Defocus-Aware Leakage Estimation and Control

Andrew B. Kahng, Swamy Muddu, and Puneet Sharma

***Abstract*—Leakage power is one of the most critical issues for ultradeep submicrometer technology. Subthreshold leakage depends nearly exponentially on linewidth, and consequently, variation in linewidth translates to a large leakage variation. A significant fraction of variation in the linewidth occurs due to systematic variations involving focus and pitch. In this paper, we propose a new leakage-estimation methodology that accounts for focus-dependent variation in the linewidth. Our approach computes the pitch of each device in the design and uses it along with defocus information to predict the linewidth of the device. Once the linewidths of the devices in a cell are calculated, the cell leakage is computed to be the sum of leakages of all off-devices in the cell; device leakages are found from a linewidth-leakage table that is precharacterized with SPICE simulations. The presented methodology significantly improves the leakage estimation and can be used in existing leakage-reduction techniques to improve their efficacy. To demonstrate the use of our approach for leakage reduction, we modify the previously proposed linewidth-biasing technique of Gupta *et al.* to consider the systematic variations in linewidth and further optimize the leakage power. Our method reduces the leakage spread between worst and best process corners by up to 62% compared with the conventional corner-based analysis. Defocus awareness improves the leakage reduction from linewidth-biasing by up to 7%.**

***Index Terms*—Estimation, leakage power, lithography, optimization, topography.**

I. Introduction

LEAKAGE POWER is one of the most critical design concerns in sub-100-nm technology nodes. Decreased supply voltage (and, consequently, threshold voltage) combined with aggressive clock-gating reduces dynamic power but increases leakage power, causing the leakage share of total power to increase. Leakage is composed of three major components: 1) subthreshold leakage; 2) gate leakage; and 3) reverse-biased drain–substrate and source–substrate junction band-to-band-tunneling leakages [4]. In recent technologies, the gate leakage has increased dramatically due to gate-oxide scaling. However, at room temperature, the subthreshold leakage is the dominant contributor to total leakage at the 90-nm technology. Table I shows the subthreshold leakage for Taiwan Semiconductor Manufacturing Corporation's (TSMC) 90-nm technology. At the 65-nm node, the subthreshold leakage is expected to again be the dominant contributor [4], and at the 45-nm node, the use of high-k dielectrics is expected to significantly reduce the gate leakage. With the use of high-k dielectrics, Intel projected a reduction of 100× in gate leakage at 45 nm [9]. Thus, the subthreshold leakage is likely to remain the dominant contributor to the total leakage for foreseeable technologies. In the remainder of this paper, we focus on the subthreshold leakage and refer to it as leakage.

TABLE I
SUBTHRESHOLD AND GATE LEAKAGE OF TSMC 90-nm NOMINAL $V_{\rm th}$ PMOS AND NMOS DEVICES OF 1 μm WIDTH AT TWO TEMPERATURES. SUBTHRESHOLD LEAKAGE IS GREATER THAN GATE LEAKAGE

Device	Temp (^oC)	Leakage (nW)	
		Subthreshold	Gate
PMOS	25	6.45	2.01
NMOS	25	12.68	6.24
PMOS	125	116.80	2.17
NMOS	125	115.90	6.62

Runtime-leakage-reduction techniques explore design trade-offs within performance constraints by identifying candidate devices for optimization using leakage-power estimates. Inaccurate estimation of leakage power can degrade the results of leakage reduction, and hence, accurate estimation of leakage is important. The leakage power increases exponentially with a decrease in linewidth (also known as channel length or gate length). For example, with 90-nm Berkeley Predictive Technology Model (BPTM) device models [1], [7], we observe over 5× and 2.5× increase in leakage for PMOS and NMOS devices, respectively, when the drawn linewidth reduces from 100 to 90 nm. In addition to the leakage power, manufacturers face the additional challenge of leakage variability. Data from [5] indicate that the leakage of microprocessors from a single 180-nm wafer can vary by as much as 20× for a 30% spread in performance. Due to the exponential dependence of the leakage power on linewidth, a small variation in linewidth can result in a significant variation in leakage power.

Traditional leakage-optimization techniques are either oblivious to across-chip linewidth variation (ACLV) or model it as a random variable. This results in a very pessimistic guard-banding and, hence, overdesign. In reality, the ACLV, due to process-variation sources such as focus, exposure, lens aberrations, and mask errors, is partially systematic and can be modeled. All sources of variations that occur during lithography can be lumped into effective focus and exposure dose variations for the purposes of analysis of their impact on the linewidth variation [13].

In this paper, we exploit the systematic variations in the ACLV induced by the focus variations to estimate and optimize

Manuscript received November 15, 2006; revised April 9, 2007 and June 27, 2007. An earlier version of this paper was presented at the International Symposium on Low Power Electronics and Design, San Diego, CA, August 8–10, 2005. This paper was recommended by Associate Editor L. Benini.

A. B. Kahng is with the Department of Computer Science and Engineering and the Department of Electrical and Computer Engineering, University of California at San Diego, La Jolla, CA 92093-0404 USA (e-mail: abk@ucsd.edu).

S. Muddu and P. Sharma are with the Department of Electrical and Computer Engineering, University of California at San Diego, La Jolla, CA 92093-0409 USA (e-mail: smuddu@ucsd.edu; sharma@ucsd.edu).

Digital Object Identifier 10.1109/TCAD.2007.913387

0278-0070/$25.00 © 2008 IEEE

TABLE II
EFFECT OF DEFOCUS AND PITCH (LAYOUT CONTEXT) ON THE LINEWIDTH OF DEVICES IN A CELL, NOR2X2. THE CHANGE IN DEVICE LINEWIDTH WITH DEFOCUS WHEN THE CELL IS IN THE ISOLATED AND DENSE CONTEXTS IS SHOWN. THE CHANGE IN DEVICE LINEWIDTH WITH PITCH AT DEFOCUS VALUES OF 0 AND 100 nm IS ALSO SHOWN. THE DRAWN OR TARGET LINEWIDTH IS 100 nm

Device	Change with defocus (through-focus variation)		Change with pitch (through-pitch variation)	
	Isolated	Dense	0*nm*	100*nm*
M0	6nm	5nm	0nm	1nm
M1	5nm	3.5nm	0.5nm	2nm
M2	2nm	3nm	1.5nm	0.5nm
M3	1nm	3nm	1nm	1nm

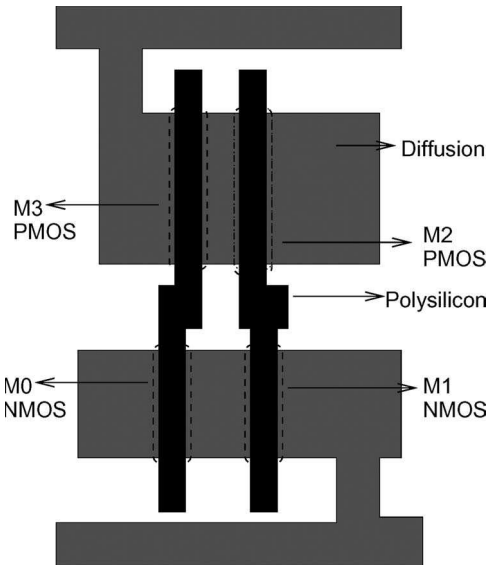

Fig. 1. Layout of the two-input NOR gate in 90-nm technology with poly and diffusion layers only. Devices M0, M1, M2, and M3 are labeled on the layout.

chip leakage power. A similar methodology can potentially be developed to exploit the systematic variations induced by the exposure dose variations. The impact of focus variations on linewidth is strongly dependent on the pitch of the line within the optical radius (OR) of influence. Line pitch is dependent on the physical layout of the circuit, and focus depends primarily on optical column parameters in the wafer stepper. Focus variation, which is also known as defocus, occurs due to wafer topography variation, lens aberrations, tilt of wafer stage during processing, etc. [11]. Among different sources of defocus, topography and lens aberrations are systematic and can be modeled.

