Using Test Data to Improve IC Quality and Yield

Anne Gattiker
IBM Research
Austin Research Lab
Austin, TX, USA
gattiker@us.ibm.com

***Abstract*—The complexity of interactions in today's manufacturing processes makes test structures and experiments inadequate as sole drivers of yield-learning and design-for-manufacturing [DfM]. They must be driven by product impact. Product-impact-oriented test-based learning provides insight into the nature of model-hardware mismatches and variability that exist on and impact real products. That insight can be used to drive both parametric and defect-oriented process actions and DfM.**

I. Introduction

To achieve good quality and yield, good design-for-manufacturing [DfM] must be applied, but to achieve good DfM, knowledge is needed of what to do and what to avoid doing in design. To achieve that knowledge, we need to know the weaknesses in our manufacturing processes – and, importantly, the weaknesses in our manufacturing processes that *impact* the final product. We need to consider both parametric and defect-related process weaknesses.

Parametric processes weaknesses are those that cause mismatches between the manufactured circuit and designer intent, e.g., because physical shapes manifest dimensions different from their intended width and length, or because electrical parameters, such as threshold voltage, Vt, or interconnect resistivity do not match their expected values. One approach to identifying parametric weaknesses is to create test structures aimed at measuring these effects. The structures can be in the scribe line or on test chips. Doing so is challenging, however, because experiments are inevitably limited in the number of physical and electrical configurations they can cover. Instance-specific effects based on, for example, neighborhood/density and stress-related geometries, as well as operating condition effects, such as history and self-heating, make it very difficult to cover all relevant cases. In addition, today's processes are complex enough that not few but many parameters may be out of their specified range, or perhaps not even covered by a specification, and those specifications that do exist may not correlate well to true requirements for products successfully meeting their power-performance objections. As a result, assessment of product impact is needed to identify and prioritize those weaknesses that truly need to be rectified in manufacturing or compensated for in design. Test, e.g., at wafer and final package levels, including operating-speed and power-draw measurements, assesses product power and performance directly and therefore can be a key source of guidance for directing corrective actions and DfM efforts. Here the term DfM is used to mean action taken at any step in the design process to ensure the quality or yield of the outgoing product. It includes both physical and electrical design.

The shortcomings of parametric test structures and the need to drive process actions and DfM by product impact makes product-based learning attractive. However learning based on test results provides challenges as well. One reason is that test results are coarse. Typically at most a few large circuit blocks are independently testable. In addition, even speed-measuring test results typically return just a maximum operating frequency without information on which specific path failed. To a diagnostic engineer trying to determine the reason for a failing parametric test, neither the tools comprising the design flow nor scribe line FETs, likely to be distant from the failing block, are very helpful. The design flow loses a great deal of information through simplifications that are compensated for by guardbanding. The design flow's orientation toward bounding circuit behavior, rather than predicting it, limits its use in understanding product failure. For example, it is not possible to adjust FET parameters and produce a meaningful new static timing report. In addition, the design flow tools are not necessarily prepared to handle the types of process imperfections that really impact the product. Examples of such variations, including mistracking of diverse circuit types and systematic across reticle variation, are highlighted in the following sections.

The challenges to test-based learning highlight several needs. One is greater test-based visibility into product behavior to combat coarseness of test results. Another is a design flow that handles the imperfections that really impact the product, such as those mentioned above. Such a flow would allow diagnostic engineers to analyze the effect of those imperfections on the product. A second motivation for such a flow is that it would facilitate more aggressive designs. Currently, guardbands are put in place to handle unmodeled variability. When additional phenomena can be understood

978-1-4244-2820-5/08/$25.00 ©2008 IEEE

and modeled, guardbands can be reduced, leading to more aggressive designs. This paper describes an example of test-based learning that addresses the first need and provides insight into elements needed to address the second.

The term defects here refers to areas of extra or missing conducting or insulating material. Defects are usually the result of contamination or undesired artifacts such as scratches and usually cause shorts or opens. The defects that occur and how they affect circuit structure depend on details of the layout. As with parametrics, attempts are made to learn about defects using test structures, such as combs and serpentines, but complex interactions between layout features, including multiple-layer and neighborhood effects, make it very difficult to capture all relevant effects using an inevitably limited number of structures. Again, test, including slow-speed hard-defect-oriented tests and at-speed or out-of-specification soft-defect-oriented tests, is attractive as a basis for yield-learning because it inherently reflects the layouts and neighborhoods actually seen in the product. This paper discusses defect-oriented test-based learning, providing examples of reasons test structures alone are inadequate and discussing recent success in product-based learning.

