Body-bias-driven design strategy for area and performance efficient CMOS circuits

Citation for published version (APA):

Meijer, M., & Pineda de Gyvez, J. (2012). Body-bias-driven design strategy for area and performance efficient CMOS circuits. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 20(1), 42-51.
<https://doi.org/10.1109/TVLSI.2010.2091974>

DOI:

[10.1109/TVLSI.2010.2091974](https://doi.org/10.1109/TVLSI.2010.2091974)

Document status and date:

Published: 01/01/2012

Document Version:

Accepted manuscript including changes made at the peer-review stage

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

Body-Bias-Driven Design Strategy for Area- and Performance-Efficient CMOS Circuits

Maurice Meijer and José Pineda de Gyvez, *Fellow, IEEE*

Abstract—Worst-case design uses extreme process corner conditions which rarely occur. This limits maximum speed specifications and costs additional power due to area over-dimensioning during synthesis. We present a new design synthesis strategy for digital CMOS circuits that makes use of forward body biasing. Our approach renders consistently a better performance-per-area ratio by constraining circuit over-dimensioning without sacrificing circuit performance. An in-depth analysis of the body-bias-driven design theory is provided. It is complemented by an algorithm that enables fast reconstruction of the area-clock period tradeoff curve of the design. We validated these new concepts through industrial processor designs in 90-nm low-power CMOS. For standard- V_{th} implementations, we observed performance-per-area improvements up to 40%, area and leakage reductions up to 30%, and dynamic power savings of up to 10% without performance penalties as a benefit from our proposed body-bias-driven design strategy. The benefits are larger for high- V_{th} implementations. In this case, we observed performance-per-area improvements up to 90%, area and leakage reductions up to 40%, and dynamic power savings of up to 25% without performance penalties.

Index Terms—Circuit optimization, circuit tuning, CMOS digital integrated circuits, logic design.

I. INTRODUCTION

CONVENTIONAL and well-established digital design practices are based on a worst-case design (WCD) style to guarantee chip operation for meeting timing specifications among the process corners [1]. The circuit is designed in the slow-process corner to meet frequency specifications, while the maximum leakage target is verified in the fast-process corner. However, such extreme process corners rarely occur in most of the fabricated chips. Moreover, WCD makes high performance specifications harder to meet due to over-dimensioning of the design. Over-dimensioning leads to a larger silicon footprint, higher power consumption, and larger leakage. Fig. 1 shows the area-delay tradeoff involved during logic synthesis. Observe that circuit area depends on the process margin for high-performance circuits. If a lower process margin can be tolerated without a parametric yield penalty, circuit performance can be further increased without spending excessive area. Statistical circuit design has long been seen as a viable way to avoid

Fig. 1. Area-delay tradeoff at logic synthesis.

the use of worst-case parameters [2], [3]. However, these approaches have not totally found their way into industrial practices. This is due to, among other reasons, the moving average of process parameters, the flexibility of fabricating the same chip in multiple foundries, and the lack of appropriate EDA tools for statistical logic synthesis. In this paper, we show that a body-bias-driven (BBD) logic synthesis overcomes these drawbacks.

Alternatively, post-silicon tuning has been proposed for improving product-binning yields and for trading off power performance [4]–[7], but does not eliminate the problem of area over-dimensioning. Body biasing is typically used for leakage reduction or performance tuning [4]–[8]. Forward body biasing (FBB) is preferred over supply voltage scaling (VS) to achieve increased performance [4], [7], [8]. This is because the power penalty of FBB is lower for dynamic-power dominant designs. Reverse body biasing (RBB) can effectively achieve leakage reductions [6]. Other post silicon tuning works have been reported. For instance, a joint design-time and post-silicon tuning optimization strategy for minimizing leakage under delay constraints was proposed in [9]. This approach relies on detailed process variability inputs and is capable of reducing process-dependent delay spread. However, it consider neither a timing speed-up nor a circuit area reduction as outcome. Others propose body bias clustering at design-time for minimizing leakage under delay constraints [10], [11], or enhancing circuit performance [12]. These approaches do not consider a (joint) design-time optimization for improving performance or reducing area of the circuit.

Threshold voltage (V_{th}) assignment and gate sizing during the design synthesis phase is a known problem [13], [14]. V_{th} assignment has been used for reducing leakage or for reducing dynamic power consumption. Leakage power of digital IP blocks is mostly a concern when the circuit is in standby mode. High-performance circuits typically use low- V_{th} (LVT) devices to speed-up critical delay paths at a higher intrinsic device leakage penalty [13]. This higher leakage is unacceptable for portable

Manuscript received April 23, 2010; revised August 26, 2010; accepted October 19, 2010.

M. Meijer is with the Central R&D Division, NXP Semiconductors, 5656A Eindhoven, The Netherlands (e-mail: maurice.meijer@nxp.com).

J. Pineda de Gyvez is with the Central R&D Division, NXP Semiconductors, 5656A Eindhoven, The Netherlands, and also with the Electrical Engineering Department, Technical University of Eindhoven, 5600MB Eindhoven, The Netherlands (e-mail: jose.pineda.de.gyvez@nxp.com).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TVLSI.2010.2091974

Fig. 2. FBB utilization under BBD design.

applications since it increases standby power, and thereby reduces battery time. The use of body biasing offers several advantages: 1) it offers a continuum of V_{th} values; 2) it is V_{th} technology flavor independent; 3) it can be used on top of multi- V_{th} assignments; and 4) it can be used applied dynamically or adaptively. FBB can achieve LVT performance during operation, while it can be turned off to achieve low leakage in standby [7]. LVT circuits with RBB cannot achieve such low leakage [15]. Like multi- V_{th} and gate sizing assignments, BBD designs render smaller footprint area than WCD. Unlike multi- V_{th} and gate sizing assignments, the V_{th} choice is not technology-constrained since it is possible to characterize a standard cell library with FBB targeting a given V_{th} value within a certain range of V_{th} 's. In [16], we presented a body-bias-driven gate-level optimization method that leverages FBB to improve the performance-per-area (PPA) ratio of digital CMOS circuits. In this paper, we provide an in-depth analysis of the body-bias-driven design theory. This theory allows us to predict the design's optimum PPA with a minimum number of synthesis trials. We validated these new concepts through industrial processor designs in 90-nm LP-CMOS. In this paper, we discuss as well how our approach is fully integrated in a state-of-the-art commercial design flow.

The remainder of this paper is organized as follows. In Section II, we introduce BBD design. Section III presents the theoretical background and modeling. In Section IV, we explore the area, performance, and power trends for BBD design. Section V presents the BBD logic synthesis approach. In Section VI, we validate the proposed models. Section VII shows our benchmarked results. Finally, Section VIII presents our conclusions.