Table II shows the change in linewidth of devices in a two-input NOR gate in the 90-nm technology during the following conditions: 1) when defocus is changed from 0 to 100 nm; and 2) when layout environment (referred to as context) surrounding the gate is changed from isolated to dense. The layout of the two-input NOR gate is shown in Fig. 1. Isolated context implies that there are no layout features surrounding the cell under study. This simulates the absence of optical proximity effects from the neighboring layout features. Dense context implies that the cell is surrounded by other layout features. This simulates the significant optical proximity effects that can result in focus-induced linewidth variation. In our experiment, we place four copies of the same cell on all four sides to simulate the dense context.

Columns 2 and 3 in Table II show the linewidth deviation from drawn for all devices in NOR2X2 due to defocus variation for both the isolated and dense contexts. Columns 4 and 5 show the linewidth deviation from drawn for all devices due to change in context for nominal (0 nm) and 100-nm defocus conditions. From the table, we observe that the impact of defocus and pitch on the linewidths of devices in a cell is a large fraction of the total linewidth-variation budget, which is typically about 10% of the drawn linewidth. The corresponding variations in the cell leakage with defocus are 29% and 23% for the isolated and dense contexts, respectively. Due to the presence of a dense layout context, cell leakage variation ranges from 2% at nominal defocus to 7% at 100-nm defocus.

In this paper, we first assess the improvements in leakage estimation that can be obtained by modeling of systematic variations in the linewidth. In the context of standard-cell-based designs, we model the linewidth variation of polysilicon (poly) lines in a placement context by simulating the aerial image-transfer process during lithography after optical proximity correction [(OPC), which is discussed in more detail in Section II]. The OPC is typically performed at the best focus condition, and hence, lithography simulation on the poly lines within standard cells at this focus condition does not show any critical-dimension (CD) variation. At other defocus conditions, however, we see a systematic relationship between layout parameters of the lines and amount of defocus. To predict the leakage of a design, we first analyze its standard-cell layout and extract poly-pitch information. We then use the linewidth model constructed by simulating poly-line patterns along with the defocus map of the design to predict postlithography linewidths. This method does not require a design-level lithography simulation to compute the postlithography linewidths. The predicted linewidths are then used to determine the device and circuit leakages.

The postlithography linewidths of devices within a standard cell depend significantly on the standard-cell layout itself. The interaction between intracell devices and its surrounding optical environment is typically confined to a range of 1 μm. Exact OR, which is the radius of optical influence, depends on the wavelength (λ), the numerical aperture (NA), the partial coherence factor (σ), and the illumination settings (annular, dipole, and quadrupole) of the wafer stepper used in lithography. Since these parameters are fixed for any given process, the OR can be computed exactly.

Our second contribution is to add defocus awareness to enhance a recently proposed leakage-reduction technique, i.e., linewidth-biasing [17]. The linewidth (device gate length) biasing selectively increases the linewidth of devices (which has an effect of making the device slower but less leaky) in cells that are not on timing-critical paths. Defocus awareness enables the linewidth-biasing to positively bias any cell instances for which devices are likely to print with a smaller linewidth and be extremely leaky. With our modifications, the linewidth-biasing achieves an improved leakage reduction.

In summary, the main contributions of this paper are the following:

1) modeling of layout- and defocus-dependent systematic components of linewidth variation to better predict the leakage;
2) defocus-aware linewidth-biasing that models systematic linewidth variation for an improved leakage reduction.

Previous variation-aware leakage-analysis methods have focused on statistical analysis (e.g., [8] and [24]). Compared with the traditional corner-case-based methods, statistical approaches yield a more accurate and less pessimistic analysis. The approaches propose mathematical frameworks using which leakage distributions can be found, given the distributions of process variations and the dependence of leakage on them. These approaches assume process-variation distributions to be given. Linewidth is assumed to be one of the random variables, and systematic variations in linewidth are modeled using spatial correlations. Such frameworks cannot satisfactorily capture the ACLV which is highly context-dependent. In the absence of suitable statistical frameworks, and for simplicity and easier adoptability, we perform our analysis deterministically.

The remainder of this paper is organized as follows. In Section II, we present background on ACLV and discuss topography simulation. Section III describes our defocus-aware leakage-estimation methodology and presents its results. Defocus-aware linewidth-biasing methodology and its results are presented in Section IV. Section V concludes and summarizes this paper.

II. Background

In this section, we present background information on lithography and sources of linewidth variation.

1) **NA** of the lens used in wafer-patterning is the sine of the maximum half-angle of light that can make it through a lens, which is multiplied by the index of refraction of the media. The NA of a lens is a measure of its ability to capture and focus light from the mask across a wide range of angles.[1] The OR of interaction is inversely proportional to NA.
2) **Partial coherence factor** (σ) is the ratio of the sine of the maximum half-angle of illumination striking the mask to the NA of the objective lens.[2] The partial coherence factor determines the degree of spatial coherence between any two points on the objective lens. If the degree of spatial coherence is high, then the partial coherence factor is low. Typical values of partial coherence factor in optical lithography range from 0.3 to 0.9. The value of the partial coherence factor determines the extent of optical interaction between features. The higher the value of σ, the less the optical-interaction range.
3) **Off-axis illumination (OAI)** refers to the illumination which has no on-axis component, i.e., which has no light that is normally incident on the mask. Examples of OAI include annular and quadrupole illuminations. In subwavelength optical lithography, the OAI is commonly used.
4) **OPC** is a method of selectively changing the sizes and shapes of patterns on the mask in order to more exactly obtain the desired printed patterns on the wafer. The OPC is the most prominent resolution enhancement technique to enable patterning of layout features in the subwavelength optical lithography.
5) **OR** of influence refers to the optical-interaction range between features. If a layout feature is within the optical radius of another feature, then each feature gets influenced by the other during image transfer in optical lithography (due to diffraction effects). The value of OR depends on the NA, the σ, and the illumination settings chosen.
6) **Subresolution assist features (SRAFs)** or scattering bars (SBs) are layout features that are inserted between the poly lines in the layout to improve their printability. The SRAFs have linewidths that are less than the drawn linewidth of layout features, and hence, they do not print on the wafer. However, they enhance the printability of layout features by changing the diffraction pattern of light through the mask.
7) **Defocus** is defined as the distance, which is measured along the optical axis (i.e., perpendicular to the plane of the best focus) between the position of a resist-coated wafer and the position if the wafer was at the best focus. The defocus results in blurring of the image transferred onto the wafer and, consequently, translates to the linewidth variation. The extent of variation depends on the line pitch.

The ACLV due to defocus is one of the most significant contributors to linewidth and, consequently, to leakage-power variability. The OPC enables the control of variations due to pitch and focus and is mandatory before mask-manufacturing in very large scale integration design flows at the present technology nodes. The OPC controls printed shapes on wafer by applying corrections to design features based on proximity effects at nominal defocus conditions. Although the OPC is extremely effective in linewidth control at nominal defocus conditions, it can result in significant linewidth variation at other defocus conditions. This variation, which is caused by interaction between the pitch and the defocus, is systematic and can be modeled, predicted, and compensated.

