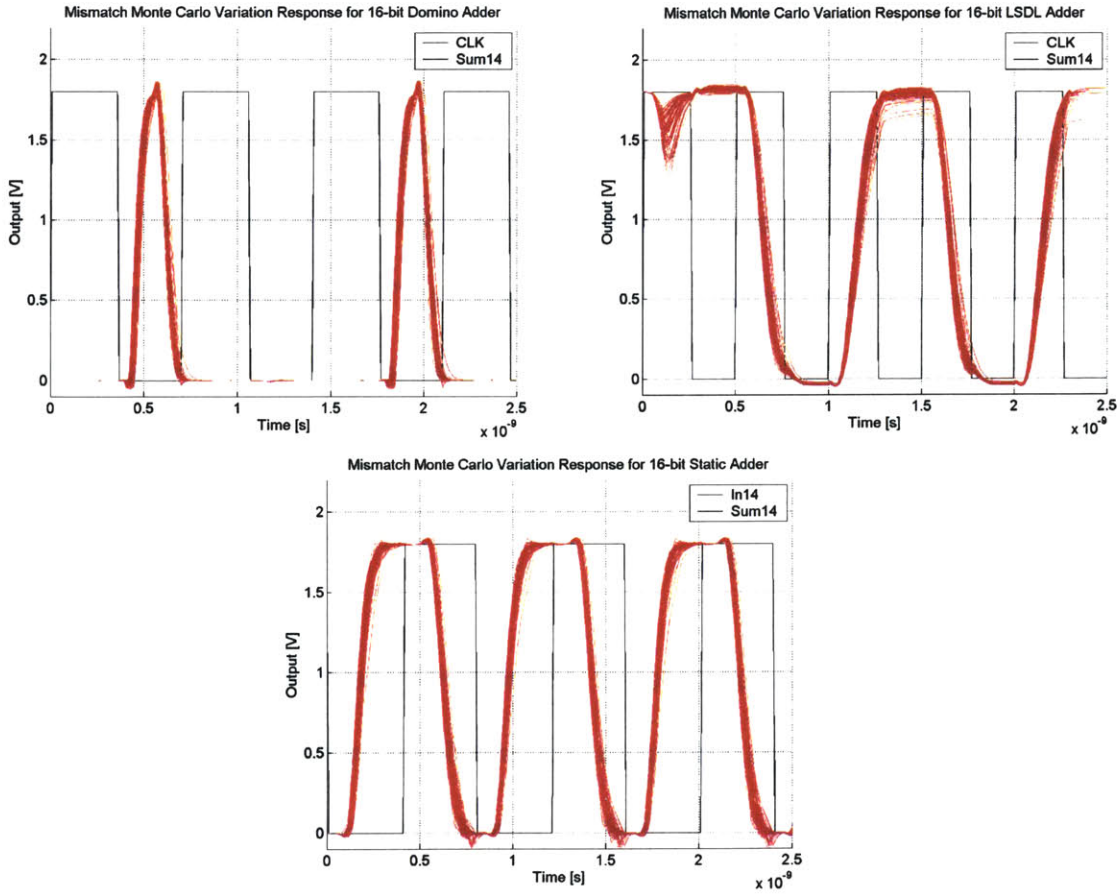

Fig. 6.2. Mismatch Monte Carlo variation analysis for adder architectures.

Mismatch effects appear minimal in both the domino and static adder designs, but more noticeable in the LSDL adder. Peak voltage variation is visible during operation of the LSDL adder when mismatch is introduced. For the LSDL design, mismatched transistor sizes may be contributing to charge sharing at the output of the stage latches, causing a reduced high value at the output.

6.3 Chip-to-Chip Variation

Chip-to-chip, die-to-die, or interchip, variation effects are not localized like mismatch errors. These are the variations caused by more macroscopic problems that may affect the entire design, such as oxide thickness, temperature, or other environmental elements acting across the entire chip as a whole. We refer to Equation 1.5 again to model our interchip variation, introducing random variables that may affect a design more globally.

$$P_{interdie} = P_{fab\text{-}to\text{-}fab} + P_{lot\text{-}to\text{-}lot}(\text{fab}) + P_{wafer\text{-}to\text{-}wafer}(\text{lot}) + P_{die\text{-}to\text{-}die}(\text{wafer}) \qquad (1.5)$$

For this analysis, the Gaussian random variables introduced were the n- and p-oxide thickness, active layer length, polysilicon channel length, threshold voltage (function of doping, a macroscopic trend), temperature, and supply voltage. The results of this analysis can be seen in Figure 6.3.

First, we see that domino logic, for the most part, is sensitive to such interchip variation elements. The shape, timing, and correctness of the output vary greatly with these large-scale variables. In several cases, logical errors can be spotted. On the other hand, LSDL looks more tolerant to this category of variation than its static competitor.

One common trend discovered during these tests was the fact that LSDL was less robust against mismatch variation than the static adder, yet it was more robust against interchip variation versus the same adder. Though these are subtle differences, it is still interesting to note that LSDL, while a dynamic logic, is competing on par with its static counterpart in robustness to process variation.

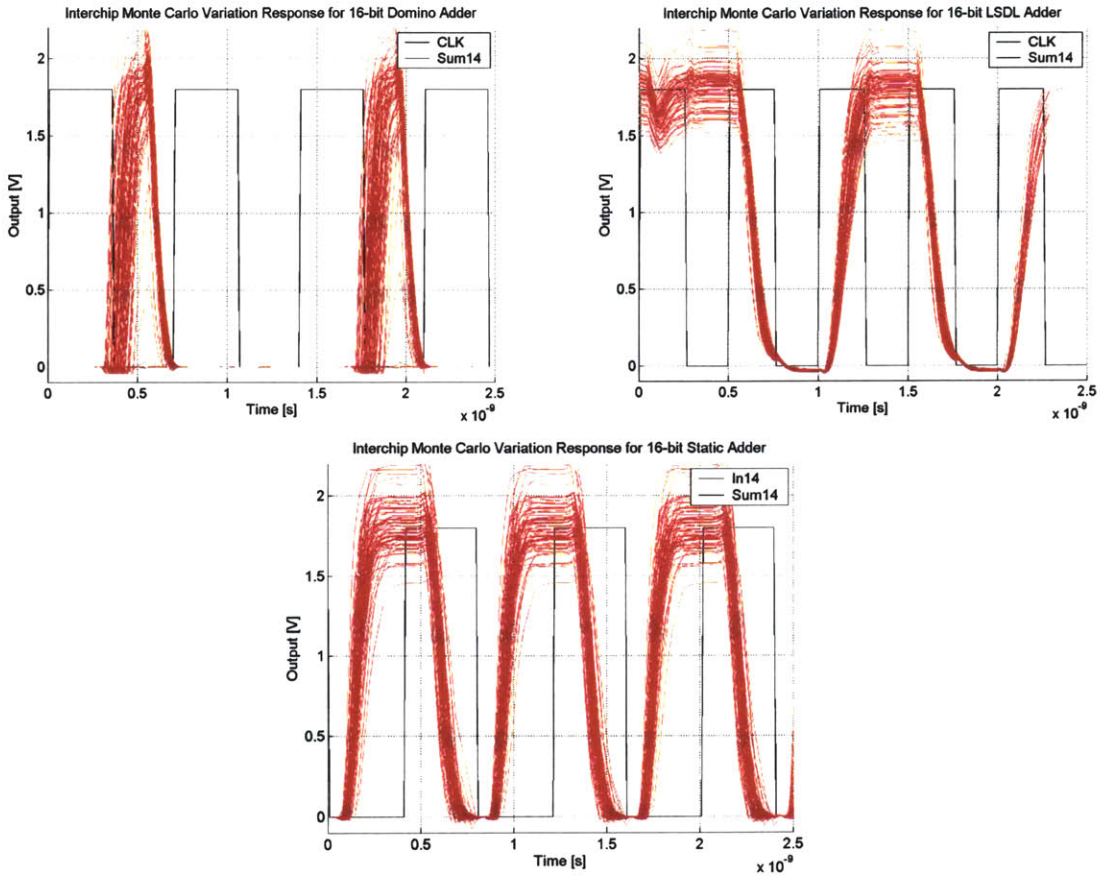

Fig. 6.3. Chip-to-chip Monte Carlo variation analysis for adder architectures.

6.4 Overall Variation

Overall variation, defined here as the combination of mismatch and chip-to-chip variations, established the simulated Gaussian random variables of n- and p-oxide thickness, active layer length, polysilicon channel length, threshold voltage, transistor

gate length and width, temperature, and supply voltage. Each of these variables (with the exception of gate length and width, which varied per transistor per run) changed per run based on a normalized Gaussian distribution with various three-sigma deviations, estimated from the historical MOSIS models. The results over two-hundred runs, can be seen in Figure 6.4.

