Body bias driven design synthesis for optimum performance per area

Citation for published version (APA):

Meijer, M., & Pineda de Gyvez, J. (2010). Body bias driven design synthesis for optimum performance per area. In *Proceedings of the 2010 11th International Symposium on Quality Electronic Design (ISQED)*, 22-24 March 2010, San Jose, California (pp. 472-477). Institute of Electrical and Electronics Engineers.
<https://doi.org/10.1109/ISQED.2010.5450531>

DOI:

[10.1109/ISQED.2010.5450531](https://doi.org/10.1109/ISQED.2010.5450531)

Document status and date:

Published: 01/01/2010

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

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 Synthesis for Optimum Performance per Area

Maurice Meijer¹, Jose Pineda de Gyvez^{1,2}

¹NXP Semiconductors, Eindhoven, The Netherlands

²Technical University of Eindhoven, Eindhoven, The Netherlands

{maurice.meijer, jose.pineda.de.gyvez}@nxp.com

Abstract

Worst-case design uses extreme process corner conditions which rarely occur. This costs additional power due to area over-dimensioning during synthesis. We present a new design strategy for digital CMOS IP 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. Dynamic power is reduced depending upon the ratio of flip-flops to logic-gates, and data activity. On a set of benchmark circuits in 65nm LP-CMOS, we observed performance-per-area improvements up to 81%, area and leakage reductions up to 38%, and total power savings of up to 26% without performance penalties.

Keywords

CMOS, logic synthesis, body biasing, performance, area

1. 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 trade-off involved during logic synthesis. Observe that circuit area depends on the process margin. If a lower process margin can be tolerated without a parametric yield penalty, circuit performance can be increased without spending excessive area. Statistical circuit design has long been seen as a viable way to avoid the use of worst-case parameters [2-3]. Yet these approaches have not totally found their way in industrial practices. This is because, among other reasons, the moving average of process parameters, the flexibility of fabrication of the same chip design in multiple foundries, and the lack of appropriate EDA tools for statistical logic synthesis. In this paper we show that a body bias driven logic synthesis overcomes these drawbacks.

A way out to avoid the previously mentioned weaknesses has been the use of post-silicon tuning. Basically, post-silicon tuning approaches have been proposed for improving product-binning yields and for trading-off power-performance [4-5], but do not eliminate the problem of area over-dimensioning. Well-known approaches are: supply

voltage scaling (VS) and body biasing (BB). VS is primarily used to reduce active power at the expense of a lower circuit performance [4]. BB is typically used for leakage reduction or performance tuning [4-5]. Forward body biasing (FBB) is preferred over VS to achieve increased performance [4]. This is because the power penalty of FBB is lower in case of dynamic-power dominant designs. Leakage power of digital IP blocks is only a concern when the circuit is in standby. Moreover, FBB needs only to be applied to those die samples with a lower speed than the nominal process outcome. Such samples have already a low intrinsic leakage power.

Figure 1: Area-Clock Period Trade-Off at Logic Synthesis.

A joint design-time and post-silicon tuning optimization strategy for minimizing leakage under delay constraints was proposed in [6]. This approach relies on detailed process variability inputs, and is capable of reducing process-dependent delay spread. However, it does neither consider a timing speed-up nor a circuit area reduction as outcome. Other works propose body bias clustering at design-time for minimizing leakage under delay constraints [7-8], or enhancing circuit performance [9]. These approaches do not consider a (joint) design-time optimization for improving performance or reducing area of the circuit.

High-performance circuits typically use low- V_{th} devices to speed-up critical delay paths at the cost of an intrinsic higher device leakage [10]. The application of FBB offers additional benefits. FBB can be used to further enhance low- V_{th} performance. Alternatively, it can eliminate the use of multiple V_{th} options. Moreover, FBB can achieve low- V_{th} performance during operation with lower standby leakage when it is used dynamically at run-time.

