

Variations-Aware Low-Power Design with Voltage Scaling

Navid Azizi[‡], Muhammad M. Khellah[†], Vivek De[†], Farid N. Najm[‡]

[‡]Department of ECE, University of Toronto, Toronto, Ontario, Canada

[†]Circuits Research, Intel Labs, Hillsboro, Oregon

{nazizi,najm}@eecg.utoronto.ca, {vivek.de,muhammad.m.khellah}@intel.com

ABSTRACT

We present a new methodology which takes into consideration the effect of Within-Die (WID) process variations on a low-voltage parallel system. We show that in the presence of process variations one should use a higher supply voltage than would otherwise be predicted to minimize the power consumption of a parallel systems. Previous analyses, which ignored WID process variations, provide a lower non-optimal supply voltage which can underestimate the energy/operation by 8.2X. We also present a novel technique to limit the effect of temperature variations in a parallel system. As temperatures increases, the scheme reduces the power increase by 43% allowing the system to remain at it's optimal supply voltage across different temperatures.

Categories and Subject Descriptors B.7.1 [Integrated Circuits]: Types and Design Styles;

General Terms: Design

Keywords: Process Variations, Parallel Systems, Low-Voltage.

1. INTRODUCTION

Power consumption is increasingly becoming the bottleneck in microprocessor design. The core of the microprocessor, which includes the datapath of the processor, has the largest power density on the microprocessor [1]. In an effort to reduce the power consumption of the datapath, the supply voltage can be reduced leading to a reduction of dynamic and static power consumption. Lowering the supply voltage, however, also reduces the performance of the circuit which is usually unacceptable. One technique, available in certain application domains, to overcome this limitation is to replicate the circuit for which the supply voltage is being reduced in order to obtain the same throughput regardless of the value of the supply voltage [2]. It can be shown that even with circuit replication and overhead there are large power benefits that can be obtained by going to a lower supply voltage [2].

While some previous studies have parallelized systems and considered die-to-die variations [3][4], these studies have not

taken into consideration the effect of WID variations during low-voltage operation of a parallel system.

As a result of technology scaling, modern integrated circuits exhibit an increased sensitivity to local variations and thus understanding the effect of WID variations becomes important in the design of high performance systems [5][6]. Local variations can affect the critical path delay significantly [5]; given that a parallel system may have thousands of critical paths, local variations can have a large effect on total throughput and consequently power.

In this paper we present a new methodology for low-power design which takes into consideration the effect of WID process variations in the design of a parallel system. It will be shown that the number of blocks needed in parallel at low voltages increases considerably, when WID process variations are considered, and consequently the optimal supply voltage that provides the lowest power at the same throughput and yield as that of the original system is higher than if not considering WID process variations. We further show how correlations affect the design and the optimal choice for the supply voltage.

We also show that changes in temperature can have a large effect on the power of parallel systems, and on the choice of supply voltage. Previous designs have used body bias to adjust for temperature variations [4]. These designs need a triple well process which may not always be available. We present a novel technique, the Temperature Dependent Deactivation Scheme (TDDS), to limit the variations in power consumption due to temperature fluctuations, allowing a lower supply voltage and lower system power.

The paper is organized as follows: In section 2, we present some background. In section 3 we present our new methodology which takes WID process variations into consideration when designing a parallel system. We then present our technique to limit the effect of temperature variations in section 4. In section 5 we present our results. Finally, we conclude the paper in section 6.

2. BACKGROUND AND MODELS

A well-known technique for low-power design, proposed by Chandrakasan and Brodersen [2], is to replicate a logic block a number of times (i.e., to use several *instances* of the same block) and to allow all instances to work in parallel at reduced supply voltage and frequency, with the aid of a demultiplexor and a multiplexor, as shown in Fig. 1. If the application domain allows this type of fine-grained parallelism, such as in DSP applications, then this allows one to maintain the same throughput (operations completed per unit time), at reduced power dissipation.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

DAC 2005, June 13–17, 2005, Anaheim, California, USA.

Copyright 2005 ACM 1-59593-058-2/05/0006 ...\$5.00.




Figure 1: Transformation into a LV Parallel System

The motivation for this transformation is that the dynamic (switching) power is given by $P = CV^2f$. In the transformed circuit, it can be shown [2] that the power is given by $P = ((1 + o/m)CV^2f)/b^2$ where m is the number of blocks in parallel, o is the overhead involved in the process and b is the amount by which the supply voltage can be reduced. If the overhead is small, then the dynamic power is reduced by b^2 .

We will refer to the original block, operating at the higher voltage, as the High-Voltage (HV) block, and to each of the blocks operating at the lower voltage, in the parallel system, as a Low-Voltage (LV) block. The number of blocks required, m , is found [2] by dividing the delay through the LV blocks, which we denote by T_{max} , by the delay of the HV block, denoted T_o :

$$m = \lceil T_{max}/T_o \rceil \quad (1)$$

This prior work considered only the dynamic (switching) power. In this work, we extend the analysis to take into account leakage current as well as statistical variations in leakage, resulting from underlying process variations. Notably, we take into account Within-Die (WID) variations. This will lead to new insights for how the choice of the reduced supply voltage should be made, and will give a methodology for how the number of blocks should be determined.