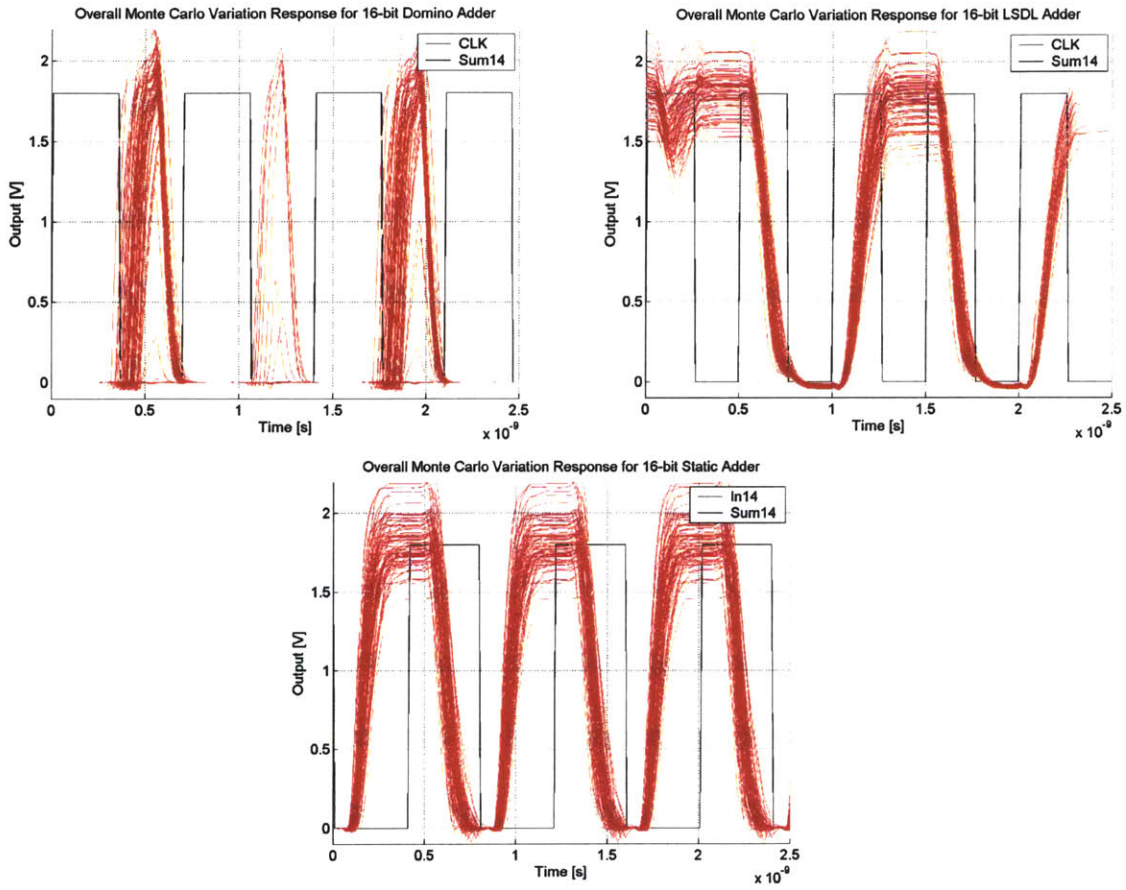

Fig. 6.4. Overall Monte Carlo variation analysis for adder architectures.

In the domino architecture, process variation-induced logical failures at high speeds were much more common than in either the LSDL or static case. The domino adder output false high signals, which can be seen near the middle of the time interval for the domino output in Figure 6.4. In the LSDL case, it is interesting to note how similar

the performance is to the traditionally robust static adder. Despite the barrage of process variation, LSDL manages to holds state well and surprisingly avoid logical failure, similar to the static adder. Static logic has traditionally been the benchmark for robustness and reliability in digital design, as we can see from the Monte Carlo analysis above. However, it can also be said that LSDL appears equivalently robust in this particular examination.

6.5 Sensitivity Regression Analysis of Monte Carlo Data

Output measurements were taken for Monte Carlo runs for the adders and were designated to be used for multiple regression and correlation analysis for the various designs. Adders were tested under the same conditions and speeds as in the previous sections. The measurements taken were delay, fall time, rise time, signal skew, total power, peak-to-peak of signal, average current, and signal width. These measurements can be identified in Figure 6.5. Power and average current should ideally follow one another linearly, though circuit operation may establish a notable deviation in average current versus the total current profile used to calculate power. The random input variables were n-active critical dimension, polysilicon critical dimension, n-oxide thickness, p-oxide thickness, n-threshold voltage, and p-threshold voltage. The goal was to calculate the relationship between the output and input variables through all adder designs and runs, and to correlate output measurements with one another.

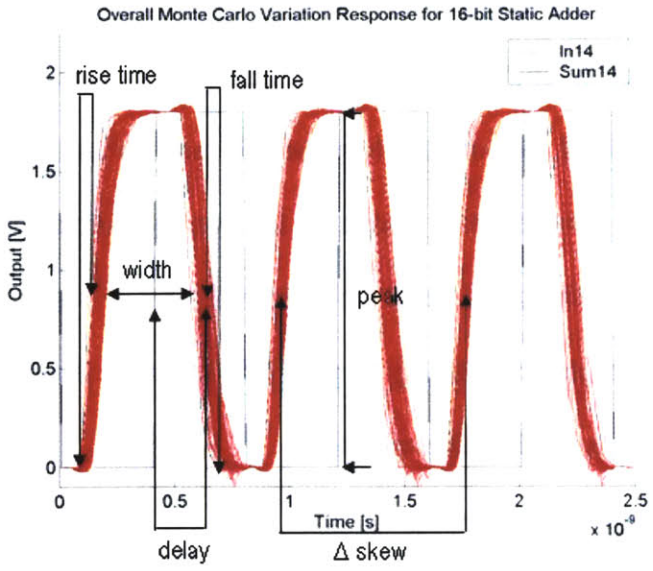

Fig. 6.5. Output measurements taken for Monte Carlo runs.

The sensitivities, as defined by Equation 6.2, to each variation source are shown in Figure 6.6. Looking closely at the figure, where the sign of the sensitivity indicates the correlation between the output and input metric, the domino adder is the most sensitive design to polysilicon critical dimension variation for the metrics measured. Output pulse width is especially more sensitive to poly CD variation in the domino adder than in the other designs. Such sensitivity may result in signal jitter during operation. As we showed in Section 5.1, jitter is a significant cause of failure in the domino style, which may explain its poor performance at higher speeds. The domino design's rise and fall times seem just as sensitive to the process variables tested as the LSDL and static designs. However, this adder appears relatively insensitive to many of the other process variables introduced, when compared to the LSDL and static adders. Because of its volatile behavior (more vulnerable to standard logical failures at high speeds, as seen in previous sections, yet relatively insensitive to process variables at moderate speeds) it can be said that the domino design lacks robustness across its full operating range.

The LSDL adder's peak voltage was generally more sensitive to the variables examined. Also, power and current are not matched in sensitivity. For simulation purposes, supply voltages are constant here, which means that the discrepancy must be caused by the current measurement technique for average current and power. In a true design, however, sensitivity in peak voltage will create supply voltage variation, compounding the problem. Sensitivity to poly CD variation is also a problem with the LSDL adder, just as it is for the other adders.

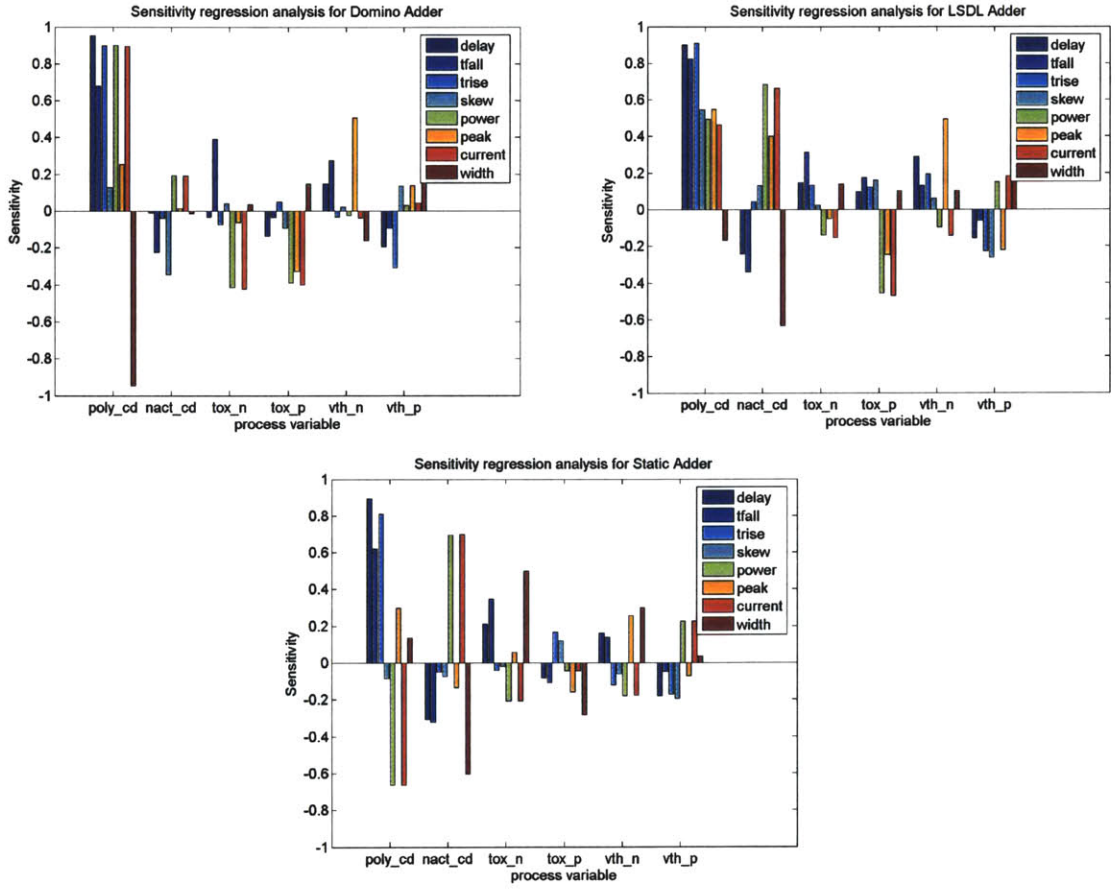

Fig. 6.6. Multiple regression sensitivity analysis for Monte Carlo adder simulations. Domino adder (upper left), LSDL adder (upper right), static adder (bottom).

The static design's peak voltage and skew metrics are relatively insensitive to the process variations examined, though output pulse width does appear to be sensitive. As mentioned before, this could introduce jitter and may be a cause for concern during operation. Overall, the static adder appears robust in performance as well as relatively insensitive to many process variables.

6.6 Adder Performance Comparisons

For purposes of side-by-side robustness comparisons for the three adder styles, boxplots were constructed to highlight the differences for each of the output metrics. The results can be seen in Figures 6.7.1 and 6.7.2. For each boxplot, the mean value is represented as a red line through the center of the box. The lower and upper quartiles of the values bound the box, the whiskers represent 150% of the interquartile range, and the red crosses represent outlying values beyond this range. The notches in each box represent an estimate of the uncertainty of the median; therefore tighter notches signify more robust data.

