Center-of-delay: a new metric to drive timing margin against spatial variation in complex SOCs

Christian Lütkemeyer

Marvell Semiconductor, Inc. Irvine, CA, USA clutkemeyer@marvell.com

Abstract—Complex VLSI SOCs are manufactured on large 300mm wafers. Individual SOCs can show significant spatial performance gradients in the order of 10% per 10mm. The traditional approach to handling this variation in STA tools is a margin look-up table indexed by the diagonal of the bounding box around the gates in a timing path. In this paper we propose a new approach based on the concept of the Center-of-Delay of a timing path. We justify this new approach theoretically for linear performance gradients and present experimental data that shows that the new approach is both safe, and significantly less pessimistic than the existing method.

Index Terms-static timing analysis, CMOS, variability

I. INTRODUCTION

Manufacturing¹ today's advanced nanometer-scale FinFET CMOS systems-on-chip (SOC) requires amazing accuracy and resolution. Foundries are able to integrate many billions of transistors on a single chip and connect them over many metal layers, creating systems of amazing complexity and sophistication that provide the communication infrastructure for today's omni-connected society.

Modern SOCs are created on large 300mm wafers in many dozens of manufacturing steps. To ensure high yield for all SOCs on a wafer, manufacturing is optimized to maintain tight control of the material properties and physical device parameters that drive the dynamic performance of digital and analog circuits: gate length, gate oxide thickness, fin height, threshold voltage, and interconnect capacitance. Despite a strong focus on manufacturing uniformity of these parameters across entire wafers, performance measurements show significant dynamic performance gradients (see [1] and later in this paper). Digital circuit designers who want to ensure robust yield for all locations on a wafer must take spatial performance variation into consideration.

Spatial performance variation has plagued CMOS manufacturing already since early 2000s [2]–[6] and the established industry practice is to account for this variation during Static Timing Analysis (STA). Although a few Statistical STA (SSTA) based approaches to this problem have been proposed initially [7], [8], a simple margining scheme driven by the length of diagonal of the *bounding-box* around the gates of a timing path Anton Belov Synopsys, Inc. Dublin, Ireland anton.belov@synopsys.com

[9]–[11] was adopted by the industry and remains dominant to this day.

In this paper we argue that the bounding-box based approach is often unnecessarily pessimistic, thus causing over-design and degradation of PPA. As an alternative, we propose an approach based on the novel concept of the *Center-of-Delay* of a timing path. We provide theoretical justification of the new approach and demonstrate its effectiveness empirically.

This paper is structured as follows: after providing the necessary background and motivation in Section II, we introduce the new approach and prove its theoretical soundness in Section III. In Section IV we present the results of our experimental study. We conclude by outlining some of the directions for future work in Section V.

II. BACKGROUND AND MOTIVATION

We assume that the reader is familiar with the basic concepts of semiconductor design and manufacturing, as well as basic understanding of Static Timing Analysis (STA).

A. Spatial performance gradient

As highlighted in Section I, despite a strong focus on uniformity of process parameters across entire wafers, performance measurements of manufactured silicon show significant performance gradients. Figure 1 shows a wafer diagram with digital supply voltage data that maintains constant performance for the slowest location on SOCs placed across the wafer, averaged over multiple manufactured wafers. Higher supply voltage implies slower silicon and vice versa. We see a performance profile that resembles a Mexican hat with a slower center (VDD up to 0.67V), a faster ring (VDD reduced down to 0.61V), and a slow outer rim (VDD increases up to 0.66V). Looking for the steepest voltage gradient, we find a case with 0.05V difference for a pair of SOCs that are touching at their corners in the North-West section of the rim. This voltage differential translates to a relative dynamic spatial performance change of about 10% over a distance of 10mm. For SOCs with a die size of $10 \times 10 mm^2$ in the process used to manufacture the wafers this translates to about 2.5 sigma of global variation in a single die when the gradient aligns with the diagonal of the SOC.

Although the impact of spatial variation can be significant, further analysis of the wafer map in Figure 1 shows that large spatial performance gradients are in fact infrequent. Therefore, it is important that a spatial margin model does not increase

¹It would be more appropriate to use the term "Machinefacturing" for the high precision "machine"-pulation of matter and materials that the foundries do. Human touch is completely forbidden.