II. BBD DIGITAL DESIGN

Here, we will introduce BBD design and present body-bias silicon-tuning capabilities for a 90-nm low-power CMOS process technology.

A. Concept

Under WCD, digital CMOS circuits are implemented to meet timing specifications for slow process conditions. Observe, however, that FBB enhances circuit speed. Bearing this

TABLE I
EXPERIMENTAL RESULTS FOR 90-nm LP-CMOS RING-OSCILLATORS AT $V_{DD} = 1.1$ V AND $T = 85$ °C FOR A SLOW DIE SAMPLE

FBB [V]	Performance increase [%]		Rel. leakage increase factor	
	SVT	HVT	SVT	HVT
0	0	0	1.0	1.0
0.1	4	7	1.6	1.8
0.2	9	14	2.7	3.5
0.3	14	22	4.9	7.7
0.4	19	31	9.9	20.4
0.5	24	40	24.4	80.0

in mind, one does not need to pursue WCD. Instead, it is possible to design the circuit in between the worst and nominal process corners provided that the IC has FBB capabilities to correct performance deviations due to fabrication outcome. This creates opportunities for more cost-effective solutions without sacrificing performance specs and parametric yield. The amount of FBB required can be calibrated at test time or during boot of the chip.

Fig. 2 illustrates the parameters that are under control with BBD design. The right-hand side of Fig. 2 plots the dependency between clock period and FBB. A higher FBB value enables faster circuit operation. The amount of speed-up depends on the process technology, the used transistor threshold voltage option, and the design's power supply voltage. FBB needs only to be applied to those die samples with a lower speed than the nominal process outcome. The left-hand side of Fig. 2 plots the relationship between circuit area and clock period. For increasing FBB values, the curve shifts linearly proportional to a reducing clock period. Notice that a performance increase by FBB can be traded off against a performance decrease due to a smaller circuit area. In this way, we are able to maximize the PPA ratio of the circuit at design-time, while meeting a target performance.

B. Body-Bias Tuning Capabilities in CMOS 90-nm Process

The effectiveness of BBD design depends on the performance tuning range available with FBB. We briefly summarize our experimental results that were obtained for a set of ring-oscillator test structures in a 90-nm LP CMOS process. These test structures are similar to the ones presented in [4]. Both standard- V_{th} (SVT) and high- V_{th} (HVT) versions are available. Measurements have been performed for 61 die samples of the same 300-mm wafer at $V_{DD} = 1.1$ V and $T = 85$ °C. In our experiments, we applied FBB of, at most, 0.5 V to avoid turning on the device's junction diodes. FBB is applied simultaneously to pMOS and nMOS transistors through P- and N-well biasing, respectively.

Table I presents the measurement results for a slow die sample. A 24% and 40% performance increase is observed for the SVT and HVT ring-oscillator test structures, respectively, when 0.5-V FBB is applied to both N- and P-wells simultaneously. Contrarily, leakage increases by up to about 25 $\times$ and 80 $\times$, respectively. The leakage increase is more severe for HVT. This is because the forward-biased junction leakage at 0.5-V FBB dominates over the subthreshold leakage.

LVT circuits show a lower performance and leakage increase with FBB as compared with SVT [15]. The intrinsic leakage of LVT is about 10 $\times$ higher than in case of SVT, which has a large impact on power consumption in standby or low-activity

use-cases. Therefore, we will focus on the use of SVT and HVT in the remainder of this work.

III. THEORETICAL BACKGROUND BBD DESIGN

Here, we present the theoretical background of BBD design for achieving an optimum PPA ratio.

A. Area and Clock Period Modeling

The delay of a digital logic gate can be modeled as [13]

$$d_{\text{gate}} = \left(d_0 + \frac{d_1}{x} \right) (1 + kV_{\text{BB}}). \quad (1)$$

The first term represents the intrinsic (d_0) and load-dependent (d_1/x) gate delay [17]. Parameter x is the gate sizing factor ($x \geq 1$). The gate delay is both V_{DD} - and V_{th} -dependent. The second term models the impact of FBB on gate delay by a linear function. V_{BB} represents the FBB value: $V_{\text{BB}} = V_{\text{pwell}} = V_{\text{DD}} - V_{\text{nwell}}$. Parameter k is the polynomial coefficient, which can be different for each gate. The maximum error of using such a linear function to model the delay dependency on FBB is lower than 2% for our 90-nm LP-CMOS test-structures.

Based on (1), we model the delay and area of a CMOS digital logic circuit as

$$D_j = \sum_{i \in j} \left(d_{0i} + \frac{d_{1i}}{x_i} \right) (1 + k_i V_{\text{BB}_i}) \leq T_{\text{ck}} \quad \forall j \in \Psi \quad (2)$$

$$A_{\text{total}} = \sum_{i=1}^m x_i A_i \quad (3)$$

where i is an index that runs over all gates in the circuit, j is an index that runs over all paths in the circuit, D_j is the delay of path j , Ψ is the collection of all paths in the circuit, and A_i is the minimum area of gate i . Expression (2) constrains the delay of each circuit path to be less than the targeted clock period T_{ck} . The total circuit area is the summed gate areas.

Fig. 3 shows a typical area-clock period tradeoff curve for a given generic digital logic circuit. The curve is constructed from a multitude of synthesis runs such that the same design meets distinct clock period constraints. In Fig. 3, the area and clock period have been normalized to the best performing design ($A_{\text{max}}, T_{\text{min}}$). This design is obtained by constraining gate sizing of digital gates to their maximum size in the digital library. The faster designs ($T_{\text{ck}} < T_{\text{min}}$) are obtained for unconstrained gate sizing. Observe that high-performance circuits consume more area than slow circuits. This is due to gate upsizing and logic reordering to speed up critical circuit paths. The trend shown in Fig. 3 can be modeled by a rational function, given as follows, with χ , δ , and η as independent fitting parameters:

$$A_{\text{total}} = \frac{\chi}{\delta + T_{\text{ck}}} + \eta \quad (4)$$

The general form of (4) describes a rectangular hyperbola. Parameters δ and η model the shift of origin. The vertical asymptote is located at $T_{\text{ck}} = -\delta$, which represents the minimum clock period of the design that is theoretically possible. The horizontal asymptote is located at $A_{\text{total}} = \eta$, which represents the minimum area of the design in case of no gate upsizing or logic restructuring. Parameter χ models the hyperbola scale factor, which accounts for the impact of gate upsizing and logic re-

Fig. 3. Area and clock period tradeoff for a generic digital logic circuit.

structuring. A characteristic point is the clock period value T_A at which the slope of the hyperbola equals -1 . This clock period T_A will be used to reconstruct the design's area and clock period tradeoff curve, as will be discussed in Section V-B. T_A and the corresponding circuit area A_A can be determined as follows:

$$T_A = -\delta + \sqrt{\chi} \quad \forall T_{\text{ck}} \geq T_{\text{min}} \wedge \chi \geq 0 \quad (5)$$

$$A_A = \eta + \sqrt{\chi} \quad \forall T_{\text{ck}} \geq T_{\text{min}} \wedge \chi \geq 0. \quad (6)$$

Next, we will discuss the relationship between the fitting parameters of (4) for WCD and BBD design styles. Parameter η is identical for both design styles because of the very relaxed or unconstrained timing. Now, notice that, if both WCD and BBD circuits were optimized in the same way over the entire clock period range, then χ would be the same. However, this is not true in general since BBD libraries are faster than conventional libraries, e.g., the gate drive of a forward bias cell is larger than the one of the same cell without FBB.