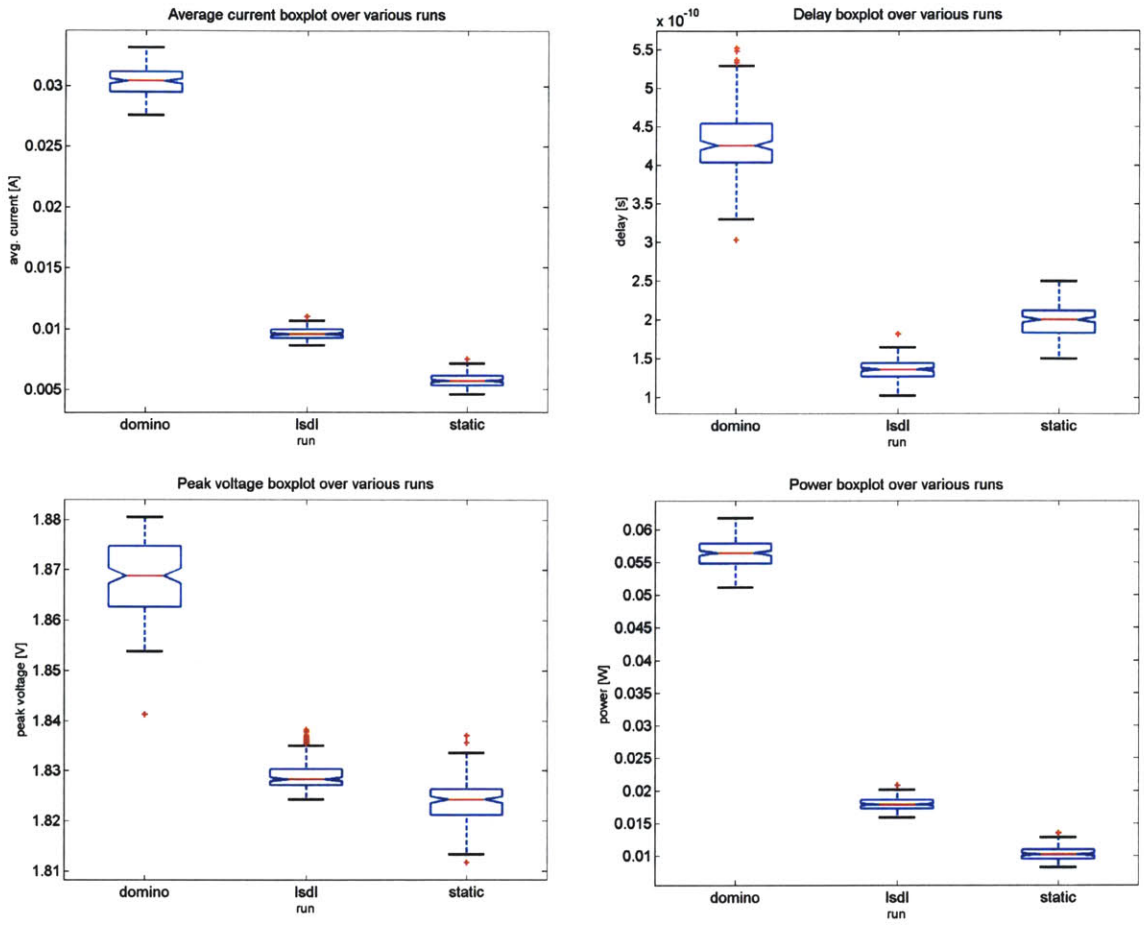

Fig. 6.7.1. Boxplots for adder output metrics. In order from uppermost left to bottommost right: current, delay, peak voltage, and power.

Many of these plots confirm ideas suggested in earlier MOSIS based tests. Based on these Monte Carlo runs, the domino adder consumes the most current, and with the static adder the least. The LSDL has the fastest delay response, lowest skew, relatively slow rise and fall times, and a strong output signal as represented by its long pulse width. However, its peak voltage is sensitive to process variations, as suggested by the sensitivity analysis above and confirmed by the presence of numerous outlying points from the whiskers in its boxplot. Overall though, the LSDL design is quite appealing from a performance perspective.

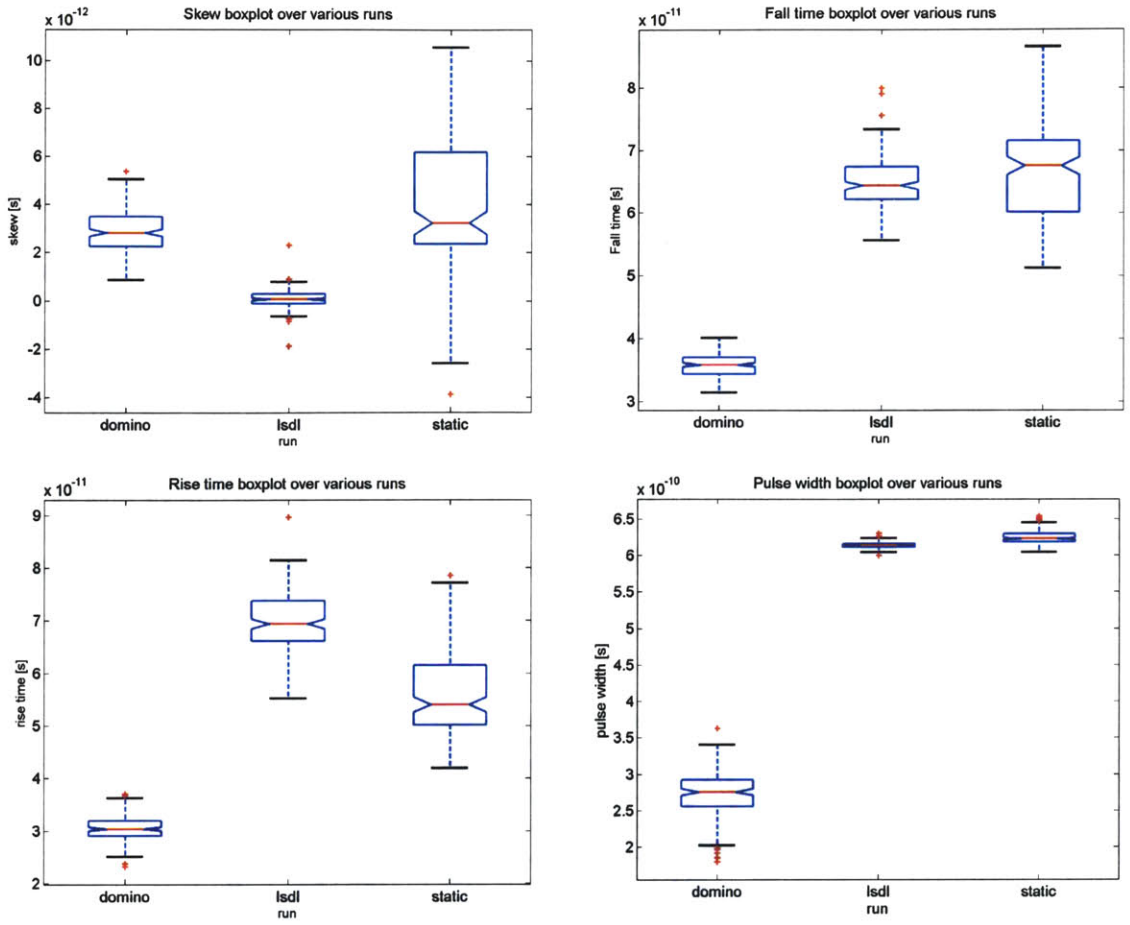

Fig. 6.7.2. Boxplots for adder output metrics. In order from uppermost left to bottommost right: skew, fall time, rise time, and pulse width.

6.7 Correlation Analysis

To complement the multiple regression analysis of the performance outputs versus process variable inputs, a group plot matrix was constructed to help correlate the relationship among the output metrics. For the domino adder, these results can be examined in Figure 6.8. For this adder, there exists a strong negative correlation between

pulse width and many of the other variables, especially delay. This is greatly attributed to domino's poor output signal strength.

From this analysis, and in comparison to other adders examined, the outputs of the domino adder were the most correlated in general when process variables are introduced. When the sum of the absolute values of the relative correlation matrices is used as a metric, the value for the domino style is 17.7% greater than the LSDL style and 15.5% greater than the static style. Such strong correlation among the output metrics may or may not be desirable, as it may be easier to create performance and robustness improvements when metrics all move together, though this may also create more sensitivity to process variations.

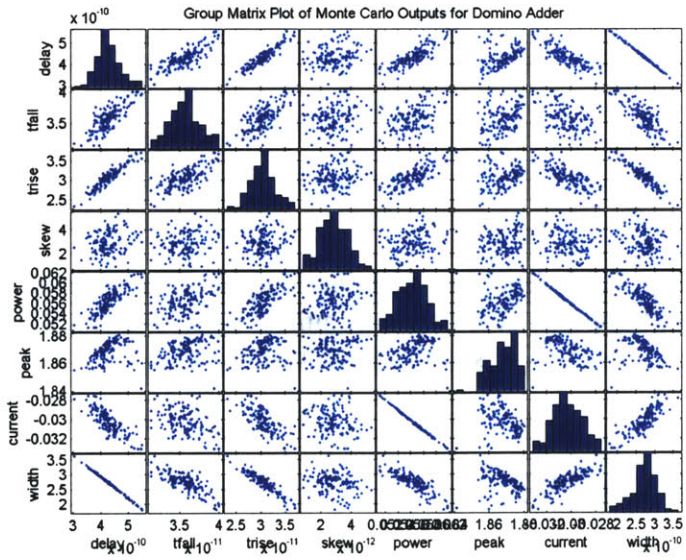

domino correlation matrix

	delay	t_fall	t_rise	skew	power	peak	current	width
delay	1	0.77752	0.93131	0.21318	0.82529	0.2925	-0.81146	-0.99818
t_fall	0.77752	1	0.68499	0.17525	0.41847	0.25012	-0.40207	-0.76154
t_rise	0.93131	0.68499	1	0.16838	0.74187	0.17234	-0.72944	-0.92456
skew	0.21318	0.17525	0.16838	1	0.19561	0.1734	-0.19175	-0.22032
power	0.82529	0.41847	0.74187	0.19561	1	0.31753	-0.99874	-0.82866
peak	0.2925	0.25012	0.17234	0.1734	0.31753	1	-0.31985	-0.29729
current	-0.81146	-0.40207	-0.72944	-0.19175	-0.99874	-0.31985	1	0.81467
width	-0.99818	-0.76154	-0.92456	-0.22032	-0.82866	-0.29729	0.81467	1

Fig. 6.8. Group plot matrix (above) and correlation matrix (below) for domino adder.