The Bossung plot in Fig. 2 shows the variation of post-OPC printed linewidth at different pitch and defocus conditions. We observe that dense lines tend to "smile" with defocus, whereas isolated lines "frown." The decrease in linewidth for isolated patterns is greater than the increase in linewidth for dense patterns over the same range of defocus despite the use of SBs (or SRAFs) [6] in our OPC recipes. Hence, the isolated lines tend to become more leaky than their dense counterparts. Defocus is caused by several sources such as variation in shallow-trench-isolation (STI) layer thickness during chemical–mechanical

[1] Light from the mask gets diffracted through the gratings (i.e., transparent regions between features). The diffracted pattern is captured by the reduction lens to focus on the wafer. The NA of the lens determines the number of diffraction orders captured by the lens.

[2] Objective lens focuses light that is passing through the mask onto the wafer.

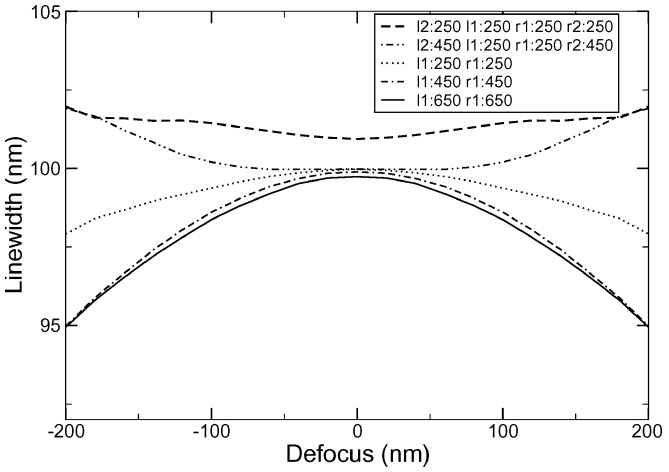

Fig. 2. Dependence of linewidth on defocus for patterns with different pitches. Each curve in the plot corresponds to a specific device-pitch configuration. The value of $l1$ ($r1$) specifies the distance to the first left (right) neighbor. The value of $l2$ ($r2$) specifies the distance between the first left (right) neighbor and the second left (right) neighbor. The densest configuration corresponds to a pitch of 250 nm between all poly lines, and the sparsest configuration corresponds to poly pitch of 650 nm for a three-bar test pattern. The linewidth increases with defocus for the dense patterns, and it decreases for the isolated patterns.

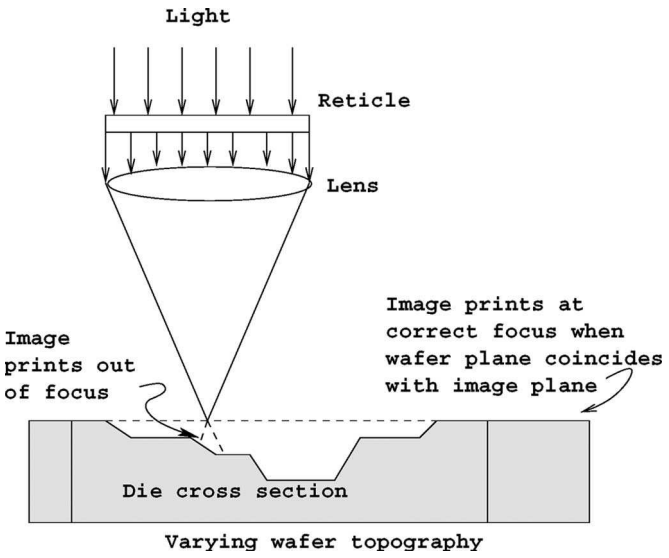

Fig. 3. Vertical cross section of a wafer showing topography nonuniformity. Focus variation due to nonplanar wafer topography is illustrated. Substrate thickness due to CMP effects varies according to the density of the STI (active) layer. Lens aberrations, misalignment of wafer-plane and lens axes, and variation in distance between the lens and wafer planes (stage error) add to the defocus.

planarization (CMP), lens aberrations, wafer-stage misalignment, and resist thickness variation. The linewidth variation caused by defocus due to thickness variation can be modeled systematically by layout density analysis. For a particular value of defocus, the linewidth depends on layout proximity effects.

A schematic of topography-dependent defocus during lithography is shown in Fig. 3. If the image plane of the reticle and lens system coincides with the wafer plane, the image prints with high resolution. However, in the regime of topography variation, which is caused predominantly by erosion and dishing effects during the CMP, the image prints out of focus, leading to topography-dependent linewidth variation. Other

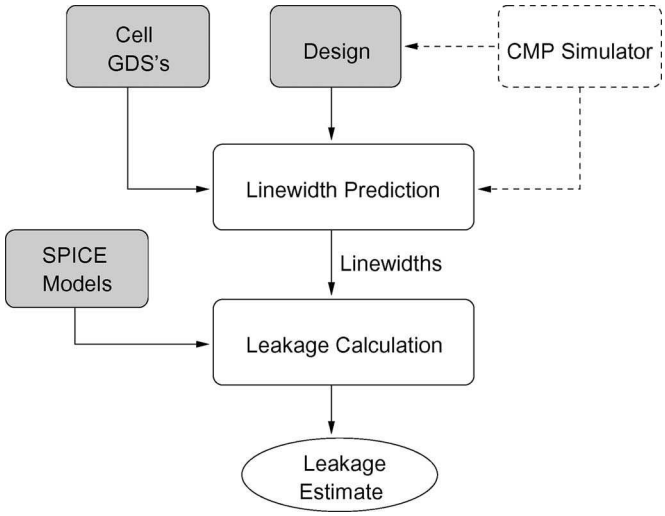

Fig. 4. Our defocus-aware leakage-estimation methodology.

optical and mechanical effects, such as wafer-stage misalignment, substrate flatness, and field tilt variation, result in additional variations that are random or difficult to model.

Topography simulation was the focus of several recent papers, e.g., [20] and [22]. These works present and calibrate analytical models that account for the underlying pattern and various CMP process parameters, such as planarization length, pad bending, slurry selectivity, etc., to predict the post-CMP thickness variation at all locations of a chip. Since the CMP simulation is a complex task involving several process parameters which may not be available, we also propose alternative analysis and optimization flows that do not rely on the CMP simulation and just consider the sensitivity of linewidth variations to defocus. In our experiments, we assume that a full-chip topography map is given as input.

III. Defocus-Aware Leakage Estimation

Our defocus-aware leakage-estimation methodology is composed of two modules: 1) linewidth prediction; and 2) leakage calculation. Fig. 4 shows the methodology. The linewidth-prediction module uses placement information of the design along with the locations of devices within each cell in the cell library [from the cell Graphic Data Systems (GDSs)] to compute the pitches of all devices in the design. It then uses the Bossung table, which captures systematic variation of linewidth with defocus and pitch, to compute the linewidths of all devices. The leakage-calculation module computes the leakage of all devices given their linewidths and finds the leakage of the design. We propose the following two flows depending on the availability of die-topography information.

1) Defocus-aware topography-oblivious flow. We do not rely on a CMP simulator and assume the defocus (due to topography and other sources) to be random. In this flow, we use the fact that the linewidth variation is greater for devices with dense or sparse pitches. Devices that have medium pitches, or high pitch on one side and sparse pitch on the other, are self-compensating and print with less linewidth variation.

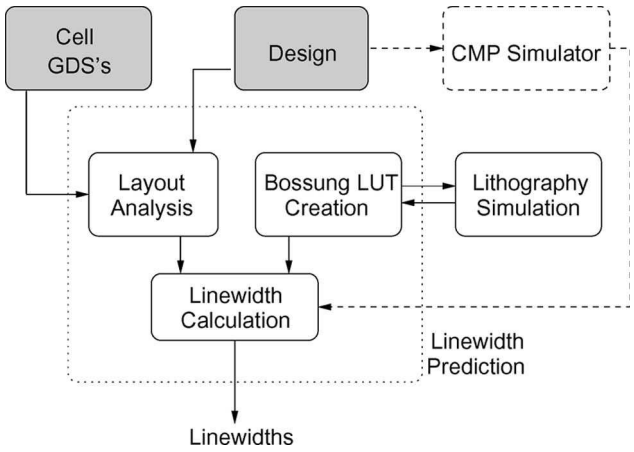

Fig. 5. Proposed linewidth-prediction flow.