Fig. 1. Voltage distribution on a FinFET wafer to maintain constant performance. Higher voltage (red) corresponds to slower silicon.



Fig. 2. Launch- and capture clock path cell placement examples with spatial gradient concern. Large sensitivity (left) and small sensitivity (right). Performance gradient is represented by background color, bounding-boxes are drawn in red.

the power or area of a design substantially to only recover a fairly small number of chips per wafer.

B. Bounding-box based margining

The bounding-box based spatial margining methodology has been described first in the literature circa 2009 [9], [10], but to our knowledge had already been a "folk-knowledge" in the design industry by then. In this methodology a STA tool is provided with a table that represents timing *derate*, i.e. cell or net delay multiple, as a function of a distance on chip. A separate table is given for early and late, cell and net derates. A fragment of a distance derate table is shown in Figure 3.

ocvm_type: pocvm object_type: design rf_type: rise fall delay_type: cell derate_type: late object_spec: distance: 0 5000000 10000000 20000000 30000000 40000000 table: 1.00000 1.05403 1.07633 1.10796 1.13228 1.15265

Fig. 3. Example distance derate table (fragment).

For a given timing path, the STA tool computes the coordinates of the smallest rectangle that encloses all of the devices of the path starting from the CRP common point – this is the so-called *bounding-box* of the timing path, see Figure 2. The length of the diagonal of the bounding-box is the quantity that is used to look up the timing derates from the late (resp. early)

distance derate tables to derate the delays of devices on launch (resp. capture) segments of the timing path. A similar operation is performed for nets to create margin for RC delays – note that the bounding-box for nets can be different from that for cells.

It is beneficial now to highlight the following subtlety. Consider a chip subjected to some spatial performance gradient. Clearly two topologically identical timing paths in different parts of the chip are likely to have different timing slack. Yet, since the diagonals of bounding-boxes for these paths are equal, both paths will be subjected to the same distance-based timing derate. This illustrates that distance-based derating is applied in order to margin for *difference* in performance shift of launch vs capture segments of a path, rather than a global performance shift due to spatial gradient. The latter is accounted in practice via corner library models.

With the above in mind, consider a chip with a linear performance gradient with faster devices in the top-left corner and slower devices in the bottom-right corner. Figure 2 shows two timing paths with tight hold time constraints and long divergent clock paths. On the left, the spatial performance gradient can introduce significant skew because clock paths are placed far from each other, with the launch path tracking along the faster sides of the bounding-box while the capture tracks along the slower bottom and right edges. On the right, launch and capture clock paths track each other closely and so the clock skew sensitivity to spatial gradients on a wafer will be small. Yet, both timing paths have exactly the same boundingbox, and hence will be subjected to the *same* distance derate.

This example demonstrates that the bounding-box based margining methodology is agnostic to the detailed placement of cells on the timing paths and therefore can be unnecessarily pessimistic.

Furthermore, even the task of translating the performance gradients measured on manufactured silicon into spatial derate values for the bounding-box model poses a challenge. As demonstrated above, the required margin depends on the specific placement of the gates on the launch and capture paths. Although a theoretical worst case, with the launch delay concentrated in the fastest corner of the box and the capture delay concentrated in the slowest corner, can be constructed this would be a rather unrealistic scenario given the fact that the launch and capture paths originate from the common launch point, and both terminate at the common capture point². Thus in practice distance derate tables are frequently adjusted in an *adhoc* manner to reduce the artificial pessimism imposed by the bounding-box model. Clearly such adjustments cannot be done correctly since the required margin depends on path topology.

These fundamental deficiencies of bounding-box based spatial margining motivates the development of the novel, *Centerof-Delay* based methodology, described in the remainder of this paper.

²For larger distances the implementation of the clock distribution network typically uses repeaters which make the theoretical worst case configuration still more unlikely.

III. CENTER-OF-DELAY

The Center-Of-Delay (COD) calculation of a timing path uses the well-known equations of the center-of-gravity calculation of distributed weights. Cell delays replace weights in the locations of the individual cells. We will formally prove below that the delay of a distributed timing path on a linear gradient performance derating plane can be calculated by adding the underated cell delays and multiplying them with a fixed derate from the location of the center-of-delay of the path. This property dramatically simplifies the calculation of the timing slack shift between launch and capture paths in STA in the presence of linear spatial gradients.