The correlation matrix for the LSDL design, in Figure 6.9, is weaker than its domino counterpart. Strong positive correlations exist among the delay, rise time, fall time, and skew metrics. Compared to the domino style, there is very little correlation between output width and other metrics, suggesting its signal strength and robustness. The LSDL adder stands up very well, considering it is being tested at a higher speed than both the domino and static adders (1.61 GHz vs. 0.91 GHz and 0.88 GHz, respectively). Power is also relatively uncorrelated to the other output metrics tested. The relatively low correlation suggests relative independence between output metrics. In other words, improvements may be made to one aspect of the design without as much risk of a tradeoff in another metric.

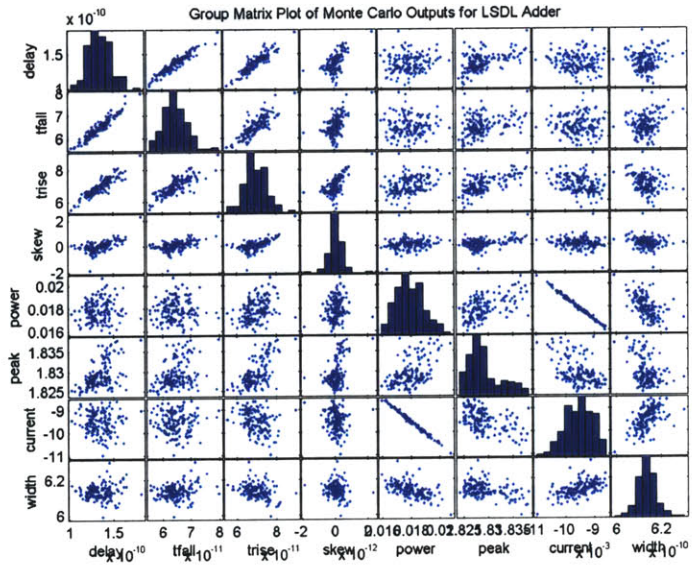

lsdl correlation matrix

	delay	t_fall	t_rise	skew	power	peak	current	width
delay	1	0.94352	0.90272	0.5637	0.14142	0.51311	-0.09743	0.076514
t_fall	0.94352	1	0.80868	0.51207	0.034685	0.3321	0.002862	0.22746
t_rise	0.90272	0.80868	1	0.61135	0.34609	0.61835	-0.29955	-0.32713
skew	0.5637	0.51207	0.61135	1	0.2131	0.40917	-0.17643	-0.18552
power	0.14142	0.034685	0.34609	0.2131	1	0.53526	-0.99504	-0.55014
peak	0.51311	0.3321	0.61835	0.40917	0.53526	1	-0.48617	-0.3932
current	-0.09743	0.002862	-0.29955	-0.17643	-0.99504	-0.48617	1	0.53845
width	0.076514	0.22746	-0.32713	-0.18552	-0.55014	-0.3932	0.53845	1

Fig. 6.9. Group plot matrix (above) and correlation matrix (below) for LSDL adder.

Figure 6.10 suggests that the static adder's power metric is negatively correlated to several time measurements, unlike the other adders. However, the current versus power discrepany is essentially nonexistent given its straightforward unclocked design. Overall, the outputs are only slightly more correlated with one another than in the LSDL adder.

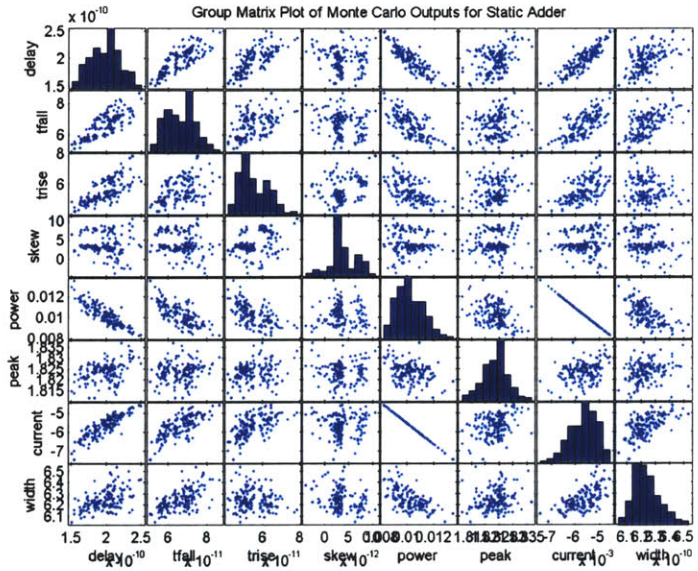

static correlation matrix

	delay	t_fall	t_rise	skew	power	peak	current	width
delay	1	0.76415	0.71181	-0.0651	-0.86288	0.36615	0.86255	0.489
t_fall	0.76415	1	0.42367	-0.2038	-0.69649	0.17471	0.69646	0.58146
t_rise	0.71181	0.42367	1	0.32104	-0.57459	0.22593	0.5748	-0.1446
skew	-0.0651	-0.2038	0.32104	1	-0.02749	0.029251	0.027025	-0.27579
power	-0.86288	-0.69649	-0.57459	-0.02749	1	-0.33258	-0.99998	-0.58242
peak	0.36615	0.17471	0.22593	0.029251	-0.33258	1	0.33297	0.21139
current	0.86255	0.69646	0.5748	0.027025	-0.99998	0.33297	1	0.58211
width	0.489	0.58146	-0.1446	-0.27579	-0.58242	0.21139	0.58211	1

Fig. 6.10. Group plot matrix (above) and correlation matrix (below) for static adder.

6.8 Results Summary

The Monte Carlo methodology introduced in this chapter has established a quantitative method to expose sensitivities in designs by using random distributions for input process variables and collecting key output measurements. From this analysis, the domino

adder's overall weaknesses are evident in performance, power, pulse width variation that may cause jitter, correlation of outputs, and poly critical dimension variation.

The LSDL adder, tested at a much higher frequency than the domino adder, exhibited relatively robust behavior from the Monte Carlo analysis, despite its slow rise and fall times and peak voltage sensitivity. It had the weakest correlation of the outputs measured, and was overall a robust design under variation conditions. From this analysis, the LSDL adder was much closer to the static adder in robustness than the domino.

The static adder's overall performance was the most robust with the lowest power, despite its slightly higher skew measurements. Also, it was tested at much lower speeds, which must be considered when compared against the faster dynamic designs. From the sensitivity analysis, the static design is generally the least sensitive to the process variations tested.

Chapter 7

Common Centroid-Based DFM

The centroid of a set of weighted points is analogous to a center of mass. Applied to circuit design layout, we can calculate and compare the centroids of groups of transistors to determine how balanced a design is. The idea is that a balanced-centroid design will be less vulnerable to process variation, since its blocks are positioned strategically to cancel out local variation gradients.

The common centroid methodology has long been used in analog integrated circuits, for which mismatch is a critical problem [26]. Much of the motivation for extending centroid analysis to the digital domain stems from analog circuits, where accurate behavior and biasing is essential for designs like differential pairs. This chapter addresses the use of centroid analysis in the digital domain, to determine the effectiveness of the method in canceling out trends in variation and improving design for manufacturability.

7.1 Common Centroid-Based Layout Methodology

A simple first order model of correlated chip-to-chip process variation in integrated circuits is a linear process variation gradient, visualized in Figure 7.1 and discussed in Section 5.5. Using this model, one powerful layout technique to improve manufacturability is the common centroid-based layout methodology.

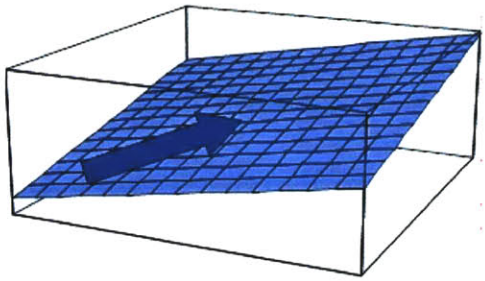

Fig. 7.1. Suggested linear process variation gradient (slanted plane effect).

To implement this methodology, we divide up the design into a grid of smaller cells, based on Equation 7.1, where S represents the underlying integrated circuit layout. To each section, we assign a pair of coordinates, which represents the center of that section, as well as a mass, which is calculated based on the area of the section relative to the mean area of all sections, as suggested in Equations 7.2 and 7.3 respectively. A mass of one indicates that the section is of average size. Once the design has been discretized into weighted coordinates, we may use a centroid matching technique for fast analysis of variation robustness.

$$S = \bigcup_{i=1}^{n} S_i \tag{7.1}$$

$$\{x_i = \frac{1}{S_{i_x}} \int_{S_{i_x}} S_i , y_i = \frac{1}{S_{i_y}} \int_{S_{i_y}} S_i \} \tag{7.2}$$

$$m_i = \frac{\int_{S_i} S}{\frac{1}{n} \sum_{i=1}^{n} \int_{S_i} S} \tag{7.3}$$

An integrated circuit can contain many different types of similar building blocks. For example, there may be several hundred similar NAND gates or multiplexers in a given design. We treat each block set (NANDs, NORs, etc) in a two-dimensional

integrated circuit layout as one of the subsections S_i described above by discretizing the locations of its n components into coordinates and calculating a corresponding mass for all combined instances. We may then compute the centroid of each block set as the weighted center among the instances, better visualized in Figure 7.2 and defined in Equation 7.4.

$$\{c_x, c_y\} = \{\sum_{i=1}^{n} \frac{x_i m_i}{n}, \sum_{i=1}^{n} \frac{y_i m_i}{n}\} \tag{7.4}$$

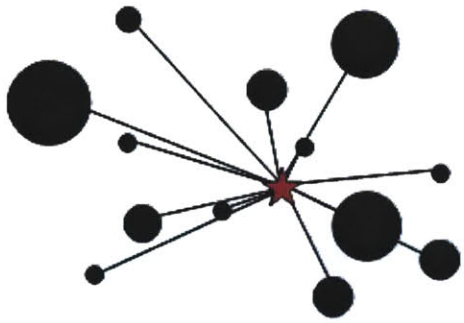

Fig. 7.2. Centroid of weighted masses.