Let us now take a look at speed and area tradeoffs between WCD and BBD design styles. Suppose first that a given circuit area ($A_{\text{total_bbd}} = \lambda \cdot A_{\text{total_wcd}}$) is desired, then the clock period of the BBD circuit can be obtained from the WCD clock period

$$T_{\text{ck_bbd}} = -\delta_{\text{bbd}} + \frac{\chi_{\text{bbd}}(\delta_{\text{wcd}} + T_{\text{ck_wcd}})}{\lambda\chi_{\text{wcd}} + (\eta_{\text{wcd}} - \eta_{\text{bbd}})(\delta_{\text{wcd}} + T_{\text{ck_wcd}})} \quad (7)$$

where λ represents the fraction $A_{\text{total_bbd}}/A_{\text{total_wcd}}$. Parameter λ equals 1 for a constant circuit area between both design styles. Alternatively, suppose now that a given clock period ($T_{\text{ck_bbd}} = \sigma \cdot T_{\text{ck_wcd}}$) is pursued, then the circuit area of the BBD circuit can be obtained from the WCD circuit area as follows:

$$A_{\text{total_bbd}} = \eta_{\text{bbd}} + \frac{\chi_{\text{bbd}}(A_{\text{total_wcd}} - \eta_{\text{wcd}})}{(\delta_{\text{bbd}} - \sigma\delta_{\text{wcd}})(A_{\text{total_wcd}} - \eta_{\text{wcd}}) + \sigma\chi_{\text{wcd}}} \quad (8)$$

where σ represents the fraction $T_{\text{ck_bbd}}/T_{\text{ck_wcd}}$ and equals 1 for a constant clock period. Notice from (7) that the speed advantage of the BBD circuit depends only on the difference between δ 's provided that $\chi_{\text{bbd}} = \chi_{\text{wcd}}$. The smaller area of BBD

in (8) is also due to the difference of δ 's. These results are expected since digital gates with FBB have a greater output drive than without FBB. Consequently, smaller area gates are employed in BBD designs. Equations (7) and (8) enable designers to estimate the effectiveness of BBD over WCD in trading off circuit speed against area. Design and process technology alternatives can be compared once the parameter values for χ , δ , and η are known. These parameters are design-dependent because of different amount and type of digital cells used as well as the logic implementation. Moreover, they are also process-technology-dependent because circuit area, performance, and body-bias sensitivity depend on technology scaling. For example, a given digital logic circuit will be smaller (lower η , different χ) and faster (lower $|\delta|$, different χ) when implemented in a next-generation CMOS technology.

B. PPA Figure-of-Merit (FOM)

Circuit performance and area are key performance metrics for digital circuit designers. We introduce a new metric (PPA) to qualify how effectively the design achieves high performance while accounting for area scaling. The PPA metric depends on the technology node, the technology's threshold voltage option, and the standard cells available for circuit synthesis. Let $f_{ck} = 1/T_{ck} = 1/\max(D_j)$. We then obtain

$$\text{PPA} = \frac{f_{ck}}{A_{\text{total}}} = \frac{1}{T_{ck} A_{\text{total}}}. \quad (9)$$

A higher PPA value indicates that the circuit design utilizes silicon area more effectively to achieve a high performance. There exists a point in (4) with a maximum PPA. This point indicates the optimum performance without circuit over-dimensioning. By combining (4) and (9), we obtain

$$\text{PPA} = \frac{\delta + T_{ck}}{(\chi + \delta\eta + \eta T_{ck})T_{ck}}. \quad (10)$$

The clock period value at which the maximum PPA occurs (T_{best}), can be determined by making the derivative of PPA with respect to T_{ck} equal to zero. By solving the equation for T_{ck} , we obtain a closed-form expression for T_{best} , namely

$$T_{\text{best}} = -\delta + \frac{\sqrt{-\delta\chi\eta}}{\eta} \quad \forall T_{ck} \geq T_{\min} \wedge \delta\chi\eta \leq 0, \quad (11)$$

$T_{ck} > T_{\text{best}}$, yields circuits without area over-dimensioning, and the contrary holds true for $T_{ck} < T_{\text{best}}$. Therefore, T_{best} identifies the minimum possible clock period without circuit over-dimensioning. The maximum PPA at $T_{ck} = T_{\text{best}}$ is obtained after substituting (11) into (10) as follows:

$$\max(\text{PPA}) = \frac{1}{\chi - \delta\eta + 2\sqrt{-\delta\chi\eta}} \quad \forall T_{ck} \geq T_{\min} \wedge \delta\chi\eta \leq 0. \quad (12)$$

Under WCD, T_{best} may be too large to meet the target frequency specification of high-performance designs. In this case, over-dimensioning cannot be avoided, thereby worsening PPA.

In the forthcoming analysis, we make use of a normalized representation of PPA. The normalization is against the highest performance under WCD ($f_{ck} = f_{\max} = 1/T_{\min}$, $A_{\text{total}} = A_{\max}$):

$$\text{PPA}_{\text{norm}} = \frac{T_{\min} A_{\max}}{T_{ck} A_{\text{total}}}. \quad (13)$$

C. Power Modeling

Power consumption of a digital gate can be modeled as

$$P_{\text{gate}} = a(xC_{\text{intr}} + C_{\text{load}})V_{\text{DD}}^2 f_{ck} + xI_{\text{leak}}V_{\text{DD}} \quad (14)$$

where a is the switching activity of the gate, C_{intr} and C_{load} are the intrinsic and load capacitance of a gate, respectively, and f_{ck} is the operating frequency. I_{leak} is the leakage current of a gate, which depends both V_{DD} and V_{th} . From experimental results, we model the normalized leakage current dependence on body biasing by a fourth-order polynomial expression

$$\text{leakage}_{\text{norm}} = 1 + \sum_{n=1}^4 l_n V_{\text{BB}}^n. \quad (15)$$

The leakage at various FBB conditions has been normalized to the case of nominal body bias. As before, V_{BB} represents the FBB value: $V_{\text{BB}} = V_{\text{pwell}} = V_{\text{DD}} - V_{\text{nwell}}$. Parameters l are the polynomial coefficients, which are different for each gate. The maximum error of expression (15) is lower than 2% for our 90-nm LP-CMOS test-structures.