In this work we leverage FBB to improve the performance-per-area (PPA) ratio of digital CMOS circuits. We enhance state-of-the-art solutions by enabling logic synthesis with FBB under bounded process variation influences. Given a FBB range, our approach finds the best PPA ratio that meets a target performance specification. Pre-

silicon design optimization is done by selecting the appropriate synthesis point in between worst-case and best-case process conditions given a FBB range. Moreover, as with other post-silicon approaches, FBB can be applied dynamically at run time to speed up slow chip samples. The reason for this is to minimize leakage overhead related to FBB during standby operation. We show that our approach renders smaller area and lower-power circuits at no performance penalty despite their fabrication in a process corner other than the nominal one. In summary, the contributions of this paper are the following:

- A new body bias driven gate-level optimization method is proposed to improve performance per area of digital integrated circuits.
- A new approach to evaluate the design's quality based on the performance per area metric.
- Full integration of our approach with a state-of-the-art commercial design flow.

The rest of this paper is organized as follows. In Section 2 we introduce body bias driven design. Section 3 presents the theoretical background and modeling. Finally, Section 4 shows our benchmarked results.

2. Body Bias Driven Digital Design 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 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.

Figure 2: FBB utilization under body bias driven design.

Fig.2 illustrates the parameters that are under control with body bias driven design (BBD). The right-hand side of Fig.2 plots the dependency between clock period and FBB. The results have been obtained experimentally for a 65nm LP-CMOS standard- V_{th} ring-oscillator test structure [4]. Up to 20% performance increase was measured when 0.4V FBB is applied to both N- and P-wells simultaneously. The left-

hand side of Fig.2 plots the relationship between circuit area and relative clock period. For increasing FBB values, the trade-off curve shifts linear 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.

3. Optimal Performance-per-Area Design

In this section we present the theoretical background of BBD design for achieving an optimum PPA ratio. We explore area, performance and power trends.

3.1. Design for Body Bias Driven Optimum PPA

The delay of a digital logic gate can be modeled as:

$$d_{gate} = \frac{(xC_{intr} + C_{load})V_{DD}}{xI_{drive}} = d_0 + \frac{d_1}{x} \quad (1)$$

where x is the gate sizing factor ($x \geq 1$), C_{intr} and C_{load} are the intrinsic and load capacitance of a gate, respectively. I_{drive} is the current drive of a gate, and depends on both V_{DD} and V_{th} . Parameters d_0 and d_1 represent the intrinsic and load-dependent gate delays, respectively, as can be inferred from expression (1). FBB impacts the delay of the circuit. From experimental results [4], we model the normalized delay dependence on FBB by a linear function as follows

$$delay_{norm} = 1 + k_1 V_{BB} \quad (2)$$

The delay at various FBB conditions has been normalized to the case of nominal body bias. V_{BB} represents the FBB value: $V_{BB} = V_{pwell} = V_{DD} - V_{nwell}$. Parameter k_1 is the polynomial coefficient, which is different for each gate. The maximum error of expression (2) was found lower than 1.5% for 65nm LP-CMOS test-structures [4].

Combining (1) and (2), we model the delay and area of a CMOS digital logic circuit as:

$$D_j = \sum_{i \in j} \left(d_{0i} + \frac{d_{1i}}{x_i} \right) \cdot (1 + k_{1i} V_{BB}) \leq T_{ck} \quad \forall j \in \Psi \quad (3)$$

$$A_{total} = \sum_{i=1}^m x_i A_i \quad (4)$$

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 (3) constrains the delay of each circuit path to be less than the targeted clock period, T_{ck} .