Centroid matching among the various block sets may be a critical component in increasing design for manufacturability in integrated circuits. Given an estimated process variation gradient, we may evaluate the variation gradient at the location of the centroid to estimate the average variation under which the instances of that block set operate. This is suggested in Equation 7.5, where v_i represents the average variation for a block set, G is the process variation gradient, and (c_x, c_y) is the centroid of the set. We consider each block type as one entity with a single variation value, and we are interested in minimizing the deviation across all block sets. One way to minimize these intra-chip deviations is to match the centroids among the blocks, through careful design and layout.

$$v_i = G(c_x, c_y) \tag{7.5}$$

In Figure 7.3, two designs with two blocks sets (purple and green) are depicted. There are eight blocks in various locations for each set. If we assume equal sizes (weights) for all of the blocks, we can estimate the centroid for each block type. For the upper design in Figure 7.3, the centroids of both block sets (depicted as circles) are matched, since all green cells are placed with x- and y-symmetry and all purple cells are placed with x- and y-symmetry. In comparison, the asymmetry in block placement in the lower design style results in unmatched centroids. By ensuring matched centroids among block sets, we can be more confident that each block set will be affected in the same way if a process variation gradient is introduced.

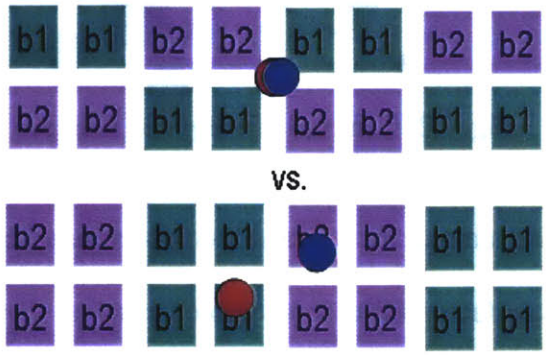

Fig. 7.3. Positioning of blocks affects centroid matching. Matched centroids (above) versus unmatched centroids (below).

7.2 Experiments with Regularity

Based on the centroid-based layout methodology described above, we can expect to minimize inter-block variation by employing a regular design and layout. In a perfectly regular design, there is only one block type, with uniformly sized cells. This suggests not only equal masses among cells, but also only a single centroid that needs no matching. As more block types are introduced, regularity deviates from the ideal case, with

unmatched centroids among blocks and unequal masses of cells. Though regularity has multiple levels of definition that may apply to integrated circuits, this common-centroid based methodology is a step forward in quantifying what it means to say a design is regular and using these characteristics to improve manufacturability.

To quantify the effects of centroid-based regularity on a design, a simulation experiment was devised. Three types of inverters were designed in the 0.18 micron process, a traditional two transistor inverter with a PMOS/NMOS W/L of 40/20, a four-transistor NAND gate with inputs tied together (inverter) of effective P/N W/L of 10/5, and a four-transistor NOR gate with inputs tied together (inverter) of effective P/N W/L of 10/5. Sixteen traditional inverters were chained together to form a standard inverter cell, 16 NAND inverters were chained together to form a NAND-inverter cell, and 16 NOR inverters were chained together to create a NOR-inverter cell.

The three cell types were grouped in three different patterns, depicted in Figure 7.4, to form a 576-inverter chain. The standard, NAND, and NOR inverter cells are respectively yellow, purple, and green in the figure. The output signal path rises up the rows in the first column and continues to the bottom of the second column, repeating this pattern going across. The final output is measured from the last inverter in the upper right cell of the design. The three patterns are used to isolate regularity based performance. Theoretically, all three styles should create the same output, considering each is composed of the same type and number of logic inverters. Layout and component grouping, however, should prove otherwise.

In the first pattern (left), like cells are grouped together for logical regularity and common-centroid based matching. The centroids of all three block types (yellow, teal,

and purple) match in the center of the design. We will call this a very regular design. In the second pattern (center), logical regularity is reduced, though the blocks still keep their matched common centroids in the center of the design. We will call this a regular design. For the final design (right), neither common centroid matching nor logical consistency is employed in this random design. This will be referred to as a non-regular design. Using the Monte Carlo methodology employed in Chapter 6 for sensitivity analysis, we can explore the intricate interplay between regularity and variation in these designs.

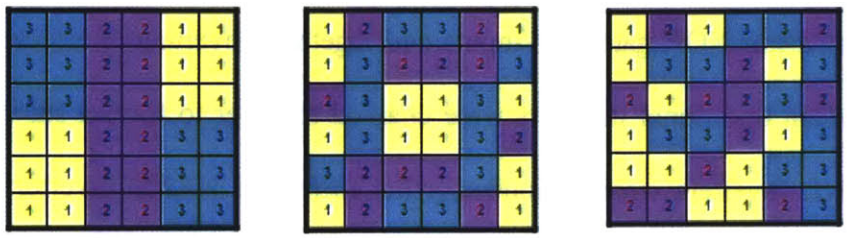

Fig. 7.4. Common centroid-based layout styles. Very regular (left), regular (center), and non-regular (right).

7.3 Tests and Data

The 576-inverter chains were instantiated as layouts as specified in the previous section, in three styles – very regular, regular, and non-regular. The layout for the very regular pattern is illustrated in Figure 7.5. Using the Monte Carlo methodology developed previously in this research, all three styles were simulated under the same range and type of process variables to which the adders in Chapter 6 were exposed. Similar analysis was collected from these layouts and formatted for consistency with previous presentation styles.

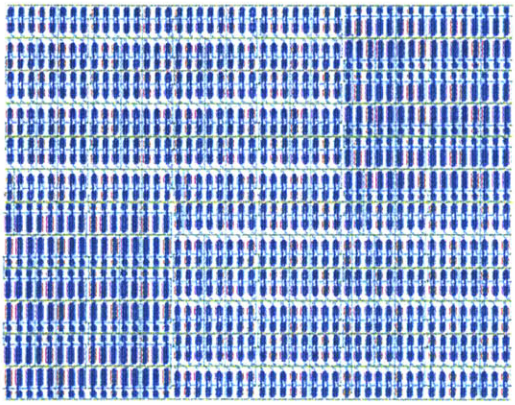

Fig. 7.5. Layout of very regular inverter chain, with grouped cells and matching centroids among blocks.

The slanted plane was used in conjunction with the Monte Carlo variation analysis to simulate conditions where centroid matching would be critical. If the centroid matching theory were correct, we would expect to see decreased performance and increased sensitivity when a slanted plane is applied to a non-regular layout style, for example. A square wave input was applied to the first inverter and metrics were taken at the output of the 576^{th} inverter. This was done over 200 Monte Carlo runs with simultaneously sampled process variables of polysilicon CD, n- and p- threshold voltage, n- and p- oxide thickness, and n-active CD. The first set of data can be visualized in Figures 7.6.1 and 7.6.2, where boxplots for current, delay, peak voltage, power, skew, fall time, rise time, and pulse width have been created.

The boxplots are classified into types, indicating the layout style, and also by the direction of the slanted planes. Type 1 is the very regular design in Figure 7.4, type 2 is the regular design, and type 3 is the non-regular style. For simplicity, two slanted planes for each type were applied, one with strength into the northeast (NE) direction and one with strength into the southwest (SW) corner. For the slanted plane, threshold voltage

variation was increased on a gradient from 0 to 50 mV across the design and temperature variation was applied from 0 to 20°C. A negative gradient (0 to -50 mV, for example) was not applied to ensure that the simulation actually ran with HSPICE.