The intrinsic (or junction) capacitance of a gate is dependent on the applied body bias [8]. In our experiments, we have extracted the junction capacitance values from the dynamic power consumption measurement results. Hence, we model the normalized junction capacitance by a second-order polynomial expression. As before, the normalization has been done against the nominal body bias case:

$$C_{\text{intr},\text{norm}} = 1 + \sum_{n=1}^2 m_n V_{\text{BB}}^n. \quad (16)$$

Parameters m are the polynomial coefficients, which are different for each gate. The maximum error of expression (16) is lower than 1.5% for our 90-nm LP-CMOS test-structures when used to model the body-bias impact on dynamic power.

By combining (14), (15), and (16), we model the total power consumption of a generic CMOS digital logic circuit as

$$P_{\text{total}} = \sum_{i=1}^m a_i \left(x_i C_{\text{intr},i} \left(1 + \sum_{n=1}^2 m_{n,i} V_{\text{BB}}^n \right) + C_{\text{load},i} \right) \times V_{\text{DD}}^2 f_{ck} + \sum_{i=1}^m x_i V_{\text{DD}} I_{\text{leak},i} \left(1 + \sum_{n=1}^4 l_{n,i} V_{\text{BB}}^n \right) \quad (17)$$

where i is an index that runs over all gates in the circuit. Observe that we are assuming that WCD and BBD circuits use the same power supply voltage and operate at the same temperature. WCD and BBD circuits have different power consumption depending on differences in circuit dimensions, circuit activity, and operating frequency. Moreover, BBD circuits utilize FBB that increases power.

IV. OPTIMUM PPA DESIGN SPACE

Here, we explore area, performance, and power trends for WCD and BBD design styles by using the previously presented models. For this purpose, we take a generic digital logic circuit with calibrated technology parameters for 90-nm LP-CMOS. The analysis was done at $V_{\text{DD}} = 1.1$ V and $T = 85^\circ\text{C}$. For BBD design, we utilized a maximum FBB of 0.5 V to explore the limits of PPA driven design. All results relate to the slow-

Fig. 4. Area, clock period, and PPA tradeoff for a generic digital logic circuit under BBD and WCD. Solid line: WCD; dashed line: BBD; overlay: PPA.

process corner. Finally, we discuss technology-scaling implications by analyzing the same circuit in 65- and 45-nm LP-CMOS. The same process and operating conditions have been used as before with the exception of $V_{DD} = 1$ V for the 45-nm case.

A. PPA Trends

Fig. 4 shows the design exploration space for circuit area, clock period, and PPA. The area–clock period trend curves are plotted for WCD (solid line) and BBD (dashed line) design. The iso-PPA curves are plotted as overlay. The intersection with the area–clock period curves represents the normalized PPA ratio of the design as defined by (13). Logic synthesis usually aims at achieving a given target speed. As way of example, all PPA values of Fig. 4 have been normalized to the maximum frequency of operation under WCD ($T_{ck} = T_{min}$). This reference point is highlighted by the triangle symbol in Fig. 4. The triangle is located at a clock period of T_{min} , while the circles relate to T_{best} which are the corresponding best PPA points. Observe from Fig. 4 that, for a given circuit area, BBD design achieves higher performance than WCD counterparts. Alternatively, BBD design enables lower area designs for a given clock period. Any FBB of less than 0.5 V results in area–clock period curves located in between the two curves plotted in Fig. 4. Therefore, it makes most sense to use BBD design with the maximum possible FBB to obtain the best PPA ratio.

Fig. 5 highlights the PPA and clock period trends under WCD and BBD design. Notice that BBD design achieves a better PPA ratio than WCD under all circumstances. From large clock periods towards smaller ones, the PPA increases to a maximum value irrespective of the chosen design style. The increasing PPA is because the decrease in clock period is greater than the increase in circuit area. This trend is reversed after the maximum PPA has been reached due to area over-dimensioning. Observe from Fig. 5 that the maximum PPA can significantly be higher than the PPA of the maximum frequency under WCD. At large clock periods, the PPA of the BBD and WCD circuits is similar (not shown in Fig. 5).

Fig. 5. PPA versus clock period for a generic digital logic circuit under BBD and WCD. Solid line: WCD; dashed line: BBD.

Fig. 6. Area, clock period, and power tradeoff for a generic digital logic circuit under BBD and WCD. Solid line: WCD; dashed line: BBD; overlay: total power consumption (solid line) and dynamic power consumption (dotted line).

B. Power Consumption Trends

Fig. 6 shows the design exploration space for circuit area, clock period and power consumption. The trend lines for total power and dynamic power have been indicated under body bias conditions. The total power includes both dynamic and leakage power consumption. The iso-power curves are plotted as overlay, and the solid and dotted lines correspond to the total power and dynamic power, respectively. The intersection of the total power curves with the area–clock period curves represents the power consumed by the design. Observe in Fig. 6 that the power increases for decreasing clock periods due to a larger circuit area, and higher frequency of operation. The dynamic power increases linearly proportional to operating frequency. Recall that the same power supply voltage of 1.1 V is used for both WCD and BBD design. Only when FBB is applied, the leakage power becomes noticeable in the total power consumption. In this case, a difference occurs between the iso-total-power and iso-dynamic-power curves due to FBB under BBD design. This difference becomes larger for larger FBB values. The snapback point of the total power trends

Fig. 7. Area, clock period, and power tradeoff for a generic digital logic circuit under BBD and WCD and different technology nodes. The values have been normalized to the maximum PPA WCD in CMOS 90-nm process. Solid line: WCD; dashed line: BBD; overlay: total power consumption; symbols: maximum PPA designs.

defines the maximum FBB value to be applied from a power point of view. In our case, this point occurs at an FBB value of 0.5 V. Notice that BBD enables lower power operation at a constant clock period. This is because of the lower circuit area for BBD design. For a given power target, BBD design offers better performance and area figures. However, BBD design consumes more power for the same circuit area as WCD. This is not only because of the higher operating frequency, but also due to the higher junction capacitance and leakage power associated to the application of FBB.