A. Definition

A timing graph (P, A) is a directed acyclic graph with a set of nodes P that represents pins or ports, and a set of edges A that represents timing arcs, i.e. either cell library arcs or nets, in the circuit. For clarity of presentation, in this paper we assume that each timing arc $a \in A$ is associated with a *unique* nominal delay value τ_a and that net arcs have zero delay. For a path $\pi = (a_1, \ldots, a_n)$ in a timing graph, we will write τ_{π} to denote the delay of π , i.e. the sum of delays τ_{a_i} of the timing arcs in π .

We assume that the die is 2D and for each pin $p \in P$, (x_p, y_p) is its x and y coordinates on the die with respect to some origin (e.g. bottom-left corner of the die). For a timing arc a = (p, q) the coordinates (x_a, y_a) of a are taken to be $(\frac{x_p+x_q}{2}, \frac{y_p+y_q}{2})$.

The COD is calculated by accumulating delay-moments $x_a \cdot \tau_a$, $y_a \cdot \tau_a$ in x and y direction over the path, and then dividing the moment sums by the accumulated path delay:

Definition 1. Let $\pi = (a_1, \ldots, a_n)$ be a path in a timing graph. The center-of-delay (COD) of π , denoted as C_{π} , is a point (x, y) on a die defined as:

$$C_{\pi} = \left(\frac{\sum_{a \in \pi} x_a \tau_a}{\tau_{\pi}}, \frac{\sum_{a \in \pi} y_a \tau_a}{\tau_{\pi}}\right).$$
(1)

Example 1. Consider the example circuit depicted in Figure 4. The launch path has 4 buffers with delay 1ns at coordinates $(1mm, 1mm), \ldots, (1mm, 4mm)$. Thus, the COD of the launch path is located at coordinate $((1mm \cdot 1ns + 1mm \cdot 1ns + 1mm \cdot 1ns)/4ns, (1mm \cdot 1ns + 2mm \cdot 1ns + 3mm \cdot 1ns + 4mm \cdot 1ns)/4ns) = (1mm, 2.5mm)$. Similarly, the COD of the capture path is located at (4mm, 1.5mm).

The key property of COD is that in the presence of a *linear* spatial performance gradient, the delay of a timing path can be computed *exactly* based on the delay of the path at the location of COD. In the rest of this sub-section we formalize this statement and prove its correctness.

To capture the impact of spatial gradient on delay of timing arcs, we will use a linear *spatial gradient function* $G(x, y) = g_x x + g_y y + 1$ to represent the timing derate, i.e. delay multiple, for the nominal delay of a timing arc at coordinates (x, y). That is, the derated delay of arc *a*, denoted by τ'_a , is computed as

$$\tau'_{a} = G(x_{a}, y_{a}) \cdot \tau_{a} = (g_{x}x_{a} + g_{y}y_{a} + 1) \cdot \tau_{a}.$$
 (2)



Fig. 4. The running example for the paper. Assume that nominal delay of every buffer is 1ns, and the rest of elements (nets and cells) have delay of 0. Distance units are mm.

For a path π we will denote by τ'_{π} the derated delay of π , i.e. the sum $\sum_{a \in \pi} \tau'_a$.

Note that points on the plane with spacial derate equal to d define a line $g_x x + g_y y + (1 - d) = 0$, and in particular, the line $g_x x + g_y y = 0$ through origin (0, 0) is the line with spatial derate equal to 1. The terms g_x and g_y represent the x and y components of the gradient vector.