2) Defocus-aware topography-aware flow. In this flow, we consider a topography map that is available from a CMP simulator. Since topography is a significant contributor to defocus variation, an improved topography prediction leads to an improved defocus prediction and, consequently, a better leakage estimation.

A. Defocus-Aware Linewidth Prediction

Both our flows analyze the layout context of each device of the design and use it with the defocus (assumed to be completely or partly random depending on the flow) at that device location to predict its linewidth. Since leakage is only affected by the dimensions of the gate in MOS devices, we are only interested in linewidth prediction of the gate regions (i.e., overlap of diffusion and poly). The gate regions are always rectangular and are always spaced by the minimum design rules from complex shapes such as line ends and poly bends. Therefore, we can expect a linewidth-prediction method that is significantly simpler and faster than the lithography simulation to be reasonably accurate.

The main components of the linewidth prediction are the following: 1) Bossung lookup-table (LUT) generation and 2) layout analysis for pitch calculation for each device. Fig. 5 shows the linewidth-prediction methodology. The Bossung LUT creation performs the lithography simulation to capture and tabulate the linewidth variation with pitch and defocus. Layout analysis calculates the pitch of each device in the design by analyzing the placement and standard-cell layouts.

1) Bossung LUT Creation: The Bossung LUT captures systematic variations in linewidth due to pitch and defocus. The Bossung LUT creation is an offline process that needs to be done only once for a given cell library and process technology. To create the LUT, we construct line-and-space patterns of gate poly with different spacings to simulate different pitches. The linewidth of gate poly in each pattern is fixed at 100 nm, which corresponds to the linewidth of TSMC 90-nm technology. Line-to-line spacing is varied from 150 nm (the minimum spacing at this technology node) to 750 nm in steps of 100 nm on both sides. In each pattern, there is one gate-poly feature that we call the "poly of interest" with two identical neighbors on each side at various spacings to get a total of five features in each pattern. Next, for each pattern, neighbors that are away from the poly of interest by more than 800 nm are removed. It is safe to discard distant neighbors because the 193-nm steppers used for patterning features in the 90-nm technology node have an optical radius of approximately 600 nm (i.e., features separated by more than 600 nm have a negligible impact on each other). We conservatively use 800 nm as the optical radius for all our experiments. We utilize symmetry of patterns to significantly cut down their number to a total of 153.

After the creation of the line-and-space patterns, we perform the OPC of the patterns with zero defocus using Calibre OPC. To measure the linewidth variation due to defocus, we then perform the lithography simulation at different defocus levels for all the patterns. We choose defocus values in the range of (−200 nm, 200 nm) in steps of 20 nm. Poly linewidth values are then extracted from all simulated printed images at each defocus level. In order to perform the OPC and the lithography simulation, we construct a model describing the optical characteristics of wafer stepper and resist-coating on wafer. The optical and resist model files are input to the OPC and the litho simulator (e.g., Calibre OPC). The optical model files are generated by specifying the NA, the partial coherence factor, the defocus, and the illumination settings. For our current experimental setup, we generated optical model files for each defocus level with an NA of 0.7 in Calibre WorkBench and set the resist threshold to 0.38; both values fall in their standard ranges for 90-nm OPC setup.

Our Bossung LUT contains rows corresponding to patterns and columns corresponding to defocus values. Entries in the table give the printed linewidth values for the feature of interest in the pattern. Linewidth change with defocus for different patterns is shown in Fig. 2. For dense patterns, we observe the linewidth to increase by up to 2 nm. For sparse or isolated patterns, on the other hand, we observe a reduction in linewidth of up to 6 nm. These observations are in line with previously reported trends [15].

2) Layout Analysis: Given the defocus and the pitch for a device, the Bossung LUT can be used to predict its printed linewidth. While the defocus is assumed to be completely or partially random, depending on the flow and as described in the experimental setup section, the pitch is computed by layout analysis. The pitch of a device is composed of two distances: 1) spacing between its right edge and the left edge of the nearest device to its right; and 2) spacing between its left edge and the right edge of the nearest device to its left.

Fig. 6 shows the pitch calculation for two devices (A0 and B0) of three neighboring cells with the intercell and device-to-boundary distances. Spacing between devices of a cell can be easily computed by taking the difference between their respective device-to-boundary spacings for a given boundary. We note that spacings between devices that belong to the same cell need to be computed only once for each standard-cell master in the cell library. Spacing computation between devices of different cells involves adding the intercell distance between the two cells and the distance of the two devices from their corresponding cell boundaries, with a careful consideration of the cell orientations. Device-to-boundary spacings are available

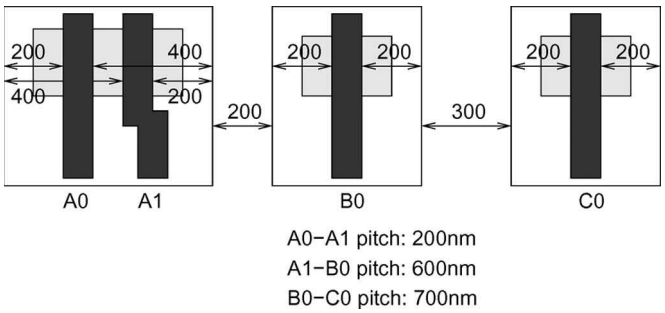

Fig. 6. Pitch computation from a design layout. Nominal linewidth of features is 100 nm.

TABLE III
DIFFERENCE BETWEEN LINEWIDTH GIVEN BY LITHOGRAPHY SIMULATION AND THAT PREDICTED BY OUR APPROACH IS SHOWN FOR DEVICES IN SEVEN CELLS, ACROSS THREE DEFOCUS VALUES (100, 200, AND 250 nm). THE DRAWN LINEWIDTH IS 100 nm

Cell	Device	Discrepancy		
		100nm	200nm	250nm
CLKINVX1	M0	0	0	0
	M1	-1	0	0
INVX2	M0	-2	-1	0
	M1	-1	0	0
NAND2X1	M0	0	0	1
	M1	-1	0	-1
	M2	-1	0	0
	M3	0	0	0
NOR2X1	M0	0	0	1
	M1	-1	-1	-1
	M2	-1	-1	-1
	M3	-1	0	-1
NAND2X2	M0	-1	0	-1
	M1	-1	-1	-1
	M2	-1	-1	0
	M3	-1	-1	-1
NOR2X2	M0	0	-1	0
	M1	-1	-1	-1
	M2	-1	-1	0
	M3	-1	-1	0
NOR4X2	M0	-1	0	-1
	M1	-2	0	0
	M2	-2	-1	-2
	M3	-1	0	-1
	M4	2	0	1
	M5	-2	-2	-2
	M6	0	-2	-1
	M7	-1	0	-1

from the cell GDSs after performing layout-versus-schematic to annotate device gate poly with device names. Information about neighboring cells, boundary-to-boundary spacings, and cell orientation can be found from the placement of the design.