C. Impact of Technology Scaling

Fig. 7 shows the design exploration space for circuit area, clock period and total power for the same generic digital logic circuit in different process technology nodes. Three groups of iso-power curves are plotted as overlay, each representing a given technology node. The symbols represent the maximum PPA designs for which the results are summarized in Table II. All values have been normalized to the maximum PPA design under WCD for 90-nm LP-CMOS. Observe in Fig. 7 the same area-clock period trends for each technology node. BBD design consistently outperforms WCD. The maximum PPA design is faster and smaller in a next-generation technology. Consequently, the PPA increases with technology scaling, as illustrated in Table II. BBD design achieves a similar PPA increase in each technology, because the performance increase with FBB is nearly constant [4], [15]. Also observe in Table II the opposing total power trends under WCD and BBD design. For WCD, the maximum PPA design operates at lower power in a scaled technology, despite the higher clock speed and the increasing leakage. This is no longer the case for BBD design, mainly due to the amplified leakage with FBB which is more pronounced in a scaled technology.

V. BBD DESIGN SYNTHESIS

From the previous sections, we saw that the PPA point indicates the optimum area-delay tradeoff for the circuit under consideration. Yet, it is necessary to construct this area-delay

TABLE II
TECHNOLOGY SCALING RESULTS OF MAXIMUM PPA DESIGNS AT $T = 85^\circ\text{C}$

		Relative PPA	Relative Clock Period	Relative Total Power
WCD	90nm LP-CMOS	1	1	1
	65nm LP-CMOS	1.45	0.85	0.87
	45nm LP-CMOS	2.49	0.61	0.74
BBD	90nm LP-CMOS	1.27	0.79	1.75
	65nm LP-CMOS	1.82	0.68	1.81
	45nm LP-CMOS	3.17	0.46	2.12

curve to find this point with minimum overhead effort. Here, we discuss the implementation of BBD design using a commercial logic synthesis tool and present an algorithm that enables fast reconstruction of the area-clock period tradeoff curve of the design.

A. BBD Synthesis With Commercial Tools

Commercial synthesis tools can target area optimization subject to delay constraints. To validate our approach, we have implemented BBD synthesis in Cadence RTL Compiler.¹ Digital cell libraries have been recharacterized to account for FBB in 90-nm LP-CMOS using Altos' Liberate library characterizer.² The library characterization uses the effective current source model (ECSM) for timing, noise, and power modeling. To enable BBD synthesis, FBB-characterized timing views have been created, utilizing 0.5-V FBB for pMOS and nMOS transistors. Both WCD and BBD digital cell libraries have been characterized for slow process conditions, $V_{DD} = 1.1\text{ V}$, and $T = 125^\circ\text{C}$ settings. Such digital cell libraries also enable static timing verification of BBD circuits.

B. Minimizing the Number of Synthesis Runs

Finding the maximum PPA design starts with an iterative process for collecting sufficient data points to reconstruct the area-clock period trade-off curve. Collecting data points evenly spaced across the clock period range requires many synthesis runs, which are time-consuming for large designs. Fortunately, the trade-off curve is a continuous function over a closed clock period interval. A reconstruction is possible with three specific data points only, namely a first point at a small clock period and large circuit area, a second point at a large clock period and a small circuit area, and a third point when the slope of the curve is -1 [at $T_{ck} = T_A$, see (5)].

Our approach to curve reconstruction is based on a greedy search algorithm. It makes use of a kind of Newton-Raphson iteration, which is known for its fast convergence [18]. The proposed algorithm searches for the clock period value at which the slope of the area-clock period tradeoff curve equals -1 ($T_{ck} = T_A$). Instead of calculating the derivative of the area-clock period explicitly, we made use of (5) to determine $T_{ck} = T_A$. Let us now address our algorithm as described in Fig. 8. As a first step, the design is synthesized at the minimum clock period bound, T_{low} . The synthesis tool returns the actual clock period T_1 and circuit area A_1 . Note that $T_1 > T_{low}$ when $T_{low} < T_{min}$, for other cases $T_1 \approx T_{low}$. Next, the design is synthesized at the maximum clock period bound, T_{high} . T_{high} should be chosen large enough to ensure the clock period range at which area

¹[Online]. Available: http://www.cadence.com/products/ld/rtl_compiler

²[Online]. Available: <http://www.altos-da.com/pdfs/liberate-ds.pdf>

Algorithm 1: Find_Fitting_Parameters(T_{low} , T_{high} , ϵ)

1. Choose minimum clock period bound T_{low}
2. Synthesize design at $T_{ck}=T_{low}$: obtain (T_1, A_1)
3. Choose maximum clock period bound T_{high} ($T_{high} \geq 10T_1$)
4. Synthesize design at $T_{ck}=T_{high}$: obtain (T_2, A_2)
5. Calculate $T_{cur} = (T_1 + T_2)/2$
6. Set iteration index $i = 3$
7. Synthesize design at $T_{ck}=T_{cur}$: obtain (T_i, A_i)
8. Perform least-squares regression by using expression (4) and all available (T, A) pairs as data points: obtain χ , δ , and η
9. Calculate $T_{ck}=T_{next}$ using expression (5)
10. **if** $|T_{next} - T_{cur}| > \epsilon$ **then**
11. $T_{cur} = T_{next}$
12. Increase iteration index: $i=i+1$
13. Go to Line 7
14. **end if**
15. Return fitting parameters χ , δ , η

Fig. 8. Greedy algorithm to obtain fitting parameters of (4) for reconstructing the area and clock period tradeoff curve of the design.

TABLE III
MODEL FITTING PARAMETERS FOR SVT MICROPROCESSOR

	χ [$\mu\text{m}^2\text{-ns}$]	δ [ns]	η [μm]	T_A [ns]	T_{best} [ns]
WCD	$11.38 \cdot 10^4$	-4.65	$2.78 \cdot 10^5$	15.3	6.0
BBD	$8.64 \cdot 10^4$	-3.65	$2.75 \cdot 10^5$	13.0	4.7

over-dimensioning occurs is captured. We have used a T_{high} value of $10T_1$. The synthesis tool returns the actual clock period T_2 and circuit area A_2 . The third synthesis point is chosen based on the bi-section to ensure proper conditioning with three points for curve fitting based on (4). We apply then the least-squares method to determine the fitting parameters χ , δ , and η . The clock period $T_{next} = T_A$, is determined at which the slope of area-clock period tradeoff curve equals -1 from (5). When the difference between the current clock period T_{cur} and the new one T_{next} is larger than the tolerated error ϵ , a new synthesis run is executed at $T_{cur} = T_{next}$. Again, the least-squares method is used to recalculate the fitting parameters using the available synthesis results. This process is repeated until $T_{next} - T_{cur}$ is smaller than ϵ . If this condition is met, the clock period at which the maximum PPA occurs can be calculated by expression (11), as shown before. A final synthesis run is required at $T_{ck} = T_{best}$ to obtain the maximum PPA design.