II. Parametric Learning and Model-Hardware Mismatch

As described above, this paper describes an example of test-based learning that addresses the need for greater visibility in test and provides insight into elements needed to address improved variability modeling in the design flow. One portion of the approach is the use of ring oscillators [ROs] that have been included on recent IBM products [1-3]. In this section, examples of specific types of variation they reveal are given. Those types are not currently well-modeled in the design flow, but because of their systematic nature, are in fact amenable to both process actions and modeling. The ROs' correlation to product power-performance metrics is considered to establish their relevance. The variation decomposition techniques applied using ROs are then applied to real product data. Two categories of variation, mistracking among diverse circuits and spatial variation among identical circuits, are addressed.

The ROs used in this work are of two types – identical ROs distributed across the product and ROs representing diverse circuit types that allow for isolating various parasitic and transistor parameters. The design of these ROs and the principles on which their design and analysis are based are described in [1-3]. In this paper, we use specific data analysis techniques and examples that have previously been described in [2,4,5]. Note that the ROs have a large number of stages, typically one hundred. RO frequencies reflect average delay over the stages, which allows us to focus on systematic, as opposed to random, variability. We also look at real product test data. Specific examples from experimental hardware fabricated in 130 nm, 90 nm and 65 nm circuit nodes are given to illustrate the approach and its utility.

A. *Mismatch in Diverse Circuits*

The design flow uses circuit simulation models, e.g., BSIM, to characterize circuits and predict their behavior in the product. Process corners are used to cover shifts of the process, e.g., from fast to slow, but it is not necessarily the case that all parts of the process move together. The main variable determining process speed is typically channel length, but other aspects of the process, such as those determining other transistor parameters, like source-drain resistance or parasitic elements such as interconnect resistivity may move in a different way, and in a way that is not predicted by circuit simulation models. Specifically, there may be mismatches in the modeled tracking of one circuit type with another.

Figure 1 illustrates this mistracking by plotting the delays of a diverse set of ROs versus the delay of a reference RO [4]. The reference RO comprises inverters driving gate loads [1]. The values on each axis are normalized to predicted values based on circuit simulation. While the x-axis provides an overall view of where the process is running (points to the right of x=1 slow; points to the left of x=1 fast), the y-axis provides a view of how each diverse RO is running relative to model expectations. Since each RO is by design sensitive to parameters different from the reference RO, the tracking gives information on how those different parameters are running in the process.

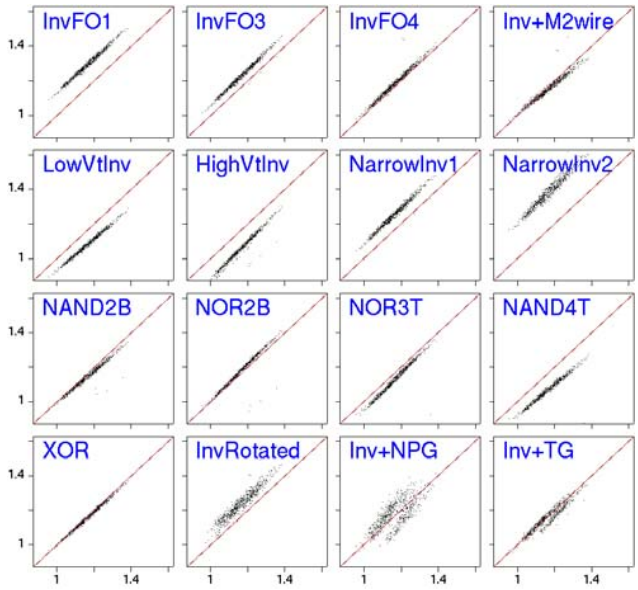

Figure 1. Normalized delays of diverse ROs versus normalized delay of a reference RO [4]. Diverse ROs comprise inverters ("Inv") unless otherwise stated. Loads include fan-outs of 1, 3 and 4 or metal-2 wire loads. Variations include low-Vt, high-Vt and two narrow widths. Also included are bottom-switching 2-input NAND and NORs, a 3-input top-switching NOR, a 4-input top-switching NAND and an XOR. Inverter ROs with rotated orientation and driving nfet-passgate and transmission gate loads are also included.

