Minimum Energy Point in Constant Frequency Designs under Adaptive Supply Voltage and Body Bias Adjustment in 55 nm DDC

Christoph Thomas Müller*†, Marc Pons*, David Ruffieux*, Jean-Luc Nagel*, Stéphane Emery*, Andreas Burg†, Shuuji Tanahashi,‡ Yoshitaka Tanaka‡, Atsushi Takeuchi‡ *CSEM SA, Neuchâtel, Switzerland *{christoph.mueller, marc.pons, david.ruffieux, jean-luc.nagel, stephane.emery}@csem.ch* †EPFL, Lausanne, Switzerland *{christoph.mueller, andreas.burg}@epfl.ch* ‡Mie Fujitsu Semiconductor Limited, Yokohama, Japan *{tanahashi.shuuj, yoshi-tanaka, a-takeuchi}@jp.fujitsu.com*

***Abstract*—In this paper, we describe a systematic low-power design methodology for technologies that offer a strong body factor. Specifically, we explore both the body bias voltage and the supply voltage knobs in order to find the MEP (minimum energy point) for a constant target frequency. Our methodology accounts for process and temperature (PT) variations while charting the design space for a simple reference design. We then show how to scale the energy data of this reference design to any arbitrary design. A case study of a 32 bit RISC microprocessor achieves an energy estimation match of our significantly less complex estimation methodology within 1% of traditional signoff results.**

I. Introduction

Technologies with a strong body factor such as MIFS 55 nm deeply depleted channel (DDC) [1] provide the designer with an additional knob for energy optimisation in addition to the widely used adaptive voltage and frequency scaling (AVFS) techniques for which an analytical energy model was presented by Jain et al.[2]. The goal of the designer is to operate as close as possible to the operating point where the sum of the leakage energy and the dynamic energy is minimal (MEP). Typically this MEP is located within sub- or near threshold operation, which significantly reduces the maximum operating frequency[3], [4]. However, with the body voltage as an additonal knob, adaptive supply and body bias scaling (ASVBB) can be implemented, allowing for MEP tracking at a given constant frequency which is determined only by the application requirements[5]. Recently Lee et al.[6] presented such a scheme to automatically track the MEP at a given frequency by pinning the ratio of leakage to dynamic energy into a predefined range. In light of the efficiency of this approach and the need for real-time systems that require a given frequency, there is a need for a methodology to rapidly explore the associated design. Such a methodology rapidly identifies the achievable limits and provides near-optimal settings for adjusting supply voltage and body bias.

Contributions: This paper describes a design methodology to find the MEP for an arbitrary design using ASVBB by selecting the operating point for a given frequency, where both the power and—due to the fixed frequency—also the energy are minimized. The methodology estimates the dynamic and leakage energy properties based on a small reference design implemented with a heavily pruned standard cell library across the ASVBB design space including process and temperature. We then show how the results obtained for the reference design can be scaled to any complex design to identify the MEP for which the designer can then characterize the full library for fine-tuning and signoff. Our analysis also shows the strong dependency of the MEP on PT variations and illustrates the need for controlling both the body bias and the supply voltage.

The paper is organized as follows: Section II defines the design space. Section III presents our method to map dynamic and leakage power as well as speed for a simple reference circuit. In Section IV we scale the reference design to an arbitrary design and discuss the influence of PT variation. In Section V we apply that method to a 32 bit microprocessor before concluding the paper.

II. Mapping the design space

We start by charting the ASVBB design space, which defines the spread of leakage, dynamic energy, and frequency operating points.

Using ASVBB we may vary the supply voltage with the lower limit set by the reliability of the circuit and the upper limit set by the process. The legal body bias range $\mathcal{V}(V_{dd})$ is set by the supply voltage as well as technology parameters, most notably the structure and forward voltages of the built-in body diodes between NWELL and PWELL. Furthermore, the valid range under nominal conditions is reduced when PT variations have to be compensated with legal settings. For the 55nm DDC process used in this work, the strong body factor of 375 mV/V, allows to fully compensate across PT while still retaining a large frequency design space for MEP tracking [7].