3) Validation: Since the leakage calculation depends on the linewidth prediction, it is important to validate our linewidth-prediction flow and assess its accuracy. To this end, we compare the linewidths predicted by our method with the linewidths predicted by lithography simulation on individual cells in the isolated and dense contexts at different defocus values. For the isolated context, we assume no neighbors of the cell. For the dense context, we assume the cell to be surrounded by four

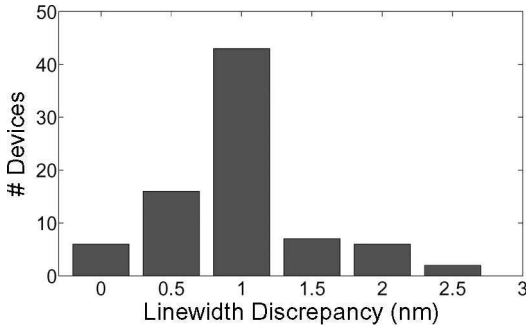

Fig. 7. Distribution of average CD discrepancy of all devices in ten standard cells. Average CD discrepancy of each device is computed by taking the mean of CD difference across all defocus conditions.

TABLE IV
LEAKAGE DISCREPANCY, WHICH IS THE DIFFERENCE IN LEAKAGE ESTIMATES WHEN LINEWIDTHS ARE COMPUTED WITH LITHOGRAPHY SIMULATION AND WITH OUR APPROACH, IS SHOWN FOR TEN STANDARD CELLS

Cell	% change in leakage power		
	100*nm* defocus	200*nm* defocus	250*nm* defocus
CLKINVX1	0.57	0.00	0.00
INVX2	3.27	1.20	0.00
MX2X1	2.74	0.55	2.82
NAND2BX1	3.08	0.59	2.37
NAND2X2	2.75	2.89	2.36
NOR2X1	1.85	0.88	1.56
NOR2X2	1.50	3.64	1.31
NOR4X2	3.04	0.62	1.88
OA21X2	0.46	0.39	0.42
OR3X2	1.31	1.19	1.90

copies of itself, one on each side. We create the Bossung LUT to have pitches with right and left spacings in the range of 150–750 nm in steps of 100 nm, as described in Section III-A. If the Bossung LUT does not contain the exact device-pitch configuration that exists in the layout, then we snap to the closest configuration in the LUT. This can result in inaccuracy between actual printed CD and predicted CD from the LUT. This source of inaccuracy can be reduced at the cost of the Bossung LUT creation time. We define discrepancy as the difference in device linewidth given by lithography simulation and that predicted by our approach. Table III shows the discrepancy for all devices in seven cells in the isolated context and for defocus values of 100, 200, and 250 nm. Discrepancy, which is averaged over the three defocus values, is shown as a histogram in Fig. 7. From the table and the figure, we observe that the discrepancy is typically 1 nm and under 2 nm for most devices. This validates our linewidth-prediction methodology.

Since the average discrepancy is less than 2 nm, the error in leakage estimation is not expected to be large. To analyze the impact of linewidth discrepancy on leakage power of cells, we perform SPICE simulations to estimate the cell leakage with linewidths from lithography simulation and those from our approach. We define the leakage discrepancy to be the difference in leakage estimates when the linewidths are computed with lithography simulation and with our approach. The leakage discrepancy for ten cells in our library is shown in Table IV. From the table, we observe that the leakage discrepancy is

negligibly small for most cells across multiple defocus conditions. This inaccuracy is a small penalty, considering the huge runtime savings obtained by using our approach instead of the lithography simulation.

B. Defocus-Aware Leakage Calculation

We have adapted the methodology proposed by Rao *et al.* [23] to compute the cell gate leakage to calculate the cell subthreshold leakage. The subthreshold leakage in a PMOS (NMOS) device occurs only when the gate terminal is in the high (low) state and the source and drain terminals are in opposite states. For each state applied to the inputs of a cell, we propagate the states to all internal terminals of the cell and find the leaky devices. To calculate the leakages of the leaky devices, we use an LUT, which is characterized with SPICE simulations, that gives the leakages of NMOS and PMOS devices for different linewidths (gate lengths) that we are likely to encounter. We then sum the leakages of all leaky devices to find the cell leakage for the state. To calculate the average cell leakage, we average the cell leakage over all states; if state probabilities are available, an average weighted by the state probabilities improves the accuracy.

In our cell-leakage methodology, we ignore the leakage of stacked devices since it is in the orders of magnitude less than that of nonstacked devices due to self-reverse-biasing of stacked devices [18]. Narrow-width effects can be accounted for by characterizing the leakage LUT for multiple device widths along with multiple linewidths. To compute the design leakage, we sum up the leakages of all cells. With respect to SPICE, our approach has a cell leakage estimation error of less than 5% for all cells in our library. Rao *et al.* [23] also reported similar maximum estimation error for the gate leakage.

C. Experimental Study

Having validated the accuracy of the linewidth-prediction and the leakage-estimation flows, we now assess the improvement in circuit leakage estimation from our flow with respect to the traditional corner-based flow.

1) Experimental Setup: We perform our experiments on the following circuits: c5315 (2077 cells), c6288 (4776 cells), and c7752 (3155 cells) from the ISCAS'85 test suite, and alu128 (11 724 cells) from opencores.org. The circuits were synthesized using Synopsys Design Compiler v2003.06-SP1 using a small standard cell library of 20 cells under tight delay constraints. Our library is composed of the 20 most frequently used cells in our test cases.[3] To create the Bossung LUT, we use Mentor Calibre v9.3_5.9 for the OPC and the lithography simulation. Our industry-strength OPC and lithography-simulation recipes are for 100-nm linewidths using 193-nm stepper. We insert SBs (assist features) to improve the process window. We use Synopsys HSPICE vU2003.09 for all our SPICE simulations and Cadence SignalStorm v4.1 for library characterization with BPTM BSIM3 SPICE models [1], [7].

[3]To identify the most frequently used cells, we first synthesize our test cases using the entire TSMC 90-nm standard cell library.

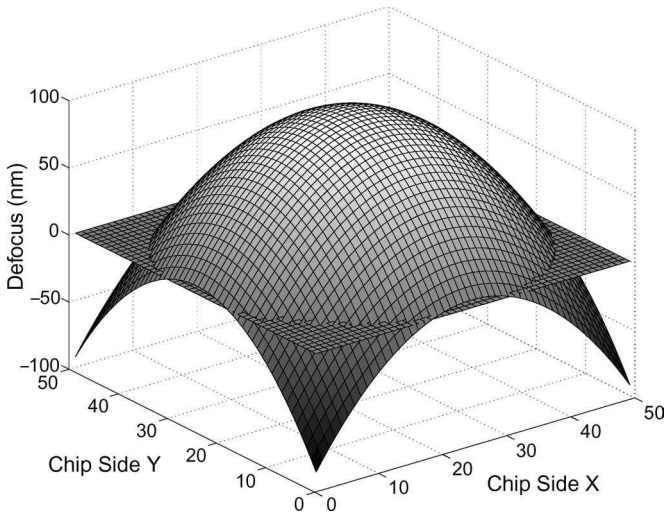

Fig. 8. Die topography used in our experiments. Maximum height is 100 nm higher than nominal (illustrated by the plane) at the center and decreases quadratically with distance from the center to become 100 nm below nominal at the die corners.

Temperature and voltage were assumed to be 25 °C and 1.2 V, respectively, in all experiments. We place the designs with Cadence SOC Encounter v3.2.

We compare the following: 1) the traditional flow; 2) the proposed defocus-aware topography-oblivious flow; and 3) the proposed defocus-aware topography-aware leakage-estimation flow. Traditional leakage estimation is corner-based and assumes devices to have the smallest, nominal, and largest linewidths for the worst, nominal, and best cases, respectively. The flow involves library characterization with a tool such as Cadence SignalStorm [3] to calculate the leakages of all cells in the library with SPICE simulations. Then, a circuit-level leakage-analysis tool such as Synopsys PrimeTimePX [2] sums the leakage of all cells in the design to calculate the design leakage. In the comparisons of the three flows, we consistently assume the smallest, nominal, and largest linewidths to be 86, 100, and 110 nm, respectively.