Circuit performance and area are key performance metrics for digital circuit designers. Therefore, we based our design synthesis on the PPA metric to qualify the design for performance while accounting for over-dimensioning. This metric depends on the CMOS technology and available standard cells in which the circuit is synthesized. Let $f_{ck} = 1/T_{ck} = 1/\max(D_j)$. We obtain

$$PPA = \frac{f_{ck}}{A_{total}} = \frac{1}{T_{ck} A_{total}} \quad (5)$$

A higher PPA value indicates that the circuit design utilizes silicon area more effectively to achieve a high performance. In our analysis, we made use of a normalized representation of PPA. The normalization has been done against the highest performing circuit under WCD ($f_{ck}=f_{max}=1/T_{min}$, $A_{total}=A_{max}$).

$$PPA_{norm} = \frac{f_{ck}}{f_{max}} \cdot \frac{A_{max}}{A_{total}} = \frac{T_{min}}{T_{ck}} \cdot \frac{A_{max}}{A_{total}} \quad (6)$$

The actual value for T_{min} can be found by correlating the targeted clock period and the one obtained from static timing analysis of the synthesized design. Two regions can be clearly identified, namely, a region where a good correlation occurs, and a region where the actual clock period can no longer meet the targeted clock period. T_{min} is found at the border of these regions. Our criterion for T_{min} is a maximum deviation of 5% between targeted clock period and the one obtained after synthesis.

Figure 3: Area, and clock period trade-off for a generic digital logic circuit.

Fig.3 shows a typical trade-off curve for a generic digital logic circuit. The curve is composed out of a multitude of designs that are synthesized to meet a distinct clock period constraint, T_{ck} . The area and clock period have been normalized to the best performing design (A_{max} , T_{min}). Observe that high-performance circuits consume more area than slow circuits. This is due to gate upsizing to speed-up critical circuit paths. The trend shown in Fig.3 can be modeled by a rational function with χ , δ , and η as fitting parameters.

$$A_{total} = \frac{\chi}{\delta + T_{ck}} + \eta \quad (7)$$

There exists a point on (7) with an optimum PPA. This point indicates the lowest clock period without circuit over-dimensioning. By combining (5) and (7), we obtain

$$PPA(T_{ck}) = \frac{1}{T_{ck} \left(\frac{\chi}{\delta + T_{ck}} + \eta \right)} \quad (8)$$

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.

$$T_{best} = -\delta + \frac{\sqrt{-\delta\chi\eta}}{\eta} \quad \forall T_{ck} \geq T_{min} \wedge \delta\chi\eta \leq 0 \quad (9)$$

$T_{ck} > T_{best}$, yields circuits without area over-dimensioning, and the contrary holds true for $T_{ck} < T_{best}$. Therefore, T_{best} identifies the minimum clock period possible without circuit over-dimensioning. Under WCD, T_{best} may be too large for high-performance designs to meet the target frequency spec. In this case, over-dimensioning cannot be avoided, thereby worsening PPA.

Figure 4: Area, clock period, and performance-per-area trade-off for a generic digital logic circuit under BBD and WCD. Solid line: WCD, dotted line: BBD, overlay: PPA.

Next, we investigate area, clock period and PPA trends for WCD and BBD design styles. For this purpose, we took a generic digital logic circuit with calibrated technology parameters for 65nm LP-CMOS. For BBD, we utilized a maximum FBB of 0.4V. Fig.4 shows the design synthesis exploration space for circuit area, clock period and PPA. The area and clock period curves are plotted for the WCD (solid line), and the BBD (dash-dotted line). The iso-PPA curves are plotted as overlay; the intersection with the area-clock period curves represents the normalized PPA ratio of the design. Since logic synthesis aims usually at a target speed, as way of example, all PPA values of Fig.4 have been normalized to the maximum frequency circuit design under WCD ($T_{ck}=T_{min}$). The triangle is located at a clock period of T_{min} , while the circles relate to T_{best} .

Observe from Fig.4 that BBD achieves a better PPA ratio than WCD under all circumstances. For a given circuit area, BBD achieves higher performance than WCD. Alternatively, BBD enables lower area designs for a given clock period. For a FBB of less than 0.4V FBB, the area-clock period curve would be located in between the two curves plotted in Fig.4. Therefore, it makes most sense to use BBD with a maximum FBB to obtain the best PPA ratio.