C. Physical Design Aspects and Body Bias Generation

The backend views of the digital standard cells have been prepared for BBD place-and-route support. N-well and P-well taps have been removed from the digital cell layouts. Instead, dedicated tap cells have been added to the library. The circuit is place-and-routed in cell rows as shown in [19], while inserting tap cells in columns at a maximum pitch of $60 \mu\text{m}$. A two-layer routing grid for connecting the tap cells to the body bias supplies has been utilized.

The body-biased cells share the same N-well and P-well. Deep N-well isolation is added to (part of) the design for separating the P-well of the body biased nMOS devices from the P-substrate. Since the nMOS devices share the same P-well, the overhead of the deep N-well is minimized. Only $2 \mu\text{m}$ extra is needed at each side of the body biased circuit part.

Fig. 9. Area versus clock period for the microprocessor design in 90-nm LP-CMOS. Lines: WCD (solid) and BBD (dashed) model; symbols: synthesis results. The normalized PPA ratio is indicated for each design.

Dynamic FBB requires a voltage generator circuit to generate the N-well and P-well bias voltages. A 90-nm LP-CMOS solution has been presented in [20]. This FBB generator occupies 0.03 mm^2 for driving digital circuits of 1 mm^2 , translating into $\sim 3\%$ area overhead. The generator's size increases by 0.01 mm^2 for each additional mm^2 digital circuit size. For digital circuits smaller than 1 mm^2 , the size of the generator needs to be adapted to minimize area overhead.

VI. MODEL VALIDATION

WCD and BBD design have been analyzed and compared for a commercial microprocessor design in 90-nm LP-CMOS. The circuit contains 3764 flip-flops and about 31 K combinational gates. It makes use of SVT devices only. This section presents correlated results obtained from logic synthesis and the presented models. As before, the analysis has been performed for slow-process conditions, $V_{DD} = 1.1 \text{ V}$, and $T = 85^\circ\text{C}$. BBD design makes use of a maximum FBB of 0.5 V .

A. Design Synthesis Approach

Design synthesis targeted reconstruction of the area-clock period tradeoff curve. We made use of the greedy algorithm as presented in Section VI-B. For each design style, only four synthesis runs were required for reconstruction. The algorithm received the following inputs: $T_{low} = 3 \text{ ns}$, $T_{high} = 30 \text{ ns}$, and $\epsilon = 250 \text{ ps}$. Table III summarizes the fitting parameters, T_A and T_{best} for WCD and BBD design styles. One more synthesis run was required to obtain the optimum PPA design.

B. Area, Clock Period, and PPA Comparison

Fig. 9 shows the design exploration space for circuit area and clock period for the given microprocessor design. The synthesis results have been indicated by circles and triangles for WCD and BBD design, respectively. The filled symbols correspond to the four synthesis cases based on our algorithm. The open symbols are additional synthesis cases for trend verification purposes only. The solid and dotted lines show the calculated tradeoff curves for WCD and BBD design, respectively, by using (4). The fitting parameters of the model are given in Table III.

Fig. 10. Area versus clock period for the microprocessor design in 90 nm LP-CMOS at $T = 85^\circ\text{C}$. Lines: WCD (solid) and BBD (dashed) model; symbols: synthesis results. The total power is indicated for each synthesized design.

Observe from Fig. 9 the close match between the modeled and the synthesized area–clock period trends. The rms error between the calculated curves and the location of each synthesis result is within 1.5%. After completing the fourth synthesis run, we calculated T_A and T_{best} values which are given in Table III. The PPA value for each synthesis point has been indicated in Fig. 9 normalized to a T_{min} of 5.5 ns under WCD. Observe the existence of a maximum PPA design for both WCD (PPA = 1.01@6 ns) and BBD design (PPA = 1.27@4.7 ns). The calculated T_{best} values match within 5% of the values obtained through synthesis (WCD: 6 ns; BBD: 4.9 ns). As expected, BBD design not only gives a better performance but also better area utilization as indicated by the PPA value.

C. Power Consumption Comparison

Fig. 10 presents the same area and clock period curve as before, but now the symbols indicate the normalized power consumption of each synthesis run. For the given microprocessor design, BBD design provides lower power operation than WCD at the same clock period. Contrarily, BBD design consumes more power at the same circuit area.

VII. BENCHMARKED RESULTS

Here, we present BBD and WCD results for three industrial processor designs in 90-nm LP-CMOS. Logic synthesis, physical implementation and power analysis has been done using Cadence’s RTL Compiler, First Encounter, and Encounter Timing System, respectively. All results have been obtained for a slow process corner, $V_{\text{DD}} = 1.1\text{ V}$, and $T = 85^\circ\text{C}$. BBD design utilizes a maximum FBB of 0.5 V. All area results account for layout effects including the overhead for deep-N-well isolation.

Each processor design has been implemented in both SVT and HVT flavors. Table IV shows the gate count summary. Two synthesis cases have been investigated, namely: 1) a maximum PPA design and 2) a maximum frequency design under WCD. In the latter case, BBD design is utilized to operate at the same speed at a lower area cost to improve the PPA ratio.

TABLE IV
EXAMPLE GATE COUNT—THREE INDUSTRIAL PROCESSOR DESIGNS

	#flip-flops	#logic gates
Digital signal processor	227	4416
Microprocessor	3764	34390
Multimedia processor	41749	252759

A. Design Synthesis for Maximum PPA

Table V presents the processor design results targeting a maximum PPA design. Five circuit parameters are presented, namely clock period, circuit area, PPA, and dynamic and leakage power. The BBD design results are presented relative to the WCD results. The PPA ratio has been normalized to the maximum performance under WCD ($T_{\text{ck}} = T_{\text{min}}$).

Let us first consider SVT results. Observe that the PPA ratio differs for each design. This depends on circuit characteristics such as circuit size, path delay distribution, and logic depth. Under WCD, the PPA ratio ranges from 1.01 to 1.10. The maximum PPA point for small circuits (low η value) tends to be located at larger clock period values ($T_{\text{best}} > T_{\text{min}}$), as can be inferred from (11). This explains the high PPA value of 1.10 for the digital signal processor. For large circuits (high η value), the maximum PPA value is located closer to, or equal to the minimum clock period, $T_{\text{best}} \approx T_{\text{min}}$. The path delay distribution of the multimedia processor is the reason for the better PPA value w.r.t. the microprocessor design (1.03 i.s.o 1.01). The multimedia processor has many (nearly) critical delay paths which are largely responsible for the area over-dimensioning of the design when requiring high performance. BBD design enables significant improvements in maximum PPA as compared with WCD, mainly due to higher clock speeds. The maximum PPA of the BBD designs ranges between 1.25–1.38. Let us look now into HVT results. The same PPA trends are observed as in the SVT case, but the increase in maximum PPA is much larger (maximum PPA: 1.52–1.90). This is because FBB has a larger impact on circuit speed for HVT. Worth noticing is that the HVT BBD processors can operate at the same speed of the SVT WCD equivalents. However, their PPA values are slightly lower due to a higher circuit area. Irrespective of the V_{th} option used, BBD design provides always a higher maximum PPA ratio than WCD. All BBD circuits operate faster than their WCD counterparts, while circuit area is comparable.