2.1 Technology

All simulation results reported in this paper are based on HSPICE, using Berkeley Predictive Technology Models (BPTM) [7] for a 70nm technology. The transistor models were expanded to include gate tunneling leakage which was modeled using a combination of four voltage-controlled current-sources, as in [8]. The resulting transistor macro-model was fitted to industrial data found in [9].

2.2 Generic Block

To determine the effect of process variations at different supply voltages requires one to estimate the statistics of the block leakage power, as well as the statistics of the total block delay. Both these sub-problems are research topics in their own right and have been the subject of various papers. Lacking a complete and universally acceptable solution to these problems, especially the timing problem, we have opted to use plain Monte-Carlo (MC) analysis to estimate both delay and power distributions, for purposes of this paper. This is not the most efficient approach, but does give us some confidence in the resulting distributions, which we need in order to demonstrate the main results of our work related to the dependence of the supply voltage setting and the block count on the underlying variations.

In order to make the MC somewhat less expensive, we have used a *generic block* as the test vehicle throughout the paper, which is meant to be representative of typical logic

blocks, whose timing is normally determined by a number of roughly equal-delay critical timing paths. Specifically, we use a generic block consisting of 1000 inverter chains, of which 100 are assumed critical. The 100 critical chains determine the block's maximum delay, while all 1000 inverter chains determine its power consumption. This allows the MC to be more efficient, and provides a means to easily vary the number of presumed critical paths, and examine the effect of that on our results.

Each inverter chain within the block is used to represent the characteristics (delay, power) of a path through a typical combinational circuit. While this may appear as a simplification, simulations on chains of NAND and NOR gates provided similar changes to inverter chains, in terms of delay and power, as the supply voltage was lowered. Thus, using an inverter chain to characterize the effect of lower supply voltages is warranted. It is, furthermore, not new to use an inverter chain to model low-voltage operation, as can be seen in [3]. We specifically use an inverter chain of length 14 and fan-out of 3, both of which are typical of modern circuits.

3. BLOCK COUNT

As a consequence of considering leakage power and WID variations, we will see below (section 5) that the power consumption at very low supply voltages increases considerably. This is an important effect that has implications for the number of blocks and the supply voltage chosen, and the reason for this behavior is as follows. Due to process variations, the maximum delay through a circuit becomes a random variable, with some distribution. While some blocks in the parallel system may be fast (i.e., they are not the delay bottleneck), other blocks may be slower. However, because all blocks operate with the same clock period, the fast blocks would spend some fraction of the cycle in idle mode, during which they dissipate only leakage power. Since the faster blocks are usually the more leaky ones, then the total leakage power of the parallel system starts to increase for larger block count (i.e., for lower supply voltages).

3.1 Independent paths

An important issue to be considered is whether the random variables representing path delays of two disjoint paths are independent or not. It simplifies the analysis to assume independence, but path delays may be correlated on silicon. In the absence of detailed information on this, which is typically the case in practice, and assuming that path delay correlations are non-negative, which also is reasonable in practice (a physical variation that slows down one path is unlikely to speed up another), the conservative approach is to assume that the path delays are independent. The reason for this is as follows.

Let \mathbf{X} and \mathbf{Y} be multi-variate normal random vectors, $\mathbf{X} \sim \mathcal{N}(\mu, \Sigma = (\sigma_{ij}))$ and $\mathbf{Y} \sim \mathcal{N}(\mu, \Gamma = (\gamma_{ij}))$. Thus, both vectors have the same mean vector, while σ_{ij} represents the covariance in \mathbf{X} and γ_{ij} represents the covariance in \mathbf{Y} . If $\sigma_{ij} \geq \gamma_{ij}$ for all $i \neq j$ (i.e., if the variables in \mathbf{X} are more correlated than the variables in \mathbf{Y}), then it was proved in [10] that the following relation holds:

$$P\{\mathbf{X} \leq \mathbf{a}\} \geq P\{\mathbf{Y} \leq \mathbf{a}\} \quad (2)$$

for any real vector \mathbf{a} . If \mathbf{Z} is obtained from \mathbf{X} by retaining the individual (marginal) distributions of the vector entries

(i.e., same means and variances) while setting all covariances to zero (i.e., all vector components become independent), then as long as the covariances in \mathbf{X} are non-negative, we have:

$$P\{\mathbf{X} \leq \mathbf{a}\} \geq P\{\mathbf{Z} \leq \mathbf{a}\} \quad (3)$$

If all the \mathbf{a}_i 's are set to one value a , then the above leads to:

$$P\left\{\max_i(X_i) \leq a\right\} \geq P\left\{\max_i(Z_i) \leq a\right\} \quad (4)$$

In other words, if the random variables in the above analysis are the path delays, and a is some time interval, then the independence assumption leads to the worst-case timing yield, hence a conservative analysis.