For the purpose of ASVBB design space exploration we use the on-current I_{ON} as a proxy for frequency, based on the following first order aproximation where f denotes the frequency, C the load capacitance, and V_{dd} the supply voltage.

$$f = {}^{I_{ON}}/_{C \cdot V_{dd}} \tag{1}$$

Contrary to the frequency, I_{ON} can easily be measured and simulated while the PMOS and NMOS strength can be trivially fixed to a constant ratio. Furthermore the I_{ON} can be kept constant across PT, as the adjustment of the threshold voltage through the body voltage counteracts the effects of variation. The design space limits are then determined through spice simulations of I_{ON} through a representatively sized transistor under application of the maximum forward ($V_{bb}^{fwd} = 0.6\,\mathrm{V}$)

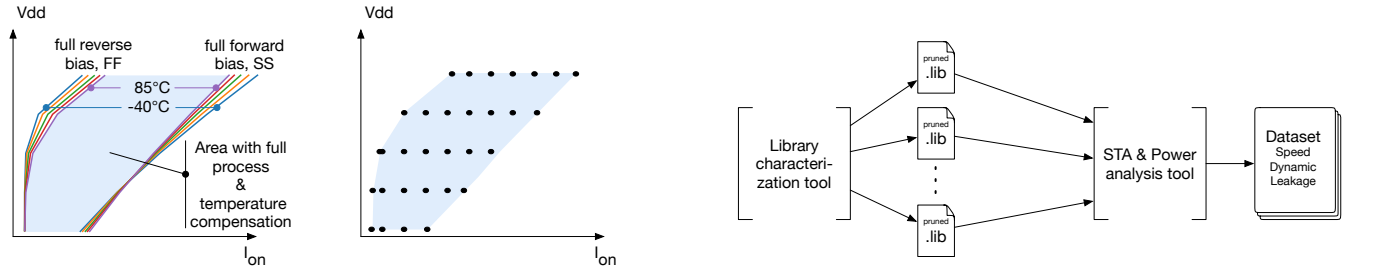

(a) Find compensation boundaries for process corners of interest

(b) Sample the design space

(c) Characterize a pruned library across the design space.

(d) Extract reference design speed, dynamic power and leakage.

Fig. 1: Steps to extract the reference data over the design space.

and reverse ($V_{bb}^{rev} = -1\,\mathrm{V}$) bias voltages, for the expected operation temperatures $T \in \{-40\,^{\circ}\mathrm{C}, 25\,^{\circ}\mathrm{C}, 85\,^{\circ}\mathrm{C}\}$, and across the process corners $P \in \{FF, SS, TT\}$ of interest. The low end of the design space is defined by the highest I_{ON} for which PT can still be compensated (i.e., maximum reverse bias):

$$I_{ON}^{min}(V_{dd}) = \max_{p \in P, t \in T} \min_{V_{bb} \in \mathcal{V}(V_{dd})} I_{ON}\left(V_{dd}, V_{bb}, p, t\right) \quad (2)$$

Similarly, the upper end of the design space is defined by the minimum I_{ON} for which PT can still be compensated (i.e., maximum forward bias):

$$I_{ON}^{max}(V_{dd}) = \min_{p \in P, t \in T} \max_{V_{bb} \in \mathcal{V}(V_{dd})} I_{ON}\left(V_{dd}, V_{bb}, p, t\right) \quad (3)$$

Figure 1a shows the boundaries set by the sweep, with the ASVBB design space spanned between the minimum I_{ON} (2) and the maximum (3) for each supply voltage. We notice that the design space across corners is very similar to the design space of the individual corners, thanks to the large body factor.

III. Reference circuit

After defining the boundaries of the ASVBB design space, we proceed to characterize normalized active energy, leakage power, and timing within this valid range. As shown in Fig. 1b, we propose to sample the ASVBB design space with a grid defined by sweeping a) the supply voltage and b) I_{ON} from minimum to maximum. For each grid point PMOS/NMOS bias voltage pairs are obtained. With these bias voltage pairs a standard cell library can be characterized for each grid-point and for each PT combination. By using these libraries, we can easily search the design space using power and timing analysis for the optimum operating points of the target design.

A. Standard cell library characterisation

Unfortunately, a full standard cell library characterisation for many operating points is often prohibitive, in particular if a dense grid is expected for high precision, due to long simulation times. We therefore propose to limit the analysis to only a few cells, representing a distribution of cells commonly found in a large design. The most significant reduction in characterisation time can be achieved by dropping sequential cells from the library. These need to be characterised for setup and hold constraints which comes with the need of extensive sweeps. With non-sequential cells accountable for the majority of the delay in the critical path of a typical design, the effect of potential scaling differences between sequentials and non sequentials can be considered neglectable.

To select a representative subset, a purely combinatoric 32 bit multiplier was synthesized with RTL compiler, constrained to a maximum delay of zero from all inputs to all outputs. The cell distribution was extracted, rarely used cells where pruned, and the circuit was resynthesized using only the top thirteen cells, which cover over 90% of the cells in the original design. This pruning approach reduced the simulation time of our characterisations from slightly short of two hours to less than three minutes, allowing to characterize 55 points across three process corners and five temperatures on a single machine over a weekend. Finally, leakage, dynamic power as well as the frequency can be extracted for the reference design after re-synthesis with only the selected cells.