3.2. Power Implications

The power consumption of a digital logic gate can be modeled as:

$$P_{gate} = a(xC_{intr} + C_{load})V_{DD}^2 f_{ck} + xI_{leak}V_{DD} \quad (10)$$

where a is the switching activity of the gate, 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 [4], we model the normalized leakage current dependence by a fourth-order polynomial expression as follows

$$leakage_{norm} = 1 + \sum_{n=1}^4 l_n V_{BB}^n \quad (11)$$

The leakage at various FBB conditions has been normalized to the case of nominal body bias. As before, V_{BB} represents the FBB value: $V_{BB} = V_{pwell} = V_{DD} - V_{nwell}$. Parameters l are the polynomial coefficients, which are different for each gate. The maximum error of expression (11) is lower than 6% for 65nm LP-CMOS test-structures [4].

Combining (10) and (11), we model the power consumption of a CMOS digital logic circuit as:

$$P_{total} = V_{DD} \sum_{i=1}^m \left(a_i (x_i C_{intr,i} + C_{load,i}) V_{DD} f_{ck} + x_i I_{leak,i} \left(1 + \sum_{n=1}^4 l_n V_{BB}^n \right) \right) \quad (12)$$

where i is an index that runs over all gates in the circuit.

We investigated the relationship between area, clock period and power for WCD and BBD. The analysis was done at $V_{DD}=1.2V$ and $T=85^\circ C$. Fig.5 shows the design exploration space for the same circuit as before. The iso-power curves are plotted as overlay; their intersection with the area-clock period curves represents the power of the design. Notice that BBD enables lower power operation at a constant clock period. For a given power target, BBD offers better performance and area figures.

Figure 5: Area, clock period, and power trade-off for a generic digital logic circuit under BBD and WCD. Solid line: WCD, dotted line: BBD, overlay: power consumption.

The application of FBB increases leakage power significantly. This is a concern when the circuit is in standby operation. Therefore, we combine BBD with dynamic FBB. No FBB is applied to the circuit during standby.

4. Benchmarked Results

Commercial synthesis tools can target area optimization subject to delay constraints. To validate our approach, we have implemented BBD in Cadence's commercial logic synthesis tool. To enable BBD, digital cell libraries are required with FBB-characterized timing views. In our case, BBD is based on 0.4V FBB for the whole design. BBD and

WCD have been analyzed and compared for sixteen circuits of the ITC99 benchmark suite [11]. The circuits have been mapped on 65nm LP-CMOS to operate at $V_{DD}=1.1V$, and $T=85^\circ C$. The area results after synthesis have been corrected with a row utilization factor of 0.9 to account for layout effects. The total and leakage power of the circuit has been determined at $V_{DD}=1.2V$, $T=85^\circ C$, and a low data activity of 5%. Two different synthesis cases have been investigated. The first case concerns design synthesis for maximum PPA independent of the chosen design style. The second case concerns the design synthesis for maximum frequency under WCD. In the latter case, BBD is done to operate at the same speed at a lower area cost to improve the PPA ratio.

4.1. Model Validation

This section provides detailed information on circuit area, clock period, PPA and power trends for ITC99 benchmark circuit b11. The circuit contains 31 flip-flops and about 700 combinational gates. Fig.6 shows the design exploration space between circuit area versus clock period. The results obtained from synthesis, have been indicated by circles and triangles for WCD and BBD, respectively. The solid and dotted lines show the corresponding results from expression (8) when combined with least-squares regression. The fitting parameters of the model are shown in Table 1.

Figure 6: Area versus clock period for the b11 circuit in 65nm LP-CMOS. Lines: WCD (solid) and BBD (dotted) model, symbols: synthesis results. The PPA ratio is indicated for each synthesized design.

Table 1: Model fitting parameters for b11 circuit

	χ	δ	η
WCD	169.34	-1.04	1900.8
BBD	181.85	-0.81	1734.9

Observe from Fig.6 the close match between the modeled and the synthesized area-clock period trends. From (10), we have calculated a T_{best} value of 1.34ns and 1.11ns for WCD and BBD, respectively. This matches with those obtained coarsely through synthesis (WCD: 1.39ns, BBD: 1.2ns). Moreover, we found similar PPA trend as presented before. The PPA value for each synthesis point has been indicated in Fig.6 normalized w.r.t T_{min} under WCD ($T_{min}=1.13ns$).