Table V also shows the dynamic power and leakage power consumption for each processor design. Notice that dynamic power dominates leakage power, even at a high operating temperature of 85°C and when FBB is applied. The ratio between dynamic and leakage power is in the range of 100–300 for SVT WCD (800–2100 for HVT WCD) for the considered processor designs. Under BBD design, this ratio is reduced to 10–30 for SVT BBD design, and 5–20 for HVT BBD design. Observe that the dynamic power for BBD is generally higher than under WCD. There are two reasons for this, namely: 1) the higher clock speed and 2) the higher junction capacitance due to FBB. Next, the BBD leakage power is significantly higher than the WCD leakage when FBB is utilized. FBB turns on the transistor’s junction diodes, which leads to a high additional leakage current, especially at higher temperature operation. This will be

TABLE V

DESIGN SYNTHESIS FOR MAXIMUM PPA—INDUSTRIAL PROCESSOR DESIGNS IN 90-nm LP-CMOS RELATIVE VALUES ARE SHOWN W.R.T. WCD FOR THE GIVEN V_{TH} OPTION. CONDITIONS: SLOW-PROCESS CORNER, $V_{DD} = 1.1$ V AND $T = 85$ °C

Design Unit		Clock period		Area		PPA		Dynamic Power		Leakage Power	
		WCD [ns]	BBD [ns]	WCD [μm^2]	BBD rel.	WCD	BBD	WCD [mW]	BBD rel.	WCD [μW]	BBD rel.
Digital signal processor	SVT	6.5	5.0	32656	1.04	1.10	1.38	2.9	2.16	9.1	21.9
Microprocessor	SVT	6.0	4.9	342705	0.97	1.01	1.25	19.1	1.60	67	21.4
Multimedia processor	SVT	3.8	3.0	3047184	0.95	1.03	1.35	67.2	1.72	519	22.3
Digital signal processor	HVT	10.5	6.5	34386	1.01	1.19	1.90	1.4	2.28	0.8	216
Microprocessor	HVT	8.5	6.0	374440	0.93	1	1.52	13.8	1.69	6.7	230
Multimedia processor	HVT	5.8	3.9	3447298	0.89	1.02	1.73	42.0	1.83	50.7	274

TABLE VI

DESIGN SYNTHESIS FOR OPTIMUM AREA—INDUSTRIAL PROCESSOR DESIGNS IN 90-nm LP-CMOS RELATIVE VALUES ARE SHOWN W.R.T. WCD FOR THE GIVEN V_{TH} OPTION. CONDITIONS: SLOW-PROCESS CORNER, $V_{DD} = 1.1$ V AND $T = 85$ °C

Design Unit		Clock period WCD,BBD [ns]	Area		PPA		Dynamic Power		Leakage Power	
			WCD [μm^2]	BBD rel.	WCD	BBD	WCD [mW]	BBD rel.	WCD [μW]	BBD rel.
Digital signal processor	SVT	5.5	42457	0.74	1	1.35	3.7	0.98	14	12.5
Microprocessor	SVT	5.5	370932	0.85	1	1.18	22.5	0.93	78	17.1
Multimedia processor	SVT	3.6	3233781	0.89	1	1.17	71.3	0.95	569	19.8
Digital signal processor	HVT	8.8	48621	0.60	1	1.72	2.0	0.75	1.4	152
Microprocessor	HVT	8.5	374440	0.80	1	1.50	12.6	0.75	6.7	184
Multimedia processor	HVT	5.7	2857787	0.84	1	1.36	43.3	0.87	55	219

also reflected in the total power which is the sum of dynamic and leakage power components. However, recall that FBB is only applied to those chip samples with a lower frequency than the targeted one due to the process outcome. Such slow samples have already an intrinsic low leakage current. Slow-process-corner samples receive the maximum FBB, while the other slow samples receive a lower FBB. When no FBB is applied, the BBD leakage power is proportional to the circuit area scaling (not shown in Table V). In addition, recall that we apply dynamic FBB during chip operation. In this way we avoid the leakage penalty associated to FBB during standby operation.

B. Design Synthesis for Optimum Area

Table VI presents the processor design results targeting a maximum WCD performance. The BBD circuits were designed to match the WCD performance. In this case, BBD circuits can enable significant area savings, irrespective of the V_{th} option used. For the SVT BBD designs, we observed area reductions between 11% and 26% w.r.t. their WCD versions. The benefits for the HVT processors are larger due to the stronger FBB dependence (14%–40% area savings). The reduced circuit area comes mostly from the area scaling of the combinatorial logic. In general, BBD circuits have fewer logic gates than WCD ones, while the amount of flip-flops is the same. The largest area savings has been obtained for the digital signal processor, which has about 19× more logic gates than flip-flops. The ratio between logic gates and flip-flops is lower for the other circuits, as

can be derived from Table IV. The lowest ratio is found for the multimedia processor, namely about 6× more logic gates than flip-flops. This explains the area scaling trends observed. The PPA for the BBD processors is not optimal, because the BBD operating frequency is not fully utilized. However, it is significantly higher than for their WCD equivalents irrespective of the V_{th} option used.

BBD design renders consistently lower dynamic power than WCD does when operating at the same maximum WCD frequency. The power reduction comes from the reduced circuit area despite the increasing junction capacitances with FBB. The dynamic power reduces up to 7% for the SVT processors, while HVT processors achieve up to 25% dynamic power reductions. We noticed that BBD design primarily affects logic gates in the data path; the clock power is not much reduced. Thus, the dynamic power savings are larger for circuits with higher data activities.

As before, the BBD leakage power is much higher than the WCD leakage when FBB is utilized. The leakage power increases up to 20× for the SVT processors, and up to 219× for the HVT ones. Recall that this leakage increase is of no concern since FBB is disabled during standby operation. The leakage power for BBD design without FBB enabled decreases by the same factor as the circuit area (not shown in Table VI). For samples that do not need FBB to achieve performance, a leakage reduction up to 26% and up to 40% is possible in case of SVT and HVT, respectively.