B. MEP for the 32 Bit Multiplier reference design

Using interpolation, we can now approximate the dynamic and leakage power across the whole ASVBB design space, which allows us to find the MEP for operating modes of interest. Figure 2a shows the tradeoffs for the multiplier reference design, tracing the minimum energy point across the range of target frequencies from 20 MHz to 140 MHz. On the low end the minimum supply voltage of our design space is the limiting factor, with the bias as far reverse as possible to still reach the frequency. Subsequently only the bias increases until an inflection point is reached around 50 MHz where the supply voltage starts to increase, causing an inversion of the bias effectively reducing the leakage. At the 80 MHz frequency cut the typical MEP curve is formed, with the higher forward bias at low supply voltage causing the leakage to dominate below 0.6 V and the dynamic dominating at a higher supply. At the 120 MHz frequency cut the design space is limited again, with the speed not being reachable at a supply voltage below 0.7 V.

Similarly, Fig. 2b shows the shift of the MEP across frequency for different process corners, when adjusting both bias and supply voltage for minimum energy at each frequency step. Note that, as long as a tradeoff exists, the absolute energy values for MEP are almost constant over PT for constant frequency, as the same I_{ON} results in the same dynamic energy while the leakage is kept under control. However, when

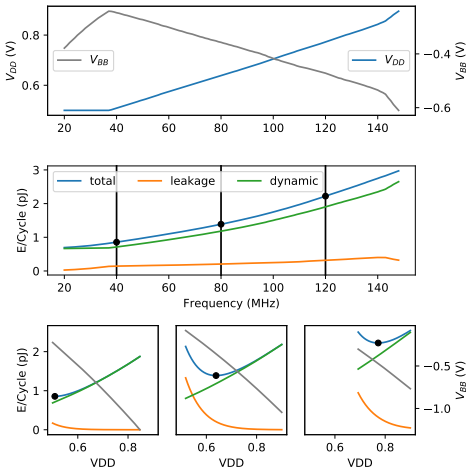

(a) Top: The bias and supply voltage pairs for minimum energy across frequency at TT/25 °C. Center: The corresponding leakage and dynamic Energy components. Bottom: minimum energy curves for the three marked frequency cuts in the center plot.

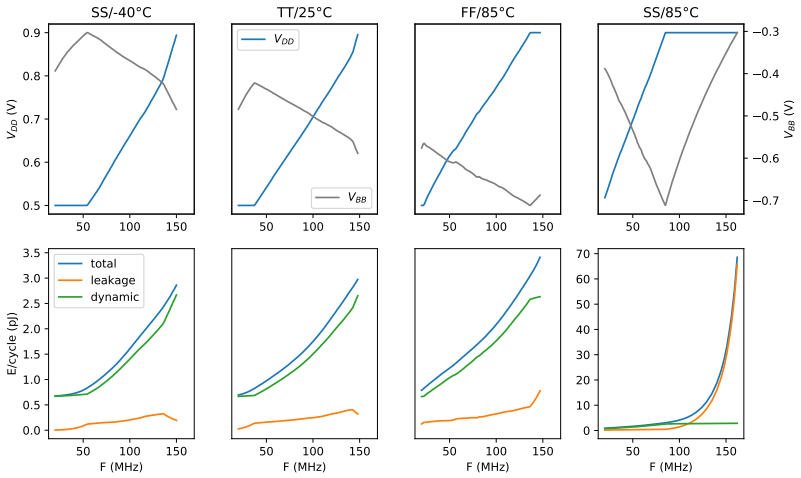

(b) Top: Supply and bias voltages at the MEP for different target frequencies in four process corners. Bottom: Energy across frequency at these points, broken down into the leakage and dynamic components.

Fig. 2: Minimum Energy point for the 32 bit multiplier shifts across frequency

the edge of the design space is reached as illustrated by the SS/85 °C case in Fig. 2b only biasing remains for frequency adjustment causing the leakage to explode.

IV. Scaling the Reference Circuit to an Arbitrary Design

Obviously, the MEP and the associated optimal settings supply voltage and body bias are highly design dependent since the share between leakage and dynamic energy is determined by the activity and the critical path length[4]. Hence the analysis has to be done for each design as well as for each utilisation profile. We propose to use the small reference design with a limited number of cells to calibrate a model that can then be used to extract the MEP and the associated ASVBB settings for different target frequencies. In the following we discuss the corresponding scaling model.