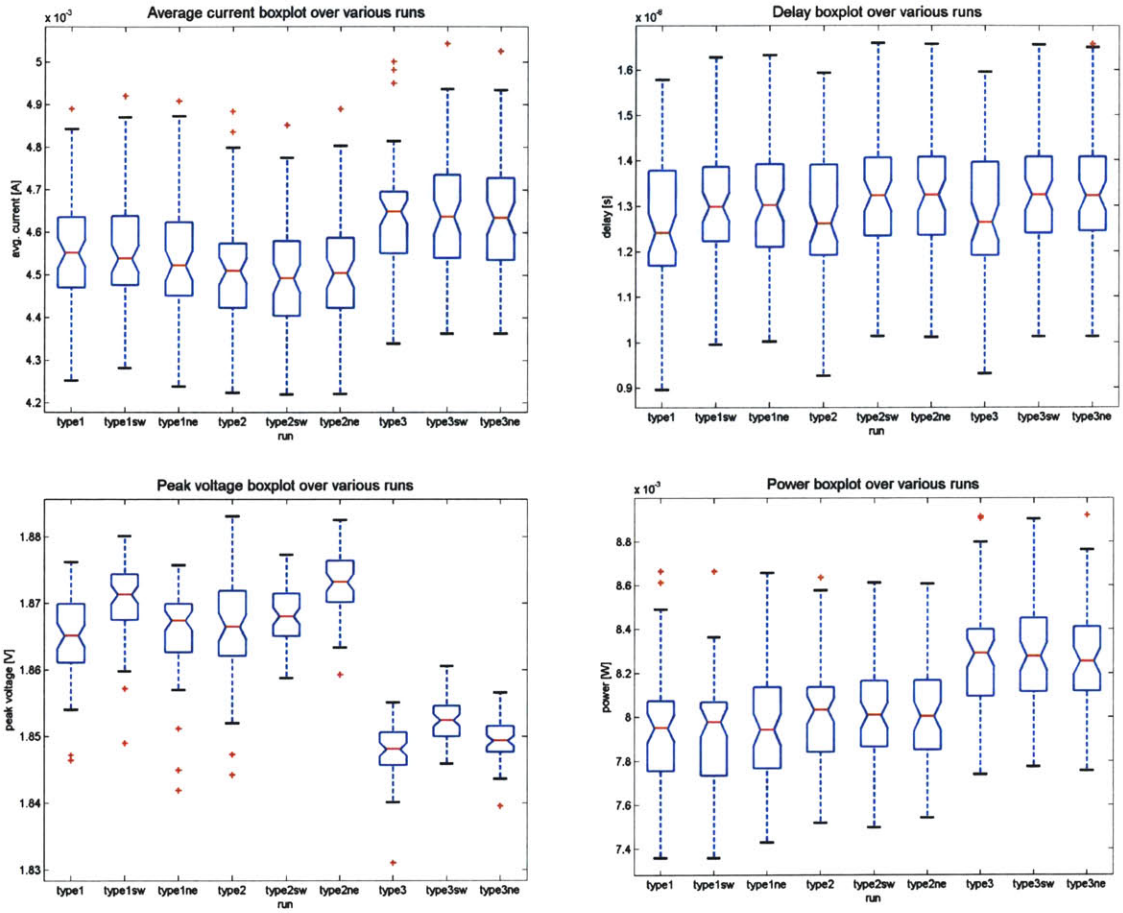

Fig. 7.6.1. Boxplots for centroid-based layout style output metrics. Grouped into types (1 = very regular, 2 = regular, 3 = non-regular). Among types, tests without slanted plane, with slanted plane into SW direction, and also into NE direction. In order from uppermost left to bottommost right: current, delay, peak voltage, and power.

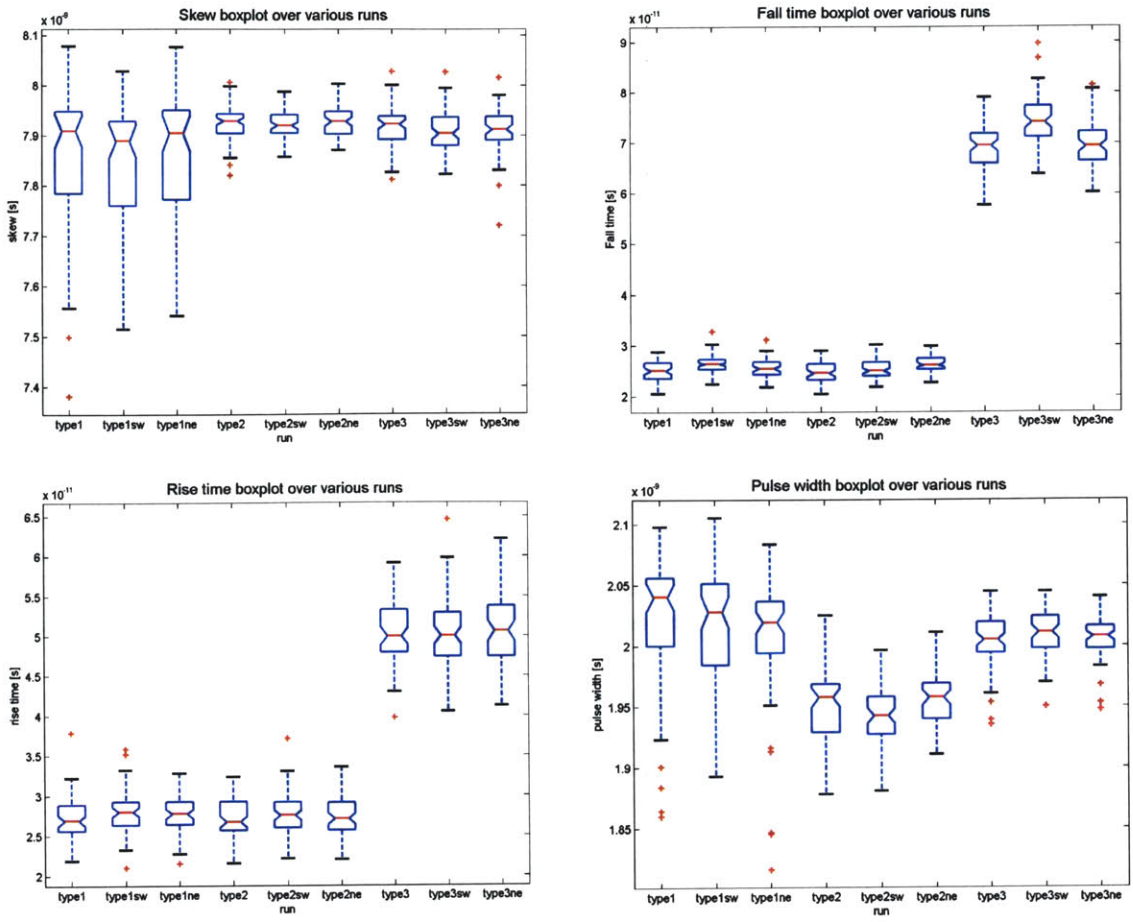

Fig. 7.6.2. Boxplots for centroid-based layout style output metrics. Grouped into types (1 = very regular, 2 = regular, 3 = non-regular). Among types, tests without slanted plane, with slanted plane into SW direction, and also into NE direction. In order from uppermost left to bottommost right: skew, fall time, rise time, and pulse width.

Many interesting trends can be extracted from this data. First, most of the performance differences are between types (very regular vs. non-regular for example), but not so much within each type (slanted plane vs. no slanted plane). We see a general degradation in performance from the type 1 (very regular) to type 3 (non-regular) styles in many instances, including peak voltage, fall time, rise time, power, average current, and pulse width. Nonetheless, the three design types maintain roughly the same delay, so

a novice designer may never notice the subtle imperfections these varying design styles induce.

For our process, we model the standard simulation as a deterministic system, given by Equation 7.6, where a standard set of inputs yields a single output. Once we introduce random variables to the inputs, the system can be modeled as in Equation 7.7, with incremental changes added to the inputs and outputs. We use our knowledge of the nominal performance numbers and mean deviation in performance ($\delta_{\bar{y}}$) to gauge how close the Monte Carlo simulations are to nominal behavior. We define the relative deviation, ΔV, as shown in Equation 7.8. The results, in Table 7.1, suggest that the Monte Carlo technique is close to nominal for all design styles, with the very regular style generally deviating the least.

$$\bar{x}_0 \rightarrow f(\bar{x}) = \bar{y}_0 \tag{7.6}$$

$$\bar{x}_0 \pm \delta_{\bar{x}} \rightarrow \bar{y}_0 + \delta_{\bar{y}} \tag{7.7}$$

$$\Delta V = \frac{\delta_{\bar{y}}}{\bar{y}_0} \tag{7.8}$$

	power	delay	rise time	fall time	skew	width	v_peak	current
type1 (very regular)	0.785	-0.317	-0.485	-0.234	-0.302	3.734	0.070	-0.316
type2 (regular)	0.144	-0.522	-5.418	-4.800	0.497	1.351	-0.062	0.324
type3 (non-regular)	0.170	-0.656	-1.863	-0.993	-0.196	0.044	-0.099	0.335

Table 7.1. Percentage deviation ($\Delta V \cdot 100\%$) of mean Monte Carlo runs over nominal run.

When we introduce a slanted plane effect, or additional spatial process variation gradient to the design, the results also degrade, visible in such metrics as pulse width, power, peak voltage, and delay within each design type. The very regular design style, simulated with a slanted plane effect, still outperforms the non-regular design without any slanted plane effect on several levels based on this data. This preliminary data has demonstrated that centroid-based matching in digital integrated circuits is worthy of further investigation.

Robustness to sensitivity is another factor that was quantified in this centroid-based design analysis. Using the same regression engine as in Chapter 6, we find that the irregular designs created under the common centroid methodology do indeed seem to be more sensitive to the process variable distributions introduced. In Figures 7.7.1 through 7.7.3, the sensitivity of certain output metrics to process variables is quantified. Figure 7.7.1 measures the sensitivity of the designs using the same Monte Carlo technique as in Section 6.5. As before, the random interchip process variables used were oxide thickness, active layer length, polysilicon channel length, and threshold voltage. Figures 7.7.2 and 7.7.3 illustrate the design sensitivities when a slanted plane gradient is introduced, in addition to the random process variables just described.

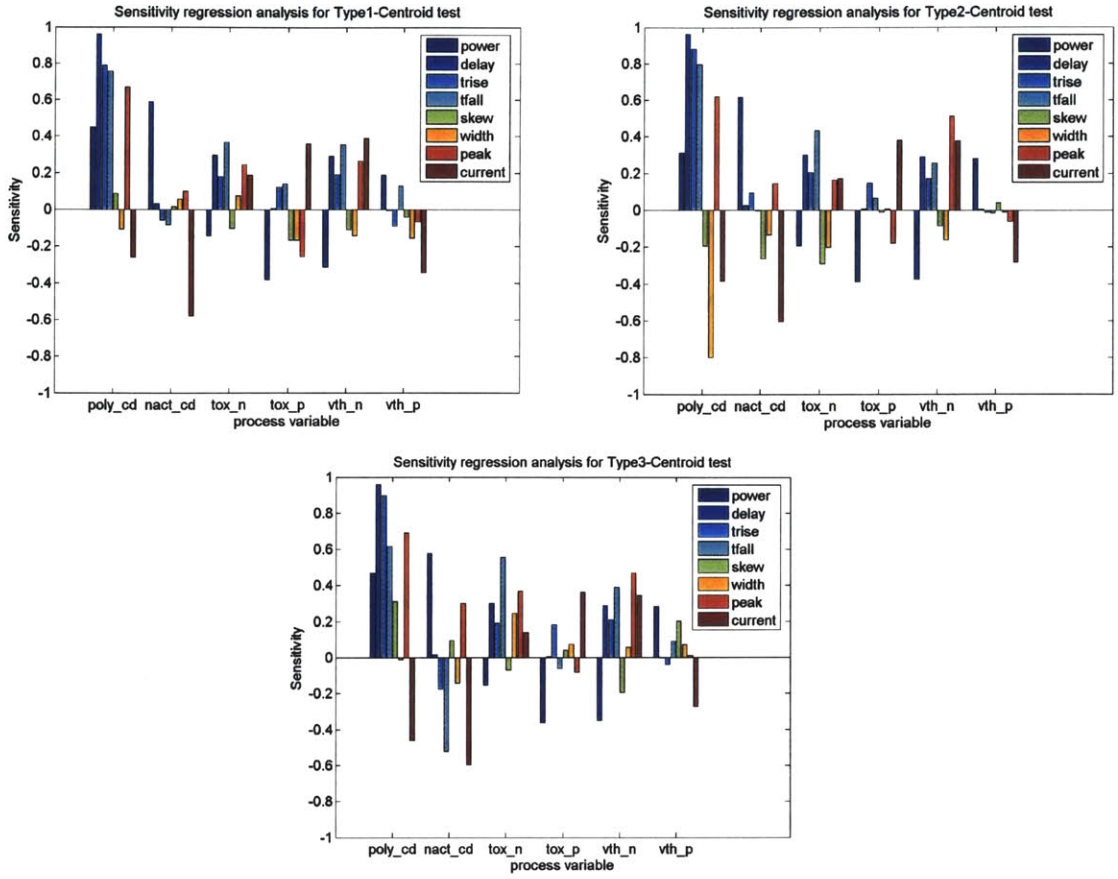

Fig. 7.7.1. Sensitivity tests for centroid-based inverter chain layouts, with no slanted plane. Type 1-very regular (upper left), type 2-regular (upper right), and type 3-non-regular (bottom).