There are both systematic offsets and scatter in the plots. The systematic offsets suggest global mistracking of the parameters that affects all manufactured chips the same way, with some ROs tracking the reference closely and others mistracking by 10% or more. Scatter in the plot, on the other hand, suggests variability in the parameter to which the diverse RO is sensitive relative to the reference RO. Our experience suggests a key step for gaining insight into that scatter is looking at across wafer variations in mistracking. Figure 2 shows wafer maps that color code the delays of a diverse RO normalized to the reference RO at various positions on the wafer. The across wafer difference in mistracking accounts for scatter in the plots of Figure 1. Note also that the two maps show different patterns (crescent versus radial), which highlights the fact that equally fast circuits may differ in the composition of their delay, e.g., high overlap capacitance with low source-drain resistance or the opposite. More detailed description and examples of parameter mistracking based on circuit simulations are given in [2].

Note that both the offsets and across wafer patterns reflect variations that are systematic rather than random in nature. The systematic nature of the variations suggests they may be amenable to process fixes or design compensation. In addition, their effects can be modeled in the design process.

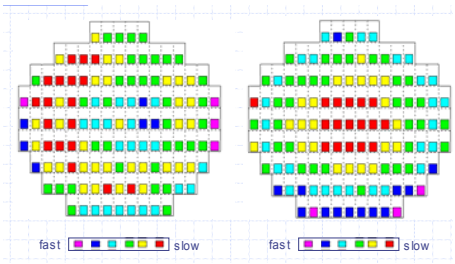

(a) (b)

Figure 2. Wafer maps color coding the delays of diverse ROs indicating source-drain resistance (a) and overlap capacitance (b) normalized to the reference RO [4].

These ROs give us insight into variations among diverse circuit types. Note that they offer a significant benefit in terms of product representativeness versus standard dc parametric measurements, e.g., of scribe line transistors, since their measurements are made under high-speed conditions like in the actual product, which helps them accurately reflect operational effects such as floating-body and self-heating [1-3]. As mentioned in Section 1, however, it is important for DfM and process actions to be driven by product impact. Although the ROs do not directly indicate product metrics, previous work has demonstrated correlation to real product metrics, specifically maximum operating frequency, FMAX, as measured by a logic built in self test. See [4] for an example of this correlation that specifically correlates behavior of different kinds of ROs to that of FMAX for different circuit blocks using delay dependence on power supply voltage. This type of correlation provides confidence in the usefulness of the ROs as proxies in our efforts to monitor variation. Using the ROs as proxies is very helpful as ROs are high-yielding and easy to measure, which leads to high data availability and robustness. That availability and robustness in turn aids our ability to map from test results to process weaknesses that matter and in turn to meaningfully guide DfM.

B. Mismatch in Identical Circuits

Timing tools generally assume a uniform model across chip, but there may be significant systematic variation across the chip even on like structures. There are many reasons for systematic variation in identical circuits. Some are related to temporal and/or spatial variation in manufacturing tools and equipment, e.g., bake plate thermal gradients or etch chamber gas flows. Others are related to physical design, such as pattern-density effects in anneal or photolithography neighbor-feature effects. When we look at product test results, electrical design effects also come into play, such as power-grid droop. See [6-10] for further discussion and references.

In addition variations occur at various repeat levels in the process, including lot-to-lot, wafer-to-wafer, across-wafer, and across-reticle/chip. Here we focus on across-chip variation which is currently not well-modeled in the design flow. Specifically we focus on systematic sources of across chip variation, using both ring oscillators and product data.

1) Ring-Oscillator-Based Characterization

Systematic across-chip variation has been identified and separated from across-wafer variation based on special purpose measurement structures, such as linewidth critical dimensions [CDs] [7,9,11,12] and chemical-mechanical polishing effects on inter-level dialectric thickness [6]. Prior work also has recognized systematic reticle-level variations using ring oscillators [4]. In this sub-section we focus on understanding and quantifying systematic across-chip variation owing to both reticle-level and wafer-level sources.

Again, many-stage ROs are used, which helps ensure focus on systematic, rather than random variation . Also, as in Section 2A, note that while ROs do not directly reflect product power-performance metrics, previous work has shown they correlate well to such metrics. In particular, [4] demonstrates correlation of ROs in geographic regions of an example chip with FMAX measurements on independently-testable blocks in those same regions. The correspondence between block FMAXs and RO delays allow us to study the ROs as proxies for product-performance.