A. ASVBB model

The per cycle energy of a design can be split into the sum of the dynamic contribution E_{dyn} and the leakage contribution E_{leak} so that $E_{tot} = E_{dyn} + E_{leak}$. The two energy components are obtained with the following first order approximations where α_{sw} denotes the design specific switching activity, C_L the total load capacitance, I_{OFF} the leakage current, V_{DD} the supply voltage, and t_p the cycle time:

$$E_{dyn} = \alpha_{sw} \cdot C_L \cdot V_{DD}^2 \quad (4)$$

$$E_{leak} = V_{DD} \cdot I_{OFF} \cdot t_p. \quad (5)$$

The dynamic energy ratio r_{dyn} between the dynamic energy of the target design E_{dyn}^{target} and the dynamic energy of the reference design design E_{dyn}^{ref} is defined as

$$r_{dyn} = \frac{E_{dyn}^{target}}{E_{dyn}^{ref}} = \frac{\alpha_{sw}^{target} \cdot C_{tot}^{target}}{\alpha_{sw}^{ref} \cdot C_{tot}^{ref}}, \quad (6)$$

where α_{sw}^{target}, and C_{tot}^{target} as well as α_{sw}^{ref}, and C_{tot}^{ref} denote the load and activity of the target and reference design, respectively.

Similarly, we define the leakage energy ratio between the leakage of the target design E_{target}^{ref} and the leakage of the reference design E_{leak}^{ref} as

$$r_{leak} = \frac{E_{leak}^{target}}{E_{leak}^{ref}} = \frac{g_{target} \cdot \overline{I_{OFF}^{target}} \cdot t_p^{target}}{g_{ref} \cdot \overline{I_{OFF}^{ref}} \cdot t_p^{ref}} \quad (7)$$

The off currents I_{OFF}^{target} and I_{OFF}^{ref} depend on the threshold voltage V_t which is a function of V_{bb}. However, the off currents can be approximated by the number of gates g_{target} and g_{ref} and the average cell leakage currents $\overline{I_{OFF}^{target}}$ and $\overline{I_{OFF}^{target}}$ in the two designs. With a sufficiently representative reference design, with reasonable constraints and a comparable cell distribution to the target design we can assume $\overline{I_{OFF}^{target}} \approx \overline{I_{OFF}^{ref}}$ to obtain

$$r_{leak} = \frac{g_{target} \cdot t_p^{target}}{g_{ref} \cdot t_p^{ref}} \quad (8)$$

While the *minimum* cycle times depends on V_{DD} and V_{bb}, we note that we are only interested in optimizing for a given cycle time t_p^{target} for the target design. Further, the cycle time t_p^{ref} for the reference design at the characterization points is anyway known. The resulting independence of (8) and (7) from V_{DD} and V_{bb} is the foundation for the subsequent target-design specific optimization.

B. Finding the MEP for a constant frequency

To scale the reference design dynamic and leakage characteristics to the target design, the latter is first synthesized at any valid V_{DD} and V_{bb} combination for the desired target

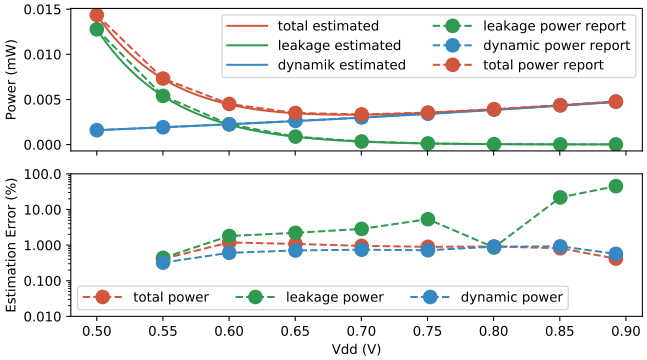

Fig. 3: 32 Bit RISC processor: Constant frequency power estimation based on the 32 bit multiplier reference design vs. power report of the design at given point (circle).

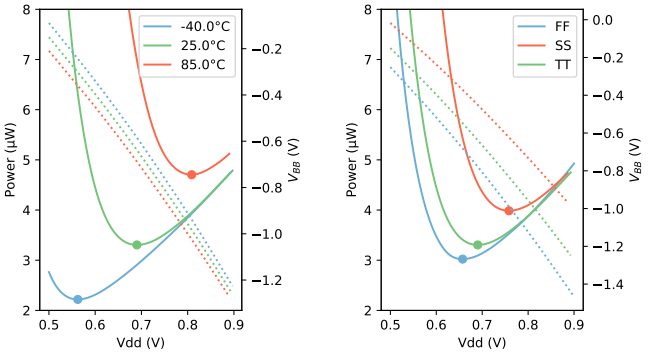

(a) Across temperature, TT corner (b) Across process, 25 °C

Fig. 4: Power estimation for the 32 bit RISC processor.