VIII. CONCLUSION

We presented a design synthesis strategy for digital CMOS integrated circuits that makes use of FBB. Our approach renders consistently a better PPA ratio by constraining circuit over-dimensioning without sacrificing circuit performance. An in-depth analysis of the BBD design analysis was provided, which enables designers to predict the design's optimum performance per area with a minimum number of synthesis runs. We validated these new concepts through industrial processor designs in 90-nm LP-CMOS. For SVT implementations, we observed PPA improvements up to 40%, area and leakage reductions up to 30%, and dynamic power savings of up to 10% without performance penalties. The benefits are larger for HVT implementations. In this case, we observed PPA improvements up to 90%, area and leakage reductions up to 40%, and dynamic power savings of up to 25% without performance penalties as a benefit from our proposed BBD design strategy.

ACKNOWLEDGMENT

The authors would like to thank A. Kumar, Central Research and Development, NXP Semiconductors, Eindhoven, The Netherlands, for his support regarding timing library generation.

REFERENCES

- [1] J. Zhang, "Worst case design of digital integrated circuits," in *Proc. of ISCAS*, London, U.K., Jun. 1994, pp. 153–156.
- [2] S. Duvall, "A practical methodology for the statistical design of complex logic products for performance," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 3, no. 1, pp. 112–123, Mar. 1995.
- [3] A. Nardi *et al.*, "Impact of unrealistic worst case modeling on the performance of VLSI circuits in deep submicron CMOS technologies," *IEEE Trans. Semicond. Manuf.*, vol. 12, no. 4, pp. 396–403, Nov. 1999.
- [4] M. Meijer and J. Pineda de Gyvez, "Technological boundaries of voltage and frequency scaling for power performance tuning," in *Adaptive Techniques for Dynamic Processor Optimization*, A. Wang and S. Naffziger, Eds. Berlin, Germany: Springer, 2008, pp. 25–47.
- [5] J. Tschanz *et al.*, "Adaptive body bias for reducing impacts of die-to-die and within-die parameter variations on microprocessor frequency and leakage," in *Proc. ISSCC*, San Francisco, CA, USA, Feb. 2002, pp. 344–345.
- [6] T. Chen and S. Naffziger, "Comparison of adaptive body bias (ABB) and adaptive supply voltage (ASV) for improving delay and leakage under the presence of process variation," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 11, no. 5, pp. 888–899, Oct. 2003.
- [7] A. Keshavarzi *et al.*, "Forward body bias for microprocessors in 130 nm technology generation and beyond," in *VLSI Circuits Dig. Symp. Tech. Papers*, Honolulu, HI, Jun. 2002, pp. 125–128.
- [8] S. Huang *et al.*, "Scalability and biasing strategy for CMOS with active well bias," in *Proc. Symp. VLSI Technol.*, Kyoto, Japan, Jun. 2001, pp. 107–108.
- [9] M. Mani *et al.*, "Joint design-time and post-silicon minimization of parametric yield loss using adjustable robust optimization," in *Proc. ICCAD*, San Jose, CA, Nov. 2006, pp. 19–26.
- [10] S. Kulkarni *et al.*, "A statistical framework for post-silicon tuning through body bias clustering," in *Proc. ICCAD*, San Jose, CA, Nov. 2006, pp. 39–46.
- [11] R. Teodorescu *et al.*, "Mitigating parameter variation with dynamic fine-grain body biasing," in *Proc. MICRO-40*, Chicago, IL, Dec. 2007, pp. 27–39.
- [12] A. Sathanur *et al.*, "Physically clustered forward body biasing for variability compensation in nanometer CMOS design," in *Proc. DATE*, Nice, France, Apr. 2009, pp. 154–159.
- [13] M. Hirabayashi *et al.*, "Design methodology and optimization strategy for dual-VTH scheme using commercially available tools," in *Proc. ISLPED*, Huntington Beach, CA, Aug. 2001, pp. 283–286.
- [14] Y. Liu and J. Hu, "A new algorithm for simultaneous gate sizing and threshold voltage assignment," *IEEE Trans. Comput.-Aided Des. Integr. Circuits*, vol. 29, no. 2, pp. 223–234, Feb. 2010.
- [15] M. Meijer *et al.*, "Post-silicon tuning capabilities of 45 nm low-power CMOS digital circuits," in *VLSI Circuits Symp. Dig. Tech. Papers*, Kyoto, Japan, Jun. 2009, pp. 110–111.
- [16] M. Meijer and J. Pineda de Gyvez, "Body bias driven design synthesis for optimum performance per area," in *Proc. ISQED*, San Jose, CA, Mar. 2010, pp. 472–477.
- [17] I. Sutherland, B. Sproull, and D. Harris, *Logical Effort: Designing Fast CMOS Circuits*. San Francisco, CA: Morgan Kaufmann, 1999.
- [18] J. F. Bonnans, J. C. Gilbert, C. Lemarechal, and C. A. Sagastizabal, *Numerical Optimization, Theoretical and Numerical Aspects*, 2nd ed. Berlin, Germany: Springer-Verlag, 2006.
- [19] L. Clark *et al.*, "Reverse-body bias and supply collapse for low effective standby power," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 12, no. 9, pp. 947–956, Sep. 2004.
- [20] M. Meijer *et al.*, "A forward body bias generator for digital CMOS circuits with supply voltage scaling," in *Proc. ISCAS*, Paris, France, Jun. 2010, pp. 2482–2485.

Maurice Meijer received the B.Eng. degree in electrical engineering from Eindhoven Polytechnic, Eindhoven, The Netherlands, in 1999, and the M.Sc. degree in electrical engineering from Eindhoven University of Technology, Eindhoven, in 2004.

From 1999 to 2006, he was a Research Scientist with the Digital Design and Test Group, Philips Research Laboratories, Eindhoven, The Netherlands, where he has been involved in the design of low-power digital circuits and signal integrity of deep-submicrometer CMOS designs. He is currently

a Senior Scientist with the Central Research and Development Division, NXP Semiconductors, Eindhoven, The Netherlands. His research interests are in the areas of low-power and variation-tolerant integrated circuit design.

José Pineda de Gyvez (F'09) received the Ph.D. degree from the Eindhoven University of Technology, Eindhoven, The Netherlands, in 1991.

From 1991 until 1999, he was a Faculty member with the Department of Electrical Engineering, Texas A&M University, College Station. He is currently a Senior Principal with NXP Semiconductors, Eindhoven, The Netherlands, where he is leading the Systems Power Management activities of the Research Sector. Since 2006, he has also held the professorship "Deep Submicron Integration" with

the Department of Electrical Engineering, Eindhoven University of Technology, Eindhoven. He is a member of the editorial board of the *Journal of Low Power Electronics*. He has authored or coauthored more than 100 publications and three books in the fields of testing, nonlinear circuits, and low-power design and holds a number of granted patents.

Dr. Pineda de Gyvez has served as an associate editor for the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS—PARTS I AND II and for the IEEE TRANSACTIONS ON SEMICONDUCTOR MANUFACTURING.