Figure 3 shows scatter plots aimed at characterizing reticle-level variation [4,5] using 24 rings oscillators distributed across a product reticle field containing two chips. Each plot is placed in the figure at the location of the RO on the reticle and plots the corresponding RO delays on a manufactured chip versus the mean of the twenty-four RO delays for the reticle field that contains the chip. The dashed line represents the boundary between chips. Shifts in the populations relative to the green x=y reference line in each plot reflect the systematic across-reticle variation. Each plot is enhanced with a red or blue arrow whose size and darkness represent how fast (blue) or slow (red) the RO is relative to the mean for the reticle. For this example there is a fast-slow-fast pattern vertically across the reticle.

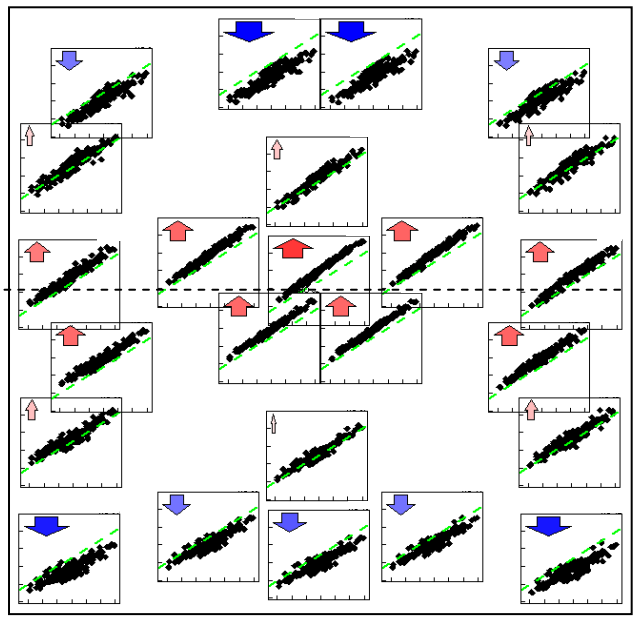

Figure 3. Scatter plots characterizing systematic across-reticle. Graphs are placed in the same locations as ROs on the reticle [5].

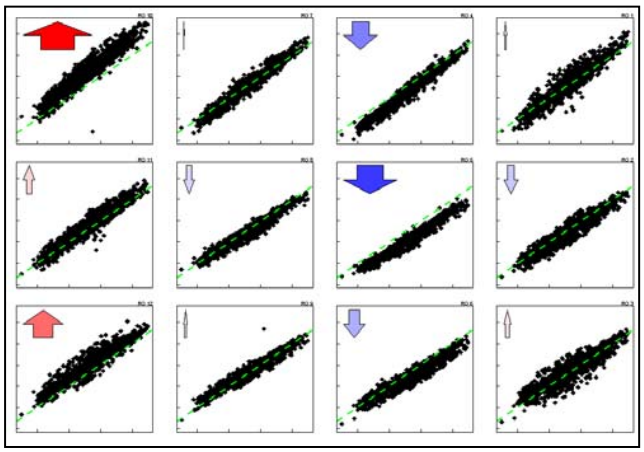

Figure 4. Scatter plots characterizing systematic across-reticle variation for a second product. Graphs are placed in the same locations as ROs on the reticle [5].

Note that this reticle-level pattern of variation means that across-*chip* variation varies depending on the chip's position on the reticle. Specifically the top chip, Chip 0, has a fast top and slow bottom, whereas the bottom chip, Chip 1, shows the opposite trend. With the chips oriented the same way in the reticle, on most chips, a core at the top of Chip 0 will be fast relative to the rest of the chip, while a core on the top of Chip 1 will be slow. Insight into the reticle-position-dependent nature of the across chip variation can be very useful for product debug and bring-up, especially for understanding the differences among manufactured chips. Figure 4 shows a similar set of plots for a second product, which this time has a single chip per reticle field and twelve distributed ROs per chip. Again, red and blue arrows reflect the reticle-level systematic across reticle variation.