The overall increase in sensitivity is subtle though apparent when traveling from the type 1 (very regular) style to the type 3 (non-regular) style, with or without slanted planes. The figures suggest, for example, that the type 2 (regular) layout pattern exhibits more sensitivity to pulse width and the type 3 (non-regular) style exhibits slightly more sensitive rise and fall times.

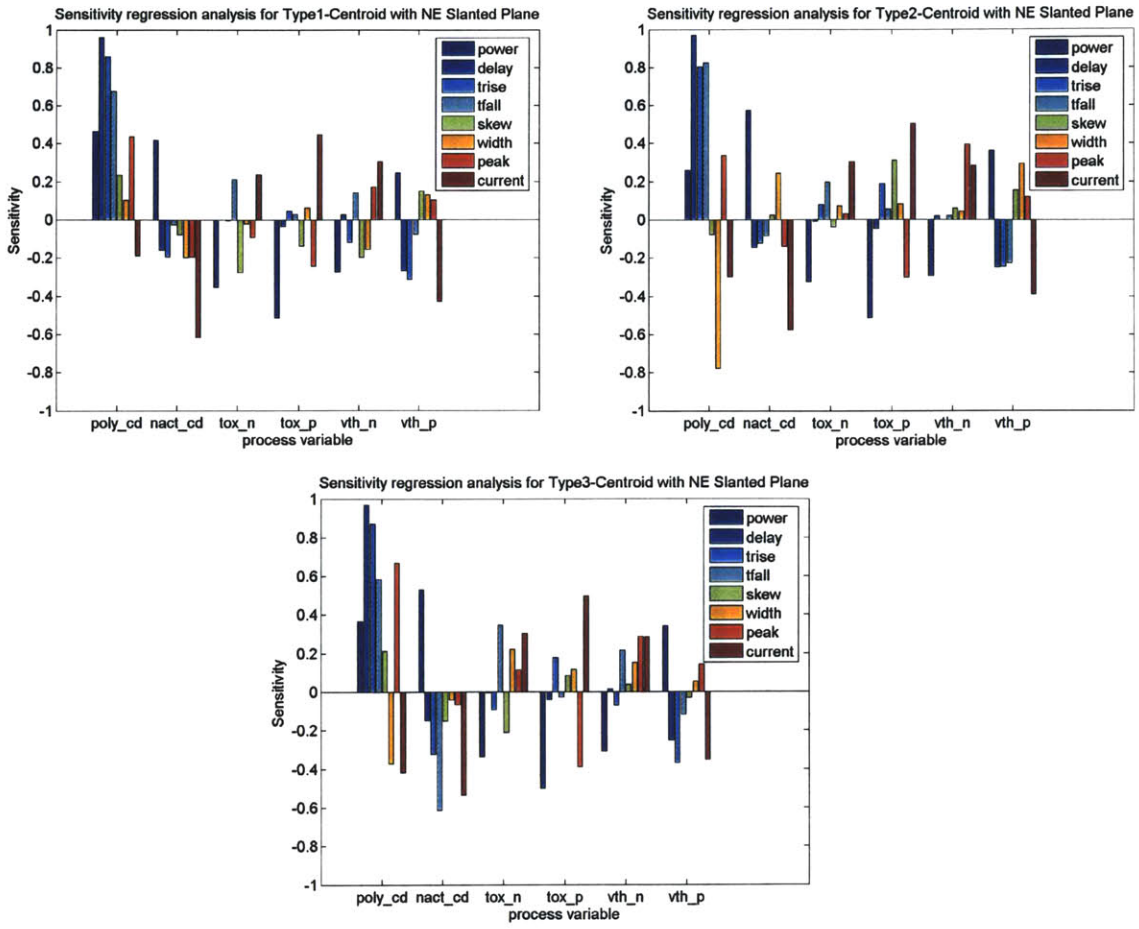

Fig. 7.7.2. Sensitivity tests for centroid-based inverter chain layouts, with northeast slanted plane. Type 1-very regular (upper left), type 2-regular (upper right), and type 3-non-regular (bottom).

Introducing a slanted plane increased overall sensitivity, though as suggested before, the type 1 style with a slanted plane still is less sensitive to the process variables that a type 3 style with no slanted plane. This strongly suggests that regularity has associated robustness characteristics in this design, considering all designs are composed of the same number and size of building blocks and have relatively similar delays.

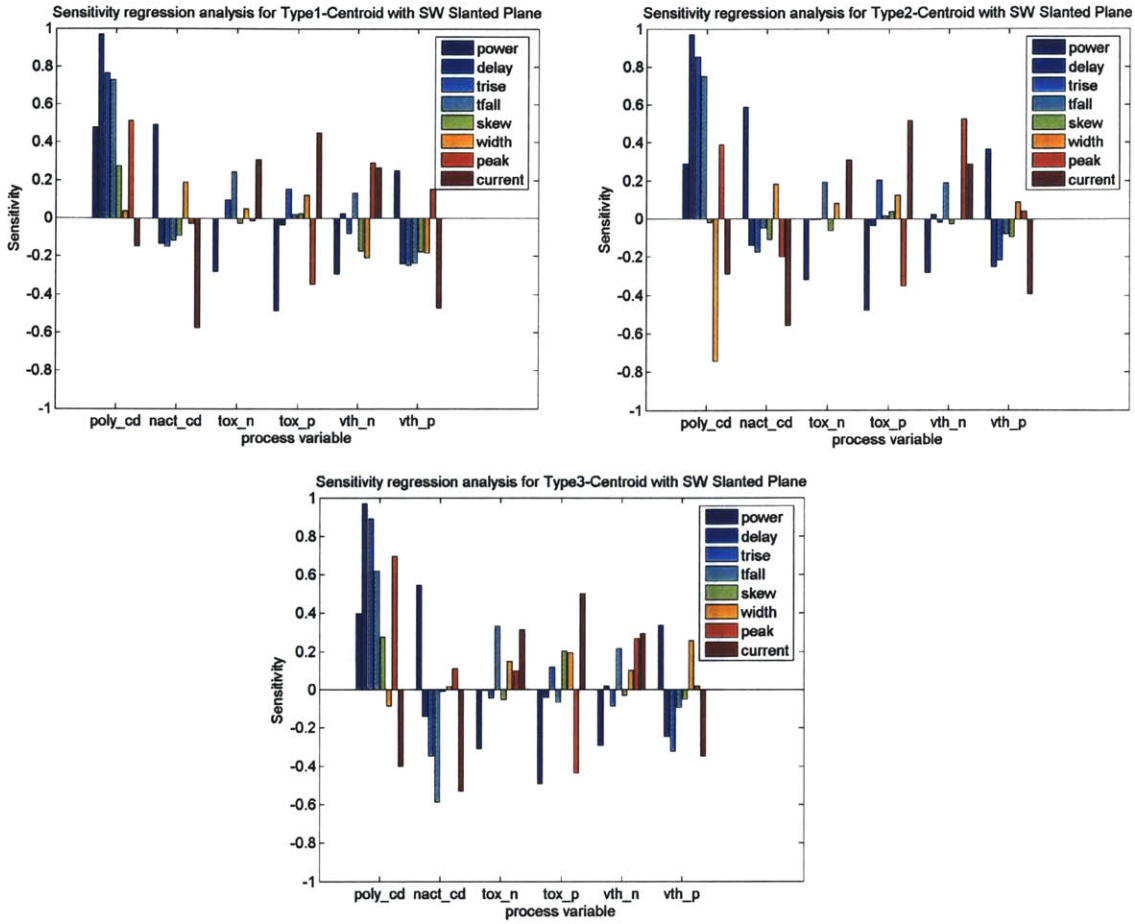

Fig. 7.7.3. Sensitivity tests for centroid-based inverter chain layouts, with southwest slanted plane. Type 1-very regular (upper left), type 2-regular (upper right), and type 3-non-regular (bottom).

7.4 Results Summary

Centroid-based analysis for digital integrated circuits may be worthy of closer examination. We assumed an underlying localized linear process variation gradient and created three inverter chain designs with varying degrees of common-centroid matching among block types. Using the Monte Carlo methodology developed before, we charted the trends of the varying centroid designs. The very regular, regular, and non-regular designs each had notable performance differences.

Simulated slanted plane effects had relatively negligible effects in comparison, suggesting that the underlying regularity of the design is much more significant in improving design robustness. Considering the manufacturing process variation gradient during design and layout and using a common-centroid style technique to increase layout regularity may enhance circuit performance and robustness, as illustrated in these simulations.