Example 2. Consider again the circuit in Figure 4. Assume that the gradient function is defined as G(x, y) = 1 + 0.01 * y. This corresponds to 1% per 1mm derate in the direction of y-axis, with derate 1.0 line coinciding with the x-axis. Considering the launch path, the buffer at coordinate (1mm, 1mm) has delay of 1.01ns, the buffer at (1mm, 2mm) has delay of 1.02ns, etc. Thus, the derated delay τ'_L of the launch path is 1.01ns +1.02ns + 1.03ns + 1.04ns = 4.10ns. Similarly, the derated delay τ'_C of the capture path is 1.00ns + 1.02ns +1.04ns = 4.06ns

Now, the main fact of this sub-section can be stated as follows:

Theorem 1. Let π be a timing path with center-of-delay C_{π} and G(x, y) be a linear spatial gradient function. Then,

$$\tau'_{\pi} = G(x_{C_{\pi}}, y_{C_{\pi}}) \cdot \tau_{\pi} \tag{3}$$

That is, the derated delay of timing path π equals its nominal delay derated at location of C_{π} .

Proof. From (2) for an arc $a \in \pi$, we have the derated delay of a as $\tau'_a = (g_x x_a + g_y y_a + 1)\tau_a$. Then, the derated delay of

 π is

$$\begin{aligned} \tau'_{\pi} &= \sum_{a \in \pi} \tau'_a = \sum_{a \in \pi} (g_x x_a + g_y y_a + 1) \tau_a \\ &= \sum_{a \in \pi} g_x x_a \tau_a + \sum_{a \in \pi} g_y y_a \tau_a + \sum_{a \in \pi} \tau_a. \end{aligned}$$

Dividing and multiplying the latter by τ_{π} and distributing the division we have

$$\tau'_{\pi} = \left(\frac{\sum_{a \in \pi} g_x x_a \tau_a}{\tau_{\pi}} + \frac{\sum_{a \in \pi} g_y y_a \tau_a}{\tau_{\pi}} + \frac{\sum_{a \in \pi} \tau_a}{\tau_{\pi}}\right) \cdot \tau_{\pi}$$
$$= \left(g_x \frac{\sum_{a \in \pi} x_a \tau_a}{\tau_{\pi}} + g_y \frac{\sum_{a \in \pi} y_a \tau_a}{\tau_{\pi}} + 1\right) \cdot \tau_{\pi}.$$

Recall from (1) that $\frac{\sum_{a \in \pi} x_a \tau_a}{\tau_{\pi}}$ is exactly the x-coordinate $x_{C_{\pi}}$ of COD of π , and similarly for $\frac{\sum_{a \in \pi} y_a \tau_a}{\tau_{\pi}} = y_{C_{\pi}}$. Therefore, taking into account the definition of G(x, y), we have

$$\tau'_{\pi} = (g_x x_{C_{\pi}} + g_y y_{C_{\pi}} + 1) \cdot \tau_{\pi} = G(x_{C_{\pi}}, y_{C_{\pi}}) \cdot \tau_{\pi}.$$

Example 3. Referring to the circuit in Figure 4, assume again that the gradient function is defined as G(x, y) = 1 + 0.01 * y. From Example 1 we know that COD of the launch path is located at (1mm, 2.5mm) - the derate at this location is 1 + 0.01 * 2.5 = 1.025. Since the nominal delay of the path is 4ns, using Theorem 1 we compute the derated delay of the path as 1.025 * 4ns = 4.10ns, that is, equal to the derated delay of this path we computed in Example 2. Similarly, COD of the capture path is at (4mm, 1.5mm). The derate at this location is 1.015, giving the derated delay of the capture path 1.015 * 4ns = 4.06ns - again, the value we computed in Example 2.

Theorem 1 implies that instead of calculating actual path delays over a large set of possible gradient directions to determine the worst-case slack shift we can immediately identify the worst gradient direction from the direction of the vector from launch path COD to capture path COD. Furthermore, the worst case difference in the derating values for the two paths equals to the distance between their CODs multiplied by the magnitude of the gradient.

These considerations suggest that it is the *distance between CODs of launch and capture paths* that should determine the amount of spatial margin in STA.

B. Application of COD to spatial derating in STA

In this work we propose to reuse the margin lookup and derating mechanism of the existing diagonal-of-bounding-box method in STA tools, but use the distance of CODs between non-common launch and capture paths to look up the corresponding derating margin values.