In addition to the systematic offsets visible in the plots, there also exists scatter. A key systematic effect that can cause scatter in the plots is wafer-level variation [4]. In order to get a better understanding of wafer level variation, we examine high-resolution wafer maps that color-code delay at each RO location on the wafer. Figure 5a shows such a wafer map for the product corresponding to Figure 4. The data at each site is averaged over a lot. Figure 5a provides some indication of wafer-level systematic variation, but it is difficult to see. In this case, the reticle-level systematic across-reticle variation is obscuring the wafer level variation. To help clarify the wafer level picture, we quantitatively characterize the systematic reticle-level variation and subtract it from Figure 5a to product Figure 5b, resulting in a "corrected" wafer map that provides a much clearer picture of the wafer level variation. Further description is available in [5].

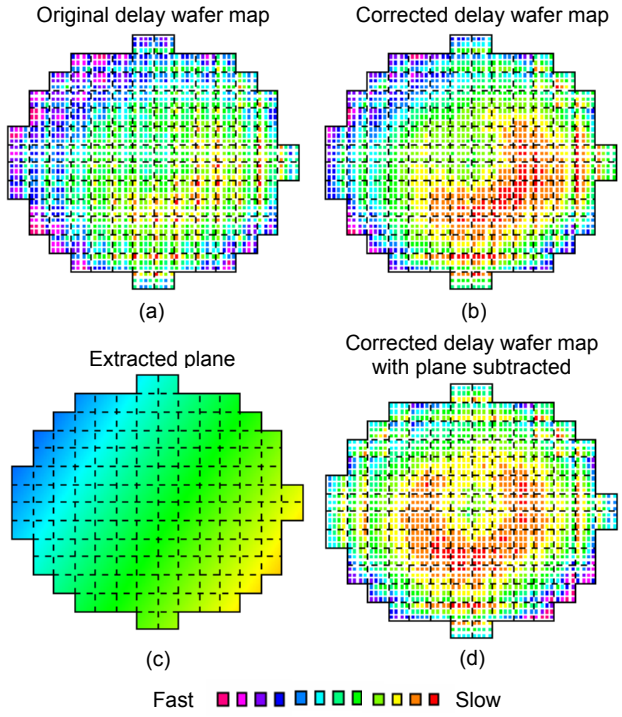

Figure 5. Original (a) and corrected (b) wafer maps color-coding RO delay for same product as Figure 5; plane model extracted from corrected wafer map (c); corrected wafer map with values modeled by the plane subtracted (d) [5].

Note that in general we focus on visual presentation and avoid making distribution assumptions. Doing so has practical advantages in accommodating complex patterns, which are especially likely to occur when looking at data from finished operating electronic circuits, which compound the effects of many process steps. Special-purpose hardware that looks at a single parameter can be expected to be influenced by fewer process steps and may have simpler patterns. The focus on visual presentation also has practical advantages for accommodating human-influenced patterns, such as intentional poly length striping and exclusion/special

processing regions in wafer processing, which are especially likely to be found on early/experimental hardware.

Figure 5b shows that after the reticle-level systematic has been removed, there remains across chip variation owing to remaining wafer-level gradients. Table 1 provides a quantitative breakdown of the within-chip variation into two different components (normalized): within-chip variation due to reticle-level variation, as characterized by Figure 4, versus within-chip variation due to interception of remaining wafer level gradients, as visually evident in Figure 5b. Often (non-random) within-chip variation is thought to occur primarily as a result of systematic across-chip variation, such as due to photolithography imperfections that get repeated site-to-site, e.g., mask or lens anomalies, or to design-density-dependent effects, such as those that occur during the chemical-mechanical-polishing step. Our analysis indicates, however that within-chip variation due to interception of wafer-level gradients can be significant as well.

TABLE I. BREAKDOWN OF WITHIN-CHIP VARIABILITY FOR EXAMPLE HARDWARE OF FIGURES 5 AND 6 (ARBITRARY UNITS) [5].

	Edge chips	Center chips
Within-chip variation due to systematic reticle-level variation	1	1
Within-chip variation due to interception of remaining wafer-level gradients	1.5-2.5	~0-1.5

Although there are advantages to focusing on visual analysis, taking advantage of the human capability for pattern recognition, further analysis of the wafer-level variation can be done. Various surfaces such as planes, radials or surfaces described by higher-order functions can be fit to the data and then subtracted to allow studying the remaining variation. As an example, Figure 5c shows a linear plane fit to the data from Figure 5b. Figure 5d shows the Figure 5b data with the plane fit subtracted. Presumably the delay variability modeled well by a plane has a different source in the process than then radial variation in Figure 5d. Subtraction of the plane helps illuminate the Figure 5d radial pattern, which can be very helpful in aiding process experts to recognize possible underlying process steps. In our normalized units, the plane contributes within-chip variation of about 0.1.