In the defocus-aware topography-oblivious leakage estimation, we assume defocus to be random with a Gaussian distribution ($\mu = 0$ nm and $\sigma = 66$ nm), leading to a 3σ value of 200 nm. Flagello *et al.* [13] use 3σ defocus of 300 nm for their study, and $\pm$200-nm defocus is considered reasonable. The focus variation assumed for our experimental setup changes between processes and can improve as the process matures. Since variations cannot be completely mitigated, the proposed methodology can be used across any range of focus settings. The assumed defocus of $\pm$200 nm induces a linewidth variation between -6 and $+2$ nm. Since the linewidth variation is caused by factors other than defocus, such as mask errors and exposure variations, we assume a random variation of $\pm$8 nm in linewidth from other sources. Thus, the contribution of linewidth variation due to defocus is 1/3 of the total linewidth variation.[4] Our assumptions are in line with the findings of Flagello *et al.* [13].

[4]It is not appropriate to find the standard deviation in linewidth due to the two sources by the "square root of sum of squares" method because contribution due to defocus is partly modeled by our approach, and the remainder is not close to Gaussian.

TABLE V
ESTIMATED LEAKAGE POWER AT WORST, NOMINAL, AND BEST PROCESS CORNERS USING TRADITIONAL, TOPOGRAPHY-OBLIVIOUS DEFOCUS-AWARE, AND TOPOGRAPHY-AWARE (ASSUMING THE TOPOGRAPHY OF FIG. 8) DEFOCUS-AWARE LEAKAGE-ESTIMATION FLOWS. LEAKAGE VALUES WHEN THE ENTIRE CIRCUIT USES ONLY LOW $V_{\rm th}$ DEVICES AND WHEN IT USES ONLY NOMINAL $V_{\rm th}$ DEVICES ARE SHOWN

Circuit	V_{th}	Traditional			Defocus-Aware, Topography-Oblivious				Defocus-Aware, Topography-Aware			
		WC (mW)	Nom (mW)	BC (mW)	WC (mW)	Nom (mW)	BC (mW)	Spread Reduction	WC (mW)	Nom (mW)	BC (mW)	Spread Reduction
c5315	Low	8.006	0.956	0.304	5.269	0.853	0.337	35.96%	4.119	0.889	0.337	50.90%
	Nom	1.481	0.125	0.036	0.931	0.111	0.040	38.34%	0.675	0.116	0.040	56.06%
c6288	Low	19.540	2.308	0.726	15.298	2.158	0.838	23.14%	11.256	2.265	0.838	44.63%
	Nom	3.625	0.302	0.086	2.827	0.282	0.101	22.97%	1.897	0.297	0.101	49.25%
c7552	Low	12.327	1.469	0.465	9.541	1.360	0.533	24.06%	7.126	1.433	0.533	44.42%
	Nom	2.281	0.192	0.055	1.757	0.177	0.064	23.94%	1.203	0.188	0.064	48.83%
alu128	Low	48.499	5.771	1.826	27.264	4.985	1.987	45.84%	22.442	5.153	1.987	56.17%
	Nom	8.978	0.754	0.217	4.574	0.644	0.238	50.51%	3.577	0.668	0.238	61.89%

For the defocus-aware topography-aware flow, we assume the topography shown in Fig. 8 as an input. The topography height is 100 nm at the center of the die and quadratically reduces with a distance from the center to become −100 nm at the die corners.[5] A topography variation of ±100 nm is within the defocus tolerance and expected to exist [14]. In practice, the topography should be predicted by a CMP simulator that models STI layer planarization (STI-CMP simulator) such as those developed in [20] and [22]. We again assume the defocus to be ±200 nm but consider only half of the defocus (±100 nm) to be random, with the other half being determined from the input topography which alters the defocus by up to ±100 nm.

2) Results: Table V shows the leakage estimation for all three leakage-estimation flows. We observe that the leakage spread between the best and worst process corners is the largest for the traditional leakage-estimation flow. The two defocus-aware flows reduce the spread by decreasing the worst case leakage and increasing the best case leakage. The defocus-aware topography-oblivious flow reduces the spread (despite assuming the defocus to be completely random) because it reduces the pessimism in leakage estimation for the cells that are less sensitive to focus variation due to the systematic nature of linewidth with focus. The defocus-aware topography-aware flow utilizes the additionally available defocus information to further reduce the leakage spread.

Leakage at the best case process corner is expected to be the highest for the defocus-aware topography-aware flow. However, we note that the leakage values for the best case process corner are identical for the two defocus-aware flows. This happens because the leakage decreases nearly exponentially with linewidth, and at large linewidths (used for the best case corner), changes in linewidth cause much smaller leakage changes. The difference between the linewidths estimated by the two defocus-aware flows is not sufficiently large to register any significant leakage difference at the large linewidths used in the best case corner. The nominal leakage for the three flows is not directly comparable as it depends on the process and assumed topography.

[5]The topography used in Fig. 8 is not unrealistic. Designs that have low device density in the center and high device density toward edges can result in the shown topography.

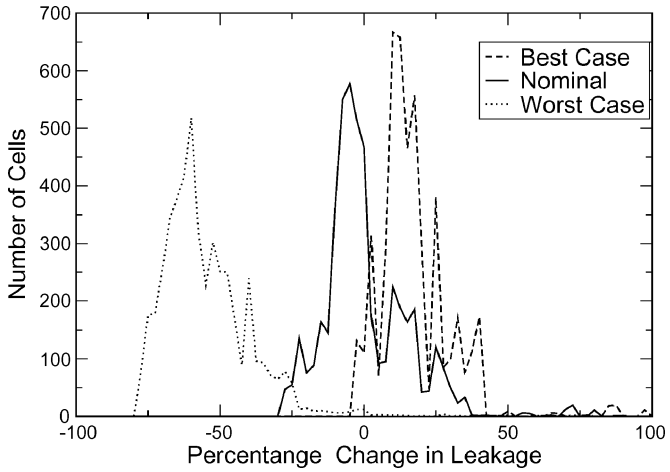

Fig. 9. Distribution of percentage change in leakage estimated with the defocus-aware topography-aware flow with respect to the traditional flow for test case c6288 for the three corners. For the nominal corner, the change in total circuit leakage is just −1.86% (traditional is higher), but individual cells have larger change.

In addition to the accurate design leakage estimation, our methodology predicts the individual cell (or device) leakages for each cell (or device) more accurately. Fig. 9 shows the distribution of the difference between the cell leakage predicted by the defocus-aware topography-aware flow with respect to the traditional method for test case c6288 for the best case, nominal, and worst case corners. While we observe large cell leakage estimation errors in the range of −29%–124% for the nominal corner, the error in overall circuit leakage estimation is only −1.86%. Our improved cell leakage prediction can be used to improve the quality of leakage-reduction techniques that selectively optimize the cells (or devices) with high leakage, such as input-vector control, $V_{\rm th}$ assignment, and linewidth-biasing.

IV. DEFOCUS-AWARE LINEWIDTH-BIASING

In this section, we begin with a primer on the previously proposed technique of linewidth-biasing for runtime leakage and its variability reduction [17]. We then describe our methodology to add defocus awareness to the linewidth-biasing and present results.