Table 2: Design synthesis results for maximum PPA - ITC99 benchmark circuits in 65nm LP-CMOS. Relative values are shown w.r.t. WCD for the process condition that is indicated in the row “Process”.

Design Process	Clock period		Area		PPA		Total power (1.2V V_{DD} , 85°C)		Leakage power (1.2V V_{DD} , 85°C)		
	WCD	BBD	WCD	BBD	WCD	BBD	WCD	BBD	WCD	BBD	
	[ns]	rel.	[μm^2]	rel.			slow,nom,fast [μW]	all rel.	slow,nom,fast [nW]	slow rel.	nom,fast rel.
b01	0.80	0.98	207	0.87	1	1.18	156, 157, 159	0.98	41.3, 161, 850	4.32	0.86
b02	0.89	0.90	129	0.93	1.17	1.39	104, 105, 106	1.11	25.8, 101, 531	4.65	0.93
b03	0.92	0.88	861	0.92	1	1.24	742, 743, 751	1.11	171, 665, 3516	4.59	0.92
b04	1.65	0.85	3460	0.93	1	1.26	1530, 1540, 1560	1.16	686, 2671, 14116	4.69	0.94
b05	1.89	0.82	3530	0.93	1.02	1.34	841, 844, 859	1.18	700, 2728, 14420	4.64	0.93
b06	0.80	0.99	260	0.98	1	1.02	255, 256, 259	1.00	51.6, 201, 1062	4.92	0.98
b07	1.30	0.84	1710	1.00	1.13	1.34	924, 928, 940	1.19	339, 1321, 6982	5.01	1.00
b08	1.28	0.78	946	1.10	1.23	1.44	701, 703, 711	1.30	188, 731, 3865	5.51	1.10
b09	0.92	0.98	868	0.73	1	1.40	974, 977, 1000	0.70	172, 671, 3547	3.65	0.73
b10	1.10	0.81	693	1.00	1.20	1.47	432, 433, 438	1.23	138, 536, 2833	5.02	1.00
b11	1.39	0.86	2254	0.94	1.30	1.60	701, 703, 714	1.13	447, 1742, 9208	4.71	0.94
b12	1.31	0.61	4218	0.96	1	1.70	2200, 2210, 2230	1.62	838, 3264, 17253	4.80	0.96
b13	1.10	0.73	1380	1.04	1.02	1.34	1079, 1080, 1090	1.38	273, 1063, 5616	5.24	1.05
b14	2.99	0.91	46739	0.80	1.32	1.81	5660, 5690, 5840	0.98	9290, 36183, 191248	4.02	0.80
b15	2.06	0.74	33671	1.06	1.03	1.31	7250, 7280, 7410	1.38	6685, 26036, 137618	5.31	1.06
b17	2.06	0.76	101667	0.99	1.06	1.42	22400, 22500, 22900	1.32	20177, 78588, 415383	4.97	0.99
Average (relative)		0.84		0.95	1.09	1.39		1.17		4.75	0.95

Fig.7 shows the same area and clock period trends as before, but now with the normalized power consumption for each design as overlay. Observe that the power consumption trend is similar as found before, as illustrated in Fig.5.

4.2. Design Synthesis for Maximum PPA

Table 2 shows the results obtained the benchmark circuits when synthesizing for maximum PPA under WCD and BBD. The process condition for which the results have been obtained is indicated as well. All BBD results are made relative to the WCD results for the corresponding process condition. For each circuit, the PPA ratio has been normalized to maximum performance design ($T_{ck}=T_{min}$).

Figure 7: Area versus clock period for the b11 circuit in 65nm LP-CMOS. Lines: WCD (solid) and BBD (dotted) model, symbols: synthesis results. The power consumption is indicated for each synthesized design.