The data in the derating tables will have to be adjusted to reflect the change in derating methodology. As explained in Section II-B, due to inherent indiscriminate pessimism imposed by the bounding-box method, distance derate tables are frequently adjusted in a heuristic manner to reduce the pessimism artificially. Our experimental data (Section IV) shows that the distance between CODs of launch and capture paths is in general significantly smaller than the diagonal of the bounding box around the paths. Therefore for the COD derating tables the derating values vs. distance can be taken directly from the measured spatial performance shift data, without any additional manipulation. Pessimism reduction is now provided by the explicit consideration of gate locations in the COD calculation. Furthermore, the worst-case placement of clock path delays in opposite corners of the bounding box would be correctly captured in our methodology, as opposed to being potentially optimistic due to the ad-hoc manipulation of bounding-box derates.

In addition, a practical methodology for COD-based spatial derating should take into account the following observations:

- Hold time violations are the primary concern created by spatial performance gradients in digital designs. If long clock insertion delays are implemented on regions of the SOC with significant differences in performance, there may be hold time failures on short data paths where the launch clock propagates substantially faster than the capture clock (see Fig. 2, left). If manufacturing is mature and well within the SPICE corners, setup violations are unlikely unless other margin deficiencies exist (e.g. underestimation of IR drop).
- The corner-focused STA analysis does not reflect the actual silicon with performance gradients accurately. To align reality and spatial margin models as much as possible, it makes sense to consider that if a strong gradient occurs on a chip with slow silicon we should assume that the gradient will lead to a speed-up. On the other hand, in case of a strong gradient with the fast silicon model the gradient should lead to slowdown. These assumptions will ensure that performance modeling overall stays inside the SPICE corners. These assumptions can be captured in different early and late derating tables for a spatial margin model.

Keeping the performance inside the library corners, and with hold time violations as the primary concern, the following derating schemes appear best aligned with silicon reality to identify endpoints with significant spatial variation hold margin exposure:

- Slow corner analysis: Speed-up of the launch path relative to the library corner model is the fundamental concern. This can be captured by assuming that the COD of the capture path is exactly at the slow library modeling corner, and distance of COD indexed derating values smaller than one are prepared in the early derating tables. Based on the data discussed at the beginning of this paper a spatial derate of 0.9 for a distance of CODs of 10mm can be justified. For distances larger than 10mm a less than linear reduction of the margin derate appears appropriate as the data does not show sustained gradients of 10%/10mm over significantly longer distances.
- 2) **Fast corner analysis**: An analysis with a fast corner library can model the hold margin concern by assuming that the COD of the launch path is at the fast library

 TABLE I

 COMPARISON OF DISTANCE-DERATING METRICS DERIVED USING COD VS BOUNDING-BOX (BBOX) METHODOLOGY.

Design	Analysis	Die diagonal	Hold analysis			Setup analysis			
	corner	(mm)	COD distance	BBOX diagonal	BBOX/	COD distance	BBOX diagonal	BBOX/	
			avg / max (mm)	avg / max (mm)	COD	avg / max (mm)	avg / max (mm)	COD	
design-1	SS	5.732	0.246 / 1.123	0.663 / 2.703	2.7 <i>x</i>	0.926 / 2.138	2.038 / 3.753	2.2x	
design-2	FF	16.127	0.105 / 2.134	0.458 / 5.663	4.4x	0.470 / 2.861	1.900 / 6.492	4.0x	
design-3a	SS	9.143	0.010 / 0.212	0.035 / 0.628	3.5x	0.061 / 2.068	0.208 / 3.531	3.4x	
design-3b	TT	9.143	0.050 / 2.628	0.342 / 5.977	6.8x	0.292 / 5.020	0.629 / 7.601	2.2x	
design-4	SS	11.773	0.007 / 0.133	0.027 / 0.348	3.9x	0.062 / 0.812	0.246 / 2.103	3.9x	
design-5	FF	21.092	4.402 / 12.595	12.141 / 20.062	2.8x	4.068 / 16.708	15.578/ 19.771	3.8x	
design-6	SS	9.866	0.207 / 1.420	1.630 / 6.619	7.9x	0.434 / 1.888	1.289 / 4.142	2.9x	

characterization corner (i.e. early analysis uses a derate of exactly 1.0), and populating the late analysis derating table with values above 1.0 to model the slowdown of the capture clock. For a distance of 10mm, the late derating table would contain a value of 1.1 to model the 10% slowdown that was observed in silicon for this distance. Larger distances should again see a less than linear increase of the derate values to reflect the spatial variation that was observed for longer distances.