TABLE VI
LEAKAGE POWER AFTER THE TRADITIONAL AND THE DEFOCUS-AWARE LINEWIDTH-BIASINGS. LEAKAGE OPTIMIZATION IS DONE FOR THE NOMINAL PROCESS CORNER AND THE TOPOGRAPHY OF FIG. 8

Circuit	Traditional Linewidth Biasing			Defocus-Aware Linewidth Biasing			Leakage Reduction		
	WC (mW)	Nom (mW)	BC (mW)	WC (mW)	Nom (mW)	BC (mW)	WC (%)	Nom (%)	BC (%)
c5315	3.948	0.855	0.326	3.838	0.838	0.321	2.78	2.01	1.63
c6288	9.363	1.923	0.730	8.958	1.861	0.712	4.33	3.23	2.56
c7552	6.678	1.350	0.507	6.212	1.280	0.485	6.98	5.17	4.21
alu128	21.258	4.908	1.907	19.968	4.663	1.827	6.07	4.99	4.19

A. Traditional Linewidth-Biasing

Linewidth-biasing exploits the fact that the leakage reduces exponentially, whereas the delay increases only linearly with an increase in linewidth. To have a minimal impact on circuit delay, the technique selectively biases only the devices that belong to cells that are not on timing-critical paths. Biasing a cell increases its delay and may cause some noncritical paths to become critical and, consequently, prevent other cells on the new critical paths from getting biased. Thus, the ordering in which cells are biased affects the quality of leakage optimization. The authors use a sensitivity-based greedy solution in which cells are iteratively biased in the order of their decreasing sensitivity. The sensitivity is defined as the ratio of leakage reduction and delay increase of a cell caused by biasing. If biasing a cell causes a timing violation, the cell is unbiased (i.e., its linewidth is set back to nominal). The algorithm continues until no more cells can be biased. Sensitivity-based algorithms have also been used for $V_{\rm th}$ assignment [25] and gate-width-sizing [12]. We improve the traditional linewidth-biasing by more accurately estimating the leakage using our defocus-aware estimation flow. This facilitates more accurate sensitivity calculation and better leakage optimization.

B. Methodology

We use the following terminologies to explain our modifications to the sensitivity function for linewidth-biasing.

1) L_p represents the leakage of cell instance p, $L_p^{\rm n}$ represents its leakage at the nominal process corner, and $\langle L_p \rangle$ represents its expected leakage.
2) L_{pt} represents the leakage of the tth device of cell instance p, and $L_{pt}^{\rm n}$ and $\langle L_{pt}^{\rm n} \rangle$ are its nominal process corner and expected leakages, respectively ($L_p = \sum_i L_{pt}$, where the summation is taken over all devices of the cell).
3) $\Delta L_p^{\rm n}$ and $\Delta \langle L_p \rangle$ represent the change in nominal and expected leakages due to biasing cell instance p (i.e., biasing all devices in cell instance p).
4) Δd_p is the change in delay of cell instance p after biasing it at the nominal process corner.

The sensitivity S_p in traditional linewidth-biasing is the ratio between the leakage reduction and the delay increase of cell p upon biasing, which is given by

$$S_p = \frac{\Delta L_p^{\rm n}}{\Delta d_p}. \tag{1}$$

The sensitivity in the defocus-aware leakage estimation is given by

$$S_p = \frac{\Delta \langle L_p \rangle}{\Delta d_p}. \tag{2}$$

To compute the expected leakage, we have two flows that are similar to the flows used for the defocus-aware leakage estimation and that depend on the availability of the topography simulation. For the defocus-aware topography-aware flow, we assume the defocus to be a Gaussian random variable centered at the topography height given as an input from the STI-CMP simulator and have a 3σ of 100 nm (50% of our defocus-variation budget). For the defocus-aware topography-oblivious flow, we consider defocus variation to be completely Gaussian random with a mean of 0 nm and 3σ of 200 nm. We model the leakage as a function of linewidth which, in turn, is a function of pitch and defocus. Therefore

$$L_{pt} = L(\ell\,(D_{pt}, P_{pt})) \tag{3}$$

where D_{pt} and P_{pt} are the defocus and the pitch for device t of cell p, respectively, and $\ell(D_{pt}, P_{pt})$ represents its linewidth. We may now write the expected leakage as

$$\langle L_p \rangle = \sum_t \langle L_{pt} \rangle \tag{4}$$

$$\langle L_{pt} \rangle = \sum_t \sum_{D_{pt}} L\,(\ell(D_{pt}, P_{pt})) \cdot P(D_{pt}) \tag{5}$$

where $P(D_{pt})$ is the probability that D_{pt} is the defocus value.

C. Results

A comparison between the traditional and the defocus-aware (topography-aware) linewidth-biasings is presented in Table VI. While we assume only defocus to be random during optimization (to exploit the systematic dependence of linewidth on the defocus and the pitch), we present results for the three process corners, as described in Section III. The delay penalty for linewidth-biasing is set to 0% (i.e., it is a constraint that the circuit delay does not increase after biasing). The runtime penalty due to defocus awareness is under 10% for all our test cases.

Our results show modest leakage reductions for all three process corners from 1.63% to 6.98%. However, given that we have made only minor changes to the sensitivity function of linewidth-biasing, we consider these results encouraging. Our

approach may be used with several other leakage-optimization approaches that rely on identifying candidate cells or devices to make tradeoffs. Larger leakage reductions are expected when the impact of systematic linewidth variations on gate delays is also considered during optimization. Slacks, which are created when pessimism in delays is reduced by systematic variation-aware timing analysis, can be used toward leakage reduction. The extent of leakage reduction depends on the reduction in pessimism and the effectiveness of the leakage reduction knob to tradeoff delay versus leakage.

V. Conclusion

Due to the exponential dependence of leakage on linewidth, pessimism in linewidth translates to large leakage pessimism and overdesign. There is a need to model systematic components of linewidth variation for an improved leakage estimation.

Our leakage-estimation methodology models the pitch- and defocus-dependent systematic components of the linewidth variation. We analyze a layout to calculate device pitches and to use them with defocus and a precharacterized Bossung LUT to predict printed linewidths and to estimate leakage with an increased accuracy. Our defocus-aware topography-oblivious flow does not rely on an STI-CMP simulator and assumes defocus variations to be random. It considers device pitches to predict linewidth and, consequently, leakage with an improved accuracy. The defocus-aware topography-aware flow uses the STI-CMP simulation to better predict the defocus variation to further improve the leakage estimation. Our methodology reduces the spread between the leakage estimation at the worst and best process corners by over half and can estimate leakages of individual devices with an improved accuracy.

Leakage-optimization techniques that rely on leakage estimation of individual cells or devices can benefit from the defocus-aware leakage-estimation flow. We enhance the previously proposed linewidth-biasing methodology that relies on the leakage estimation of individual cells to determine the order in which cells are biased. The defocus-aware linewidth-biasing has larger leakage reductions than the traditional linewidth-biasing by 2%–7% on our test cases.

Our ongoing work explores several ways to improve leakage-estimation accuracy and to apply improved leakage estimation to enhance other leakage-reduction techniques. Defocus variation has a significant systematic component arising from the lens aberrations during wafer processing [13]. Modeling lens aberrations deterministically will reduce pessimism in current guardbanding of inter- and intradie variations.

Across-chip focus-dependent linewidth variation is partly due to the complex optical interactions between the layout features. To overcome this effect, foundries are proposing the use of regular layouts that restrict designer freedom. Restricted design rules (RDRs) result in better printability at the cost of device density. The RDRs allow the creation of highly regular layouts. However, even with the use of RDRs, layout patterns are still susceptible to focus variations. As the gap between the minimum feature dimension and the lithography wavelength (currently at 193 nm) increases in technology nodes below 90 nm, intracell optical interactions increase significantly. These interactions translate to linewidth variation even with RDR layouts. Hence, the methodology proposed in this paper is relevant for technologies that use RDRs.

References

[1] *Berkeley Predictive Technology Model*. [Online]. Available: http://www.device.eecs.berkeley.edu/~ptm

[2] *Synopsys PrimeTimePX*. [Online]. Available: http://www.synopsys.com/products/solutions/galaxy/power/power.htm

[3] *Cadence SignalStorm*. [Online]. Available: http://www.cadence.com/datasheets/6256_LibChar_TP_v2.pdf

[4] A. Agarwal, C. H. Kim, S. Mukhopadhyay, and K. Roy, "Leakage in nano-scale technologies: Mechanisms, impact and design considerations," in *Proc. ACM/IEEE Des. Autom. Conf.*, 2004, pp. 6–11.