The breakdown in Table 1 can be used to guide corrective process actions. Large entries in Table 1 Row 1 suggest remedies in by-reticle process steps, such as photolithography, or in design-density dependent steps, such as chemical-mechanical polishing or remedies in physical design. Large entries in Row 2 suggest attention to by-wafer process steps, such as resist bake or etch. Although a relatively small effect here, wafer level gradients well-modeled by a plane, such as the one shown in Figure 5c, suggest remedies in wafer-level processes as well. In addition the spatial nature of the variation may provide an indication as to the processing step underlying it.

The analysis in this section and in particular the breakdown in Table 1 also can be used to guide DfM and modeling efforts. The Row 1 across-chip variation is systematic, rather than random, which can be used to guide modeling. Figure 4's two chip/reticle example also demonstrates systematic variability, although the effect is different chip-to-chip. The wafer-level sources of variation also show patterns suggesting a systematic nature. Models for across-chip variation used in design could be expanded to cover both reticle-level and wafer-level systematic contributions to across reticle variation.

2) *Product Test Data Based*

Our work has suggested ROs are good proxies for studying chip-level product-performance, but application to product test data allows even more direct use of the techniques we use to understand systematic variability. As an example, Figure 6 shows reticle-level systematic variability in product static power draw, IDDQ. Again the plots in the figure are placed in the same location as the circuit blocks being measured, where here each block is a chip on a four-chip reticle. In all cases log10 of IDDQ is shown. Note that clear systematic offsets exist, with arrows indicating red (high) and low (blue) systematic offsets. In addition, Figure 7 shows an original IDDQ wafer map and a version that has been "corrected" for reticle level systematic variation, as described above for ring oscillators. Subtraction of the reticle-level systematic clarifies the wafer level patterns.

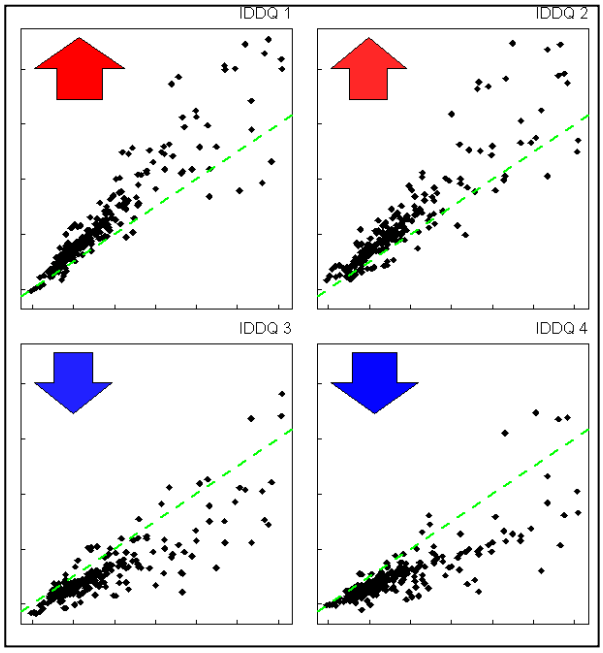

Figure 6. Scatter plots characterizing systematic across-reticle variation in log IDDQ [5].

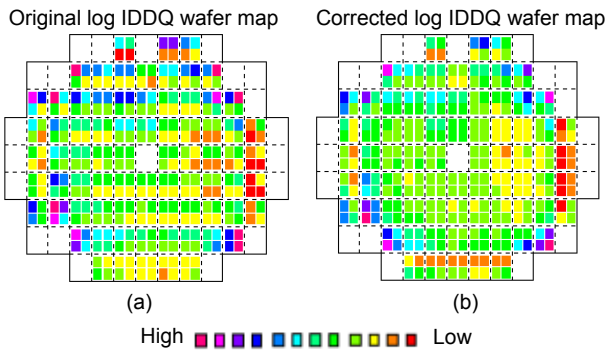

Figure 7. "Corrections" applied to wafer maps plotting log IDDQ measured on the product, original (a) and corrected (b) wafer maps [5].