Chapter 8

Conclusions and Future Work

This work has built on two key concepts in the design of digital integrated circuits: that some designs actually induce less process variation and that some designs are robust to the remaining process or environmental random variations. This research has examined both ideas, through analysis of system and circuit level designs. Methodologies were developed to verify the design impact on variation, including a sensitivity analysis and a layout model to reduce spatial variation.

8.1 Conclusions

Based on this research, regular structures appear to have notable manufacturability improvements over their non-regular counterparts. The layout based tests illustrated the reduction of pattern density dependency variation in CMP, plating, dishing, and erosion for regular (VPGA, FPGA) versus more traditional designs (Berkeley decoder). The use of dummy fill post-layout, however, significantly balances many of the process manufacturability advantages that these regular designs carry.

Circuit level regularity was explored as a method for increasing robustness to process and environmental variation. In comparing domino, LSDL, and static circuit families, we have learned that LSDL, with its minimal sizing scheme and high

performance, is a more reliable dynamic circuit architecture than traditional domino logic, though it is vulnerable to clock timing issues at its high speeds. However, reduced density variation and comparable process variable sensitivity of an LSDL adder versus a comparable static adder indicate that LSDL may rival static logic in terms of manufacturability and robustness.

A methodology was established to run designs against historical MOSIS models and has proven to be an effective technique for system level variation measurements across a design. This methodology evolved into a more comprehensive Monte Carlo engine to input process variables with their respective probability distributions for black box design testing. By recording various output metrics and using sensitivity analysis, we developed a way to expose vulnerabilities of designs. Such sensitivity analysis was used on the domino, LSDL, and static adders, and also later for the common-centroid based inverter designs. Robustness to process variation was greatest for the static architecture, though the LSDL style operated considerably well given its high performance.

The common centroid-based layout methodology traditionally used for analog mismatch avoidance was explored for digital integrated circuits. Testing the theory that more regularity ensures better matching of the numerous centroids of the different block types in a design, three sets of 576-inverter chains were designed. Each of these designs had a different degree of regularity and matching of centroids. Results showed that the more regular the blocks of inverter chain were and the better the centroids were matched, the better the design performance was. After introducing the slanted plane effect to all three designs, the more regular designs with balanced centroids again performed better.

Regularity is an important design consideration in the next era of digital integrated circuit manufacturing. Circuit design for manufacturability, balanced centroid matching of layouts, pattern density uniformity through regular blocks, and improved dummy fill techniques at edges are the trends that this research encourages in technology generations to come.

8.2 Future Work

This research has uncovered a plethora of topics for future research that deserve examination. The first includes edge uniformity for large scale digital designs. Because pattern densities near pads and corners of a design vary tremendously with the densities on a die itself, a substantial density gradient is established that is not corrected with dummy fill. For high speed designs, like the Berkeley decoder examined, an analysis on the impact of pad and edge non-uniformity on throughput may prove interesting. Dummy fill may need to be improved or more robust edge or packaging design may need to be promoted if the impact is significant.

Regular fabrics show potential for manufacturability, though must be able to differentiate themselves among ASICs with advanced dummy fill algorithms, for example, which we showed to likely exhibit comparable pattern density dependency effects. One suggestion is the inclusion of a regular logic family, like LSDL, into the logic blocks of a regular fabric. Robustness of blocks, better performance, and tighter pattern density ranges may result, and may improve the commercial viability of such projects.

The Monte Carlo engine developed for this research has further potential for refinement. More process variables, output metrics, and confidence intervals are just a few of the elements that could be added to deliver more comprehensive results. One element that limited the analysis done for this research was the physical simulation time for each analog run. Adding variation functionality, with process variable inputs and probability distributions, to commonly used digital simulators like NanoSim, would be a tremendous advancement in integrated circuit computer aided design.

Finally, the common centroid layout methodology may be a significant opportunity for exploration, based on the promising results in this research. For example, one may devise a methodology for discretizing a layout by the block types, into a network of nodes and nets, based on a complex model of dependencies, like proximity, block size, etc., to establish the mass of each of the blocks. Calculating the centroids of each of the block types may then be a powerful predictor of the potential magnitude of cross-block variation, given the process variation design gradient. This analysis may be used to optimize yield, for example, across designs generations. Ultimately, this knowledge may be integrated into place and route algorithms and design rules for enhanced yield.

References

[1] M. Rieger, J. Mayhew, and S. Panchapakesan, "Layout Design Methodologies for Sub-Wavelength Manufacturing." *Proceedings of the 38th ACM/IEEE Design Automation Conference*, pp. 85-88, June 2001.

[2] T. Park, "Characterization and Modeling of Pattern Dependencies in Copper Interconnects for Integrated Circuits." Ph.D. Thesis, MIT Dept. of EECS, May 2002.

[3] K. M. Gonzalez-Valentin, "Extraction of Variation Sources Due to Layout Practices." S.M. Thesis, MIT Dept. of EECS, May 2002.

[4] W. Maly, H. Heineken, J. Khare, and P. K. Nag, "Design for Manufacturability in the Submicron Domain." *Proceedings of the 1996 IEEE/ACM International Conference on Computer-Aided Design*, pp. 690-697, Nov. 1996.

[5] L. T. Pileggi, H. Schmit, A. J. Strojwas, P. Gopalakrishnan, V. Kheterpal, A. Koorapaty, C. Patel, V. Rovner, K. Y. Tong, "Exploring Regular Fabrics to Optimize the Performance-Cost Trade-Off." *Proceedings of the 40th ACM/IEEE Design Automation Conference*, pp. 782-787, June 2003.

[6] D. S. Boning, T. Maung, J. Chung, K.-J. Chang, S.-Y. Oh, and D. Bartelink, "Statistical Metrology of Interlevel Dielectric Thickness Variation." *Proceedings of the SPIE Symposium on Microelectronics Manufacturing,* SPIE Vol. 2334, pp. 316-327, Austin, TX, Oct. 1994.

[7] B. Stine, D. Boning, J. Chung, D. Bell, and E. Equi, "Inter- and Intra-die Polysilicon Critical Dimension Variation," *Manufacturing Yield, Reliability, and Failure Analysis session, SPIE 1996 Symposium on Microelectronic Manufacturing,* SPIE Vol. 2874, Austin TX, Oct. 1996.

[8] F. Honore and A. Chandrakasan. "Low-Power FPGA Design." MIT MTL MARC Conference, Jan. 2005.

[9] A. Strojwas, "Design for Manufacturability and Yield." *Proceedings of the 26th ACM/IEEE Design Automation Conference*, pp. 454-459, June 1989.

[10] R. Montoye, W. Belluomini, H. Ngo, C. McDowell, J. Sawada, T. Nguyen, B. Veraa, J. Wagoner, M. Lee, "A Double Precision Floating-Point Multiplier." *International Solid-State Circuits Conference,* 2003.

[11] W. Maly, "IC Design in High-Cost Nanometer-Technologies Era." *Proceedings of the 38th ACM/IEEE Conference on Design Automation*, pp. 9-14, June 2001.

[12] S. Borkar, "Design Challenges of Technology Scaling." *IEEE Micro* Vol. 19, Issue 4, pp. 23-29, 1999.

[13] MOSIS Test Results for TSMC 0.18 Micron Runs (CL018, CM018). http://www.mosis.org/products/fab/vendors/tsmc/tsmc018/

[14] E. Yeo and B. Nikolic. "High Throughput VLSI Architectures for Iterative Decoders," Ph.D. Thesis, UC Berkeley Dept. of EECS, 2003.

[15] J. Rabaey, A. Chandrakasan, and B. Nikolic, *Digital Integrated Circuits—A Design Perspective*, Englewood Cliffs, NJ: Prentice Hall 2003.

[16] N. Drego, "A Low-Skew, Low-Jitter Receiver Circuit for On-Chip Optical Clock Distribution." S.M. Thesis, MIT Dept. of EECS, 2003.

[17] S. Kao. DCDG Progress Report. UC Berkeley, 2003.

[18] A. Saltelli, K. Chan, and E. M. Scott, *Sensitivity Analysis*, Chichester, NY: Wiley 2000.

[19] H. Cai, T. Park, D. Boning, Y. Kang, J. Lee, S. K. Kim, and H. Kim, "Coherent Chip-Scale Modeling for Copper CMP Pattern Dependence," Paper K2.4, *Chemical-Mechanical Polishing Symposium, MRS Spring Meeting*, San Francisco, April 2004.

[20] D. Boning, S. Nassif, "Models of Process Variations in Device and Interconnect," in *Design for High Performance Microprocessor Circuits*, Eds. A. Chandrakasan, W. Bowhill, and F. Fox. Piscataway, NJ: IEEE Press, 2000.

[21] S. Nassif, "Modeling and Forecasting of Manufacturing Variations," *Proceedings of the 2001 Asia and South Pacific Design Automation Conference*, pp. 145-149, Jan. 30-Feb. 1, 2001.

[22] A. Koorapaty, V. Kheterpal, P. Gopalakrishnan, M. Fu, and L. Pileggi, "Exploring Logic Block Granularity for Regular Fabrics," *Proceedings of the 2004 Design, Automation and Test in Europe Conference,* Vol. 1, pp. 468-473, Feb. 2004.

[23] K. A. Bowman, S. G. Duvall, and J.D. Meindl, "Impact of Die-to-Die and Within-Die Parameter Fluctuations on the Maximum Clock Frequency Distribution for Gigascale Integration," *IEEE Journal of Solid-State Circuits*, Vol. 37, Issue 2, pp. 183-190, Feb. 2002.

[24] C. Chang, K. Kuusilinna, B. Richards, A. Chen, N. Chan, R. W. Brodersen, B. Nikolić, "Rapid Design and Analysis of Communication Systems Using the BEE Hardware Emulation Environment," *Proceedings of the 14th IEEE International Workshop on Rapid System Prototyping*, pp. 148-154, June 2003.

[25] H. Ngo, W. Belluomini, and R. Montoye, "Circuits and Systems for Limited Switch Dynamic Logic." U.S. Patent No. 6,650,145 B2.

[26] E. Sackinger and L. Fornera, "On the Placement of Critical Devices in Analog Integrated Circuits," *IEEE Transactions on Circuits and Systems*, Vol. 37, Issue 8, pp. 1052-1057, Aug. 1990.