Observe that the PPA ratio can differ for each benchmark circuit. This depends on circuit characteristics such as path delay distribution, and logic depth. Under WCD, we found a maximum PPA ratio ranging from 1 to 1.32 (1.09 on average). The benefits for BBD are higher (1.02-1.81; 1.39

on average). For a given circuit, BBD provides always a higher maximum PPA ratio than WCD. All BBD circuits operate faster than their WCD counterparts. Moreover, most BBD circuits are smaller.

The total power is dominated by dynamic power consumption, even in the fast process corner and $T=85^{\circ}C$. It is not much process-dependent. Observe that the total power for BBD is generally higher than under WCD. This is mainly because of the higher operating frequency for BBD. In case of a lower total power for BBD, the circuits operate at a similar frequency but have a smaller area. For the considered circuits, the BBD total power ranges from 0.7 to 1.62 times the total power of the WCD. The leakage power for BBD decreases by the same factor as the circuit area for nominal and fast process conditions. For slow process and active mode (non-standby) operation, the BBD leakage power is higher than the WCD leakage power due to utilization of FBB (3.65x-5.51x higher). Recall that we apply dynamic FBB during chip operation. In this way we avoid the leakage penalty associated to FBB during standby operation.

4.3. Design Synthesis for Optimum Area

Table 3 shows the results for the benchmark circuits when synthesizing for maximum performance under WCD. The BBD circuits are synthesized to match the WCD performance. Table 3 uses a similar set-up as Table 2.

Observe that BBD circuits enable large area savings when designed for maximum WCD frequency. The area reduction ranges from 2% to 35% as compared to the WCD circuit (21% on average). The lower area comes mostly from the area scaling of the combinatorial logic. In general, BBD circuits have less logic gates than WCD ones, while the amount of flip-flops is the same. The largest area savings have been obtained for the b11 and b14 circuits, which have 21-28x more logic gates than flip-flops. This ratio is lower for the other circuits. The PPA ratio scales inversely proportional to area. For BBD, the PPA ranges from 1.02 to 1.61 for the benchmark circuits (1.28 on average).

Table 3: Design synthesis results for maximum frequency with WCD - ITC99 benchmark circuits in 65nm LP-CMOS. Relative values are shown w.r.t. WCD for the process condition that is indicated in the row "Process".

Design Process	Clock	Area		PPA		Total power (1.2V V _{DD} , 85°C)		Leakage power (1.2V V _{DD} , 85°C)		
		WCD	BBD	WCD	BBD	WCD	BBD	WCD	BBD	
	[ns]	[μm^2]	rel.			slow,nom,fast [μW]	all rel.	slow,nom,fast [nW]	slow rel.	nom,fast rel.
b01	0.80	208	0.87	1	1.15	156, 157, 159	0.96	41.3, 161, 850	4.32	0.86
b02	0.80	169	0.72	1	1.39	126, 126, 127	0.91	33.4, 130, 688	3.59	0.72
b03	0.92	861	0.85	1	1.18	730, 743, 751	0.97	171, 665, 3516	4.26	0.85
b04	1.65	3460	0.86	1	1.17	1510, 1540, 1560	0.97	686, 2671, 14116	4.29	0.86
b05	1.87	3631	0.77	1	1.29	862, 871, 880	0.89	720, 2805, 14824	3.86	0.77
b06	0.80	260	0.98	1	1.02	255, 256, 259	1.00	51.6, 201, 1062	4.92	0.98
b07	1.13	2235	0.76	1	1.31	1160, 1160, 1170	0.92	442, 1723, 9107	3.84	0.77
b08	1.12	1333	0.75	1	1.34	855, 858, 868	0.95	263, 1024, 5414	3.76	0.75
b09	0.92	868	0.73	1	1.37	707, 709, 717	0.95	173, 671, 3547	3.65	0.73
b10	1.02	895	0.72	1	1.39	503, 505, 511	0.90	178, 692, 3658	3.61	0.72
b11	1.13	3617	0.65	1	1.54	1130, 1130, 1150	0.78	718, 2795, 14774	3.26	0.65
b12	1.31	4219	0.87	1	1.15	2200, 2210, 2230	0.96	838, 3264, 17253	4.36	0.87
b13	1.00	1549	0.83	1	1.20	1210, 1210, 1230	0.97	307, 1197, 6324	4.17	0.83
b14	2.77	66502	0.62	1	1.61	7680, 7720, 7930	0.74	13197, 51403, 271694	3.12	0.62
b15	1.95	36678	0.84	1	1.19	7990, 8020, 8170	0.92	7274, 28334, 149761	4.21	0.84
b17	2.00	111372	0.80	1	1.26	24100, 24200, 24700	0.90	22118, 86150, 455353	3.99	0.80
Average (relative)			0.79	1	1.28		0.92		3.95	0.79