Again both reticle-level and wafer-level sources of variation have been identified. The relevance is clear since we are using a product metric for the analysis. Since the measured circuits are identical, electrical and physical design of the blocks are exonerated as sole causes of the systematic across reticle differences. DfM focus would need to be at the supra-chip level, considering, for example, millimeter-range front-end density effects. Process fixes could be sought in the mask or photolithography. Product data presents yield-based and measurement-difficulty-based challenges to data availability and robustness that RO data does not. Nevertheless, there appear to be patterns visible in the corrected wafer map that likely have wafer-level sources. Again, corrective actions, DfM guidance and modeling efforts can be tailored to addressing the reticle and wafer-level sources of variability.

III. Defects

A. *Types of defects and DfM*

Previous sections have related to parametric yield. This section briefly discusses defect-related yield and quality. Historically most defects have been shorts, but the introduction of copper technology makes opens more likely [13]. A key attribute of both shorts and opens is their resistance. Higher resistance makes shorts more difficult to detect, while lower resistance makes opens more difficult to detect. Most shorts are hard shorts and most opens are hard opens [14-16], but soft shorts and opens are especially important because they are difficult and/or expensive to detect at and therefore are susceptible to escaping detection. As a result, defect-related DfM efforts needs to target not just the most likely defects (for yield), but also defects that are difficult or expensive to detect at test. The latter focuses not on yield, but outgoing quality level [16-17]. The rule of thumb that detecting a bad component increases by 10X at each level of assembly (wafer, package, board, system) provides vertically-integrated companies strong incentive to avoid the occurrence of hard-to-detect defects. Customer dissatisfaction upon receiving bad parts gives IC-providers strong incentive to do the same.

B. *3.2 Defect-Related Yield-Learning and DfM*

As described in the introduction, test structures can be used to learn about the occurrence of defects in a process, but it is not possible to create test structures that cover all relevant layout configurations. Specifically, in today's processes there are both *random-contamination-related* defect and *systematic* defects. Both are spots of extra or missing conducting or insulating material, but systematic defects are layout or design dependent. As feature printing increasingly becomes a function of neighboring patterns, systematic defects become more common.

Test structures may be suitable for monitoring contamination-related defects, but some systematic defects may be much harder to monitor because of the potential complexity of the layout configurations that lead to their occurrence. For example, feature-density-dependent CMP dishing and overpolish on one layer can cause shorts on the next layer up in the process. To monitor for such defects could require multi-layer test structures. Another example is density-gradient-dependent salicide thickness variations causing poor contact formation at the edges of arrays in a given process [18]. Such a defect mechanism would be extremely difficult to discover using test structures.

The difficulty of creating adequate test structures makes defect-learning based on products very attractive. Product-based defect diagnosis has made great progress in the last decade. In particular, diagnosis strategies that avoid making strict assumptions about defect behavior conforming to fault models, such as the stuck-at fault model, have been very successful [19-23]. Fewer assumptions allow more flexibility and an attendant greater likelihood of successful results on hard-to-detect and likely hard-to-monitor systematic defects, as described above.

IV. Conclusions

Achieving good quality and yield requires DfM and DfM requires knowledge of process weaknesses. The complexity of interactions in today's processes make it very difficult to guide DfM with learning based solely on test structures and experiments. As a result, product-impact-oriented learning based on test results is attractive. This paper has provided examples of learning about model-hardware mismatch based on product test results and ROs that correlate to product test metrics. Specific variations including mistracking of diverse circuits and both reticle-level-based and wafer-level-based systematic across-reticle variation have been identified. Each are systematic rather than random in nature and so are amenable both to corrective measures and DfM actions. Specifically, modeling them in the design process would both provide a valuable aid for diagnosis and allow for decreased guardbands currently in place to cover unmodeled variation. Product-impact-based DfM guidance is important for defect-related yield as well. That guidance needs both to be based on outgoing product quality rather than just yield, and also to utilize product diagnostics to ensure coverage of complex systematic defects.

ACKNOWLEDGMENTS

The author deeply appreciates the contributions of many other individuals in IBM Research and IBM Systems & Technology Group for the design, manufacturing and measurements utilized in this work. Collaborations with Manjul Bhushan, Mark Ketchen and Koushik Das and discussions with Sani Nassif are especially appreciated.

REFERENCES

[1] M. Ketchen, M. Bhushan and D. Pearson, "High Speed Test Structures for In-Line Process Monitoring and Model Calibration," Intl. Conf. on Microelectronic Test Structures, 2005.