3) **Typical corner analysis:** If designers model a timing scenario that targets the typical device model it makes sense to split the derating values between both the early and late side of the analysis evenly. For a distance of 10mm the early derating table would contain a value of 0.95 while the late derating table sets a value of 1.05.

IV. EXPERIMENTAL STUDY

To evaluate the ideas presented in this paper we performed an experimental study with the following objectives. For one, we wanted to quantify the reduction of distances used to index derates in our COD-based derating scheme vs. the currently Dominant bounding-box based methodology on realistic timing paths in modern SOCs. Consequently, we were interested in the impact of the reduction of the distances on the timing slack, assuming a gradient with a known magnitude. Finally, we were interested in affirming theoretical claims of Section III.

To this end we altered a commercial timer to enable (i) calculation of COD and application of COD-based timing derates in path-based analysis (PBA) and (ii) simulation of the physical spatial performance gradient given a user-provided gradient vector during PBA. The latter capability allows us to obtain "golden" data by rotating the simulated gradient plane and recording the impact of various gradient directions on the slack of timing paths.

In order to study the performance of the new technique on real-world designs, we selected a set of industrial IP blocks configured for STA flow that included bounding-box based spatial gradient derating. The blocks had die diagonals in the range of 5mm-21mm, covered different analysis, and one of the blocks was available in two modes. While the evaluation was primarily focused on hold analysis (for the reasons explained in Section III-B), for completeness we also collected the data for setup analysis³. Note that data-path delay is included in

both setup and hold analysis. For each of the designs and hold / setup combination, we collected the worst *nominal* – that is, un-derated – slack path to 10000 endpoints. For each path we then computed: (*i*) the distance between launch and capture CODs, (*ii*) the diagonal of bounding-box, (*iii*) the slacks of the path with cell delays derated based on the COD distance and based on the bounding-box diagonal, and (*iv*) performed a simulation to obtain the "golden" slack and confirm that it is properly bounded by both the COD and BBOX derived slacks.

The main finding of our study is presented in Table I, where we collated and contrasted the average and maximum distance between launch and capture CODs versus the corresponding metrics derived from bounding-box diagonals. We observe that average COD distance is *significantly* smaller than diagonal of bounding-box – this holds true across all analysis corners and both the setup and hold analysis. It is also important to note that even in the outlier cases, i.e. with maximum distances, COD distances are still significantly smaller than bounding-box diagonals. The top two plots in Figure 5 provide a graphical view into the magnitude of reduction in distance metrics in COD based methodology compared to the traditional boundingbox method on two example designs.

To evaluate the impact of the reduction in distance metrics on timing, we also compared the corresponding slack data. As the comparison of the absolute value of slacks is not particularly meaningful without taking into account the nominal slack, we opted to compare the *slack margin* – the difference between the derated and the nominal slacks. The bottom two plots in Figure 5 demonstrate the impact of COD based derating on slack margin graphically, while Table II presents the summary statistics comparing both the absolute and the relative slack margin values for COD and bounding-box methodologies, where *relative* slack margin is computed as percentage of absolute slack margin with respect to path arrival time⁴.

V. CONCLUSIONS AND FUTURE WORK

In this paper we proposed a novel methodology for spatial performance gradient margining in STA, based on the concept of the Center-of-Delay (COD) of a timing path. We demonstrated both theoretically and experimentally that the proposed approach is superior to the established bounding-box derating methodology due its ability to take the path topology into account and therefore remove unnecessary blanket pessimism present in the current practice.

³The COD-based derating methodology for setup analysis can be derived analogously to hold following Section III-B.

⁴Designs with relative slack margin < 1% in BBOX are not shown.



Fig. 5. Scatter plots comparing the distance and the slack margin of COD-based versus bounding-box based margining.

 TABLE II

 Comparison of slack margins computed using COD vs bounding-box (BBOX) methodology.