frequency. At this point we then extract its dynamic and the leakage energy and compute the corresponding scaling factors r_{dyn} and r_{leak} relative to the reference design at the same supply and bias voltages. With r_{dyn} and r_{leak} only depending on design specific properties, they can be used to directly scale the dynamic and leakage energy from the reference design for any other point in the iso-frequency design space.

$$E_{target}(V_{dd}, V_{bias}) \approx r_{leak} \cdot E^{ref}_{leak}(V_{dd}, V_{bias}) + r_{dyn} \cdot E^{ref}_{dyn}(V_{dd}, V_{bias}) \quad (9)$$

A straightforward exhaustive search based on (interpolated) characterization data from the reference design yields the MEP $\min E_{target}(V_{DD}, V_{bb})$ and the corresponding voltages.

V. Case Study: 32 Bit Microprocessor

Our target design is a 32 bit RISC microprocessor, synthesized for a relaxed clock of 8 MHz in order to allow for a sufficiently large design space.

A. Modelling accuracy

First, we evaluate the accuracy of the scaling approach. The microprocessor is implemented at 0.5 V with forward bias to achieve the target frequency. The leakage and dynamic power are extracted and the scaling factors are derived. At constant frequency, supply/bias tradeoff power curve is then sampled in 50 mV steps at which the full library is charcterized in order to allow for power reports. This curve is shown in Fig. 3. The circles denote the energy derived from the power reports and the contineous solid lines denote the estimate based on the scaled ASVBB model. The error in terms of total energy is on the order of 1% over the voltage range (for a constant frequency). Only the error on leakage increases slightly as the design shifts towards higher supply voltages, where, however, the leakage impact becomes neglectable and dynamic energy dominates anyway.

B. Process and Temperature effects

With the verification of the modelling accuracy we can now use this model to analyze the designs tradeoffs across PT. Figure 4a shows the effect of temperature on the MEP in the typical corner: we observe a range of the optimum supply voltage from around 560 mV to 690 mV and 810 mV for −40 °C, 25 °C, and 85 °C, respectively due to the leakage penalty of the forward bias needed to achieve the operating frequency at low supply voltages. A similar picture, however less severe, appears when analyzing the range across process corners as shown for 25 °C in Fig. 4b.

VI. Conclusion

The design space mapping methodology presented in this paper allows a designer to quickly evaluate the tradeoffs of the body bias voltage against the supply voltage for constant frequency operation. We show how to apply precharacterized leakage and dynamic energy maps from a reference design with a pruned library to a larger design with a different critical path length and a different distribution between dynamic energy and leakage power. The accuracy of this model is on the order of 1% of the energy report from the signoff tool. Further, we show how the minimum energy point can shift over a wide range when taking temperature and process variation into account. Joint adjustment of both body bias and supply voltage is key to achieve the best results.

References

[1] H. Patel *et al.*, "A 55nm Ultra Low Leakage Deeply Depleted Channel Technology Optimized for Energy Minimization in Subthreshold SRAM and Logic ," in *Proceedings of ESSCIRC 2016*, Apr. 2016.

[2] S. Jain *et al.*, "Design-oriented energy models for wide voltage scaling down to the minimum energy point," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 64, Dec 2017.

[3] B. Calhoun *et al.*, "Characterizing and modeling minimum energy operation for subthreshold circuits," in *Proceedings of ISLPED 2004*, Aug. 2004.

[4] O. Andersson *et al.*, "Impact of switching activity on the energy minimum voltage for 65 nm sub-VT CMOS," in *Proceedings of NORCHIP, 2011*, Nov. 2011.

[5] M. Miyazaki *et al.*, "A 175 mv multiply-accumulate unit using an adaptive supply voltage and body bias (asb) architecture," in *Proceedings of ISSCC, 2002*, Feb 2002.

[6] J.Lee *et al.*, "A 6.4pj/Cycle Self-Tuning Cortex-M0 IoT Processor Based on Leakage-Ratio Measurement for Energy- Optimal Operation Across Wide-Range PVT Variation," in *Proceedings of ISSCC, 2019*, Feb 2019.

[7] M. Pons *et al.*, "A 0.5 V 2.5 µW/MHz Microcontroller with Analog-Assisted Adaptive Body Bias PVT Compensation with 3.13 nW/kB SRAM Retention in 55 nm Deeply-Depleted Channel CMOS," in *Proceedings of CICC, 2019*, Apr 2019.