[5] S. Borkar, T. Karnik, S. Narendra, J. Tschanz, A. Keshavarzi, and V. De, "Parameter variations and impact on circuits and microarchitecture," in *Proc. ACM/IEEE Des. Autom. Conf.*, 2003, pp. 338–342.

[6] E. T. Brist and A. J. Torres, "Model-assisted placement of subresolution assist features: Experimental results," *Proc. SPIE*, vol. 5042, pp. 99–106, 2003.

[7] Y. Cao, T. Sato, M. Orshansky, D. Sylvester, and C. Hu, "New paradigm of predictive MOSFET and interconnect modeling for early circuit simulation," in *Proc. Custom Integr. Circuits Conf.*, 2000, pp. 201–204.

[8] H. Chang and S. S. Sapatnekar, "Full-chip analysis of leakage power under process variations, including spatial correlations," in *Proc. ACM/IEEE Des. Autom. Conf.*, 2005, pp. 523–528.

[9] R. S. Chau, "Intel's breakthrough in high-K gate dielectric drives Moore's law well into the future," *Technology@Intel Magazine*, 2004.

[10] L. Chen, Y. Cao, H.-Y. Liu, W. Shao, M. Feng, and J. Ye, "Predictive focus exposure modeling (FEM) for full-chip lithography," *Proc. SPIE*, vol. 6154, p. 61541T, 2006.

[11] R. Choudhury, "Handbook of microlithography and micromachining, volume 1," in *SPIE Press Monograph*. Bellingham, WA: SPIE Press, 2002.

[12] J. Fishburn and A. Dunlop, "TILOS: A posynomial programming approach to transistor sizing," in *Proc. IEEE Int. Conf. Comput.-Aided Des.*, 1985, pp. 326–328.

[13] D. G. Flagello, H. van der Laan, J. van Schoot, I. Bouchoms, and B. Geh, "Understanding systematic and random CD variations using predictive modelling techniques," in *Proc. SPIE—Conf. Optical Microlithography XII*, 1999, pp. 162–175.

[14] Y. Gu, S. Chang, G. Zhang, K. Kirmse, D. Rogers, and L. Olsen, "Local CD variation in 65 nm node with PSM processes STI topography characterization (I)," in *Proc. SPIE—Conf. Metrology, Inspection, Process Control Microlithography*, 2006, vol. 6152, p. 615204.

[15] P. Gupta and F.-L. Heng, "Toward a systematic-variation aware timing methodology," in *Proc. ACM/IEEE Des. Autom. Conf.*, 2004, pp. 321–326.

[16] P. Gupta, A. B. Kahng, S. Muddu, and S. Nakagawa, "Modeling edge placement error distribution in standard cell library," in *Proc. SPIE—Conf. Design Process Integration Microelectronic Manufacturing IV*, 2006, pp. 61560S1–61560S12.

[17] P. Gupta, A. B. Kahng, P. Sharma, and D. Sylvester, "Selective gate-length biasing for cost-effective runtime leakage control," in *Proc. ACM/IEEE Des. Autom. Conf.*, 2004, pp. 327–330.

[18] W. Jiang, V. Tiwari, E. Iglesia, and A. Sinha, "Topological analysis for leakage prediction of digital circuits," in *Proc. IEEE Asia South Pac. Des. Autom. Conf.*, 2002, pp. 39–44.

[19] A. B. Kahng, S. Muddu, and P. Sharma, "Defocus-aware leakage estimation and control," in *Proc. Int. Symp. Low Power Electron. Des.*, 2005, pp. 263–268.

[20] B. Lee, "Modeling of chemical mechanical polishing for shallow trench isolation," Ph.D. dissertation, MIT, Cambridge, MA, 2002.

[21] B. Lee, D. S. Boning, W. Baylies, N. Poduje, P. Hester, Y. Xia, J. Valley, C. Koliopoulos, D. Hetherington, H. Sun, and M. Lacy, "Wafer nanotopography effects on CMP: Experimental validation of modeling methods," in *Proc. MRS Spring Meeting*, Apr. 2001.

[22] X. Xie, T. Park, D. Boning, A. Smith, P. Allard, and N. Patel, "Characterizing STI CMP processes with an STI test mask having realistic geometric shapes," in *Proc. Chem.-Mech. Polishing Symp., MRS Spring Meeting*, 2004.

[23] R. M. Rao, J. L. Burns, A. Devgan, and R. B. Brown, "Efficient techniques for gate leakage estimation," in *Proc. Int. Symp. Low Power Electron. Des.*, 2003, pp. 100–103.

[24] R. M. Rao, A. Srivastava, D. Blaauw, and D. Sylvester, "Statistical estimation of leakage current considering inter- and intra-die process variation," in *Proc. Int. Symp. Low Power Electron. Des.*, 2003, pp. 84–89.
[25] S. Sirichotiyakul, T. Edwards, C. Oh, R. Panda, and D. Blaauw, "Duet: An accurate leakage estimation and optimization tool for dual-$V_{\rm th}$ circuits," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 10, no. 2, pp. 79–90, Apr. 2002.
[26] A. K. Wong, *Resolution Enhancement Techniques in Optical Lithography*, vol. TT47. Bellingham, WA: SPIE Press, 2001.
[27] P. Yu, S. X. Shi, and D. Z. Pan, "Process variation aware OPC with variational lithography modeling," in *Proc. ACM/IEEE Des. Autom. Conf.*, 2006, pp. 785–790.

Andrew B. Kahng received the A.B. degree in applied mathematics from Harvard College, Cambridge, MA, and the M.S. and Ph.D. degrees in computer science from the University of California at San Diego, La Jolla.

From 1989 to 2000, he was a member of the computer science faculty with the University of California at Los Angeles, Los Angeles. He is currently a Professor with the Department of Computer Science and Engineering and the Department of Electrical and Computer Engineering, University of California at San Diego. He has published over 200 papers in the very large scale integration (VLSI) computer-aided-design literature. His research interest focuses mainly on physical design and performance analysis of VLSI as well as the VLSI design-manufacturing interface. Other research interests include combinatorial and graph algorithms, and large-scale heuristic global optimization.

Dr. Kahng has been defining the physical design roadmap for the International Technology Roadmap for Semiconductors (ITRS) since 1997, and he has been chairing the US and international working groups for design technology for the ITRS since 2001. He has been active in the MARCO Gigascale Silicon Research Center since its inception. He was also the founding General Chair of the ACM/IEEE International Symposium on Physical Design and cofounded the ACM Workshop on System-Level Interconnect Planning. He was the recipient of three Best Paper Awards and an NSF Young Investigator Award.

Swamy Muddu received the B.Eng. degree in electronics and communication engineering from the College of Engineering, Osmania University, Hyderabad, India, in 2001, and the M.S. degree from the Department of Electrical and Computer Engineering, University of California at San Diego, La Jolla, in 2003, where he is currently working toward the Ph.D. degree.

In 2005, he was an Engineer with the Blaze DFM, Sunnyvale, CA. His research interests are in the areas of very large scale integration physical design for manufacturability, modeling of manufacturing variation, and low-power design.

Puneet Sharma received the B.Tech. degree in computer science and engineering from the Indian Institute of Technology, New Delhi, India, in 2002. He has been working toward the Ph.D. degree in the Department of Electrical and Computer Engineering, University of California at San Diego, La Jolla, since September 2002.

His research interests include delay fault testing, power reduction, and design for manufacturability.