Design		Hold ana	lysis		Setup analysis			
	Abs. margin av	/g / max (ns)	Rel. margin avg / max (%)		Abs. margin avg / max (ns)		Rel. margin avg / max (%)	
	COD	BOX	COD	BBOX	CÕD	BOX	COD	BBOX
design-1	0.002 / 0.012	0.004 / 0.027	0.14 / 0.67	0.38 / 1.53	0.005 / 0.012	0.010 / 0.021	0.12 / 0.31	0.26 / 0.63
design-2	0.001 / 0.038	0.005 / 0.114	0.05 / 1.10	0.23 / 2.95	0.010 / 0.091	0.038 / 0.212	0.31 / 1.84	1.21 / 4.30
design-3b	<0.001 / 0.107	0.004 / 0.244	0.03 / 1.31	0.17 / 2.99	0.012 / 0.619	0.026 / 1.133	0.25 / 4.45	0.53 / 6.51
design-5	0.063 / 0.226	0.226 / 0.518	2.88 / 6.83	7.61 / 13.6	0.228 / 3.938	1.002 / 14.92	3.67 / 16.5	13.1 / 19.5
design-6	0.006 / 0.069	0.050 / 0.326	0.11 / 0.77	0.89 / 3.63	0.006 / 0.042	0.019 / 0.105	0.09 / 0.49	0.26 / 1.20

Our work opens a number of interesting research and engineering directions. For instance, in this paper we mostly glanced over the implementation details of the integration of COD-based methodology in an industrial-strength timer. Although path-based analysis does not pose particular challenge, the efficient implementation of COD-based derating in graphbased analysis (GBA) will be a subject of future work. Another line of research is the extension of the ideas presented in this paper to practical non-linear gradients, such as those arising from temperature and voltage variations.

REFERENCES

- C. Lütkemeyer, "Where is my Typical Chip ? Relating Silicon Back to the Timing Sign-Off Model," in *Proc. of TAU Workshop 2019 (online)*, 2019.
- [2] P. S. Zuchowski, P. A. Habitz, J. D. Hayes, and J. H. Oppold, "Process and environmental variation impacts on ASIC timing," in *Proc. of IEEE/ACM International Conference on Computer-Aided Design, ICCAD*, 2004.
- [3] P. Friedberg, Y. Cao, J. Cain, R. Wang, J. Rabaey, and C. Spanos, "Modeling within-die spatial correlation effects for process-design cooptimization," in *Proc. of International Symposium on Quality Electronic Design, ISQED*, 2005.

- [4] J. Nakanishi, H. Notani, Y. Nakase, and H. Shinohara, "Analysis technique for systematic variation over whole shot and wafer at 45 nm process node," in *Proc. of 8th IEEE International Conference on ASIC, ASCION*, 2009.
- [5] S. Reda and S. R. Nassif, "Analyzing the impact of process variations on parametric measurements: Novel models and applications," in *Proc.* of Design, Automation and Test in Europe, DATE, 2009.
- [6] A. A. Khan, Y. Ohnari, A. Dutta, S. Singh, M. Miura-Mattausch, and H. J. Mattausch, "Die-to-die and within-die fabrication variation of 65nm CMOS technology PMOS transistors," in *Proc. of IEEE International Conference on Electronics, Computing and Communication Technologies, CONECCT*, 2013.
- [7] H. Chang and S. Sapatnekar, "Statistical Timing Analysis Under Spatial Correlations," *IEEE Transactions on Computer-Aided Design of Inte*grated Circuits and Systems (TCAD), vol. 24, 2005.
- [8] B. Hargreaves, H. Hult, and S. Reda, "Within-die process variations: How accurately can they be statistically modeled?" in *Proc. of the Asia and South Pacific Design Automation Conference, ASP-DAC*, 2008.
- [9] A. Mutlu, J. Le, R. Molina, and M. Celik, "A parametric approach for handling local variation effects in timing analysis," in *Proc. of Design Automation Conference, DAC*, 2009.
- [10] "PrimeTime ® Advanced OCV Technology," Synopsys white paper (online), 2009.
- [11] S. Kobayashi and K. Horiuchi, "An LOCV-based static timing analysis considering spatial correlations of power supply variations," in *Proc. of Design, Automation and Test in Europe, DATE*, 2011.