[2] M. Bhushan, A. Gattiker, M. Ketchen and K. Das, "Ring Oscillators for CMOS Process Tuning and Variability Control," IEEE Trans. On Semiconductor Manufacturing, Vol. 19, No. 1, 2006.

[3] M. Ketchen and M. Bhushan, "Product Representative 'At-Speed' Test Structures for CMOS Characterization," IBM Journal of Res. & Dev., Vol. 4/5, 2006.

[4] A. Gattiker, M. Bhushan and M. Ketchen, "Data Analysis Techniques for CMOS Technology Characterization and Product Impact Assessment," Intl. Test Conference, 2006.

[5] A. Gattiker, "Unraveling Variability for Process/Product Improvement," IEEE International Test Conference, 2008.

[6] B. Stine, D. Boning and J. Chung, "Analysis and Decomposition of Spatial Variation in Integrated Circuit Processes and Devices," IEEE Transactions on Semiconductor Manufacturing, Vol. 10, No 1., 1997.

[7] M. Orshansky, L, Milor and C. Hu, "Characterization of Spatial Intrafield Gate CD Variability, Its Impact on Circuit Performance, and Spatial Mask-Level Correction," IEEE Trans. On Semiconductor Manufacturing, Vol. 12, No. 1, 2004.

[8] K. Abrokwah, P. Chidambaram and D. Boning, "Pattern Based Prediction for Plasma Etch," IEEE Trans. On Semiconductor Manufacturing, Vol. 20, No. 2, May 2007.

[9] P. Friedberg, "Spatial Modeling of Gate Length Variation for Process-Design Co-Optimization", Ph.D. Thesis, University of California at Berkeley, Berkeley, CA, 2007.

[10] D. Boning et al., "Variation," IEEE Trans. On Semiconductor Manufacturing, Vol. 21, No. 1, 2008.

[11] A. Wong, et al., "Linewidth Variation Characterization by Spatial Decomposition," Journal of Microlithography, Microlithography, Microfabrication, and Microsystems, Vol. 1, Issue 2, 2002.

[12] J. Cain and C. Spanos, "Electrical Linewidth Metrology for Systematic CD Variation Characterization and Causal Analysis," Proceeding of SPIE, 2003.

[13] A. Stamper, T. McDevitt, S. Luce, "Sub-0.25-Micron Interconnection Scaling: Damascene Copper Versus Subtractive Aluminum," Advanced Semiconductor Manufacturing Conference and Workshop, 1998.

[14] R. Rodriguez Montanes, J. Figueras and E. Bruls, "Bridging Defects Resistance Measurements in a CMOS Process," IEEE Intl. Test Conference, 1992.

[15] R. Rodriguez Montanes, P. Volf and J. Pineda de Gyvez, "Resistance Characterization for Weak Opens," IEEE Design & Test of Computers, Vol. 19, Issue 5, 2002.

[16] A. Gattiker, "IC Failure Mechanisms Yesterday, Today, Tomorrow: Implications from Test to DFM," IEEE Intl. Symp. On Physical Design, 2006.

[17] A. Gattiker, "Well-Targeted Design-for-Manufacturability [DFM] Through Test," IEEE Intl. Test Conference, 2007.

[18] S. Chakravarty, K. Komeyli, E. Savage, M. Carruthers, B. Stastny, S. Zachariah, "Layout Analysis to Extract Open Nets Caused by Systematic Failure Mechanisms," VLSI Test Symposium, 2002.

[19] S. Vankataraman and S. Drummonds, "POIROT: A Logic Fault Diagnosis Tool and Its Applications," IEEE Intl. Test Conference, 2000.

[20] T. Bartenstein, D. Heaberlin, L. Huisman and D. Sliwinski, "Diagnosing Combinational Logic Designs Using the Single Location at-a-time (SLAT) Paradigm," IEEE Intl. Test Conference, 2001.

[21] D. Lavo, I. Hartanto and T. Larrabee, "Multiplets, Models and the Search for Meaning: Improving Per-Test Fault Diagnosis," IEEE Intl. Test Conference, 2002.

[22] R. Desineni and R. Blanton, "Diagnosis of Arbitrary Defects Using Neighborhood Function Extraction," IEEE VLSI Test Symposium, 2005.

[23] D. Heaberlin, "The Power of Exhaustive Bridge Diagnosis using IDDQ: Speed, Confidence and Resolution," IEEE Intl. Test Conference, 2006.