BBD renders both lower total power and leakage power. Observe that the BBD total power is generally lower than in case of WCD when operating at the same frequency. For the considered circuits, one can see total power savings of up to 26% for BBD. BBD primarily affects logic gates in the data path, thus the clock power is not much reduced. We observed that the power savings are larger for higher data activities. For a data activity of 30% instead of 5%, the total power savings are up to 35% for BBD (not shown in Table 3). The leakage savings of the BBD circuits are in between 2-38% when FBB is not enabled. Observe that the leakage power reduces more than the total power for all considered circuits. For slow chip samples, the leakage power increases up to 4.92x with FBB. Recall that this leakage increase is of no concern since FBB is disabled during standby operation.

5. Conclusions

We presented a new design strategy for digital CMOS IP 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. Dynamic power is reduced depending upon the ratio of flip-flops to logic-gates, and data activity. On a set of benchmark circuits in 65nm LP-CMOS, we observed performance-per-area improvements up to 81%, area and leakage reductions up to 38%, and total power savings of up to 26% without performance penalties as a benefit from our proposed body bias driven design strategy.

6. References

- [1] J. Zhang, "Worst Case Design of Digital Integrated Circuits," *Proc. of ISCAS*, London, UK, June 1994, pp.153-156.
- [2] S. Duvall, "A Practical Methodology for the Statistical Design of Complex Logic Products for Performance," *IEEE Trans. on VLSI Systems*, Vol.3, No.1, March 1995, pp.112-123.
- [3] A.Nardi et al., "Impact of Unrealistic Worst Case Modeling on the Performance of VLSI Circuits in Deep Submicron CMOS Technologies," *IEEE Trans.*

on Semiconductor Manufacturing, Vol.12, No.4, November 1999, pp.396-403.

- [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 Ed., 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," *Proc. of ISSCC*, San Francisco, CA, USA, February 2002, pp.344-345.
- [6] M. Mani et al., "Joint Design-Time and Post-Silicon Minimization of Parametric Yield Loss using Adjustable Robust Optimization," *Proc. of ICCAD*, San Jose, CA, USA, November 2006, pp.19-26.
- [7] S. Kulkarni et al., "A Statistical Framework for Post-Silicon Tuning through Body Bias Clustering," *Proc. of ICCAD*, San Jose, CA, USA, Nov.2006, pp.39-46.
- [8] R. Teodorescu et al., "Mitigating Parameter Variation with Dynamic Fine-Grain Body Biasing," *Proc. of MICRO-40*, Chicago, IL, USA, Dec.2007, pp.27-39.
- [9] A. Sathanur et al., "Physically Clustered Forward Body Biasing for Variability Compensation in Nanometer CMOS design," *Proc. of DATE*, Nice, France, April 2009, pp.154-159.
- [10] M. Hirabayashi et al., "Design Methodology and Optimization Strategy for Dual-V_{TH} Scheme using Commercially Available Tools," *Proc. of ISLPED*, Huntington Beach, CA, USA, Aug. 2001, pp.283-286.
- [11] ITC99 benchmarks: www.cad.polito.it/tools/itc99.html