3.2 Determining the Block Count

Let us first consider the case when a conservative approach is desirable, based on an independence assumption among disjoint paths. In this case, the random delays of the various blocks are also independent. Given a desired percentage timing yield, Y , the required setting for the allowable maximum delay through the system, T_{max} , i.e., the maximum delay among the m blocks, can be easily determined if the distribution of the delay of a stand-alone block is known. Typically, this can be found using some form of Statistical Static Timing Analysis (SSTA), if available, or by Monte-Carlo sampling, as we will use on our generic block. If the block delay distribution is $F_B(t | V_{dd})$, a notation that emphasizes the fact that the distribution depends on the supply voltage setting, then T_{max} may be determined from:

$$F_B(T_{max} | V_{dd})^{\lceil \frac{T_{max}}{T_o} \rceil} = Y \quad (5)$$

which is obtained using basic probability theory, knowing that $m = \lceil \frac{T_{max}}{T_o} \rceil$. This equation can be solved, using any method for solving non-linear equations, to find T_{max} and the block count m .

For our generic block, with a known number, say 100, of critical paths in each block, and given the distribution of delay for an inverter chain as $F_I(t | V_{dd})$ (see section 3.3 below), and once again assuming independence among paths, we get $F_B(t | V_{dd}) = F_I(t | V_{dd})^{100}$, resulting in:

$$F_I(T_{max})^{100 \lceil \frac{T_{max}}{T_o} \rceil} = Y \quad (6)$$

We have solved this for T_{max} using Fixed Point Iteration, and then m was easily computed using (1). In this study a required yield of 99.7% is used, but different yields were also tested, with little change in the relative power savings.

Finally, if the path delays are not assumed independent, then the above procedure can no longer be applied. There is no simple closed-form solution in this case. Instead, if the correlations among paths are known, then SSTA can be applied on the parallel system until an acceptable T_{max} (and m) are found. In our case, we used MC analysis on the parallel system, based on a total of $100m$ critical paths, and given some distribution of each path delay, $F_I(t | V_{dd})$ (see below), to determine T_{max} and m . As for the path-to-path correlations, for our MC analysis, we used a distance based correlation function with a quadratically decaying correlation with distance. The distance metric used was the degree of separation between paths: the list of paths was ordered arbitrarily, and paths that are nearby on the list are deemed to be near, otherwise far.

3.3 Path delay

For our generic block, the above solutions require the distribution of delay of a single inverter chain. This was determined by performing MC sampling on the threshold voltages, V_t , of the transistors. Again, a more exact analysis would require other parameters be varied as well, which can be done, but focusing on V_t is enough to make the points we want to demonstrate in this paper in connection with the generic block. As part of the same MC analysis, we also compute the distributions of the leakage power and the switching power. This was performed at different supply voltages, temperatures, and transistor widths.

As for the issue of within-path correlation, since the delay through a path is the sum of random variables, it can be easily shown that, by changing the correlation assumptions, the mean of the path delay distribution does not change, but the variance of that distribution increases as positive correlations between the individual random variables becomes stronger. Thus, by assuming correlations between V_t variations within a path, the probability of a larger delay increases, the timing yield decreases, so that the conservative case is when within-path delays are strongly correlated. This is in contrast with the path-to-path case.

For our generic block, the MC analysis on the single inverter chain was performed in two ways: first it was assumed that all V_t variations in the inverter chain were independent, and then it was assumed that there was some distance based correlation between the V_t variations of each transistor in the inverter chain. Distance was measured based on degree of separation on the path.

3.4 Summary

In summary, this section: (1) describes the procedure for determining the number of blocks that are needed when process variations are considered, and (2) explains the effect of both within-path and path-to-path correlations on the timing yield. The “worst case” timing yield corresponds to where there are strong correlations within a path (Sec. 3.1), but total independence path-to-path (Sec. 3.3). Conversely, the “best case” timing yield is the reverse: strong correlation path-to-path and total independence within-path.

4. OPERATING CONDITION VARIATIONS

In addition to process variations, changes in operating conditions (typically, temperature and supply voltage) can also affect the number of blocks needed and the best voltage to run the parallel system at. Since the number of blocks in the LV system is set during the design process, the number chosen must be such that even under the worst-case operating conditions, the throughput of the LV system, at a minimum, is equivalent to that of the original HV system. We will focus on temperature variations, but the approach can be easily extended to account for supply voltage variations; it can also be extended to Die-to-Die (D2D) variations.

At traditional supply and threshold voltages, circuits operate slowest at high temperatures, but this is no longer true when supply voltages near the threshold voltage. In the process we used, at supply voltages below 0.7V, circuits operate slowest at low temperatures, because a large portion of the current through the switching cycle is provided by subthreshold leakage current which is smaller at low temperatures. Thus, the number of blocks needed and the supply voltage setting must be chosen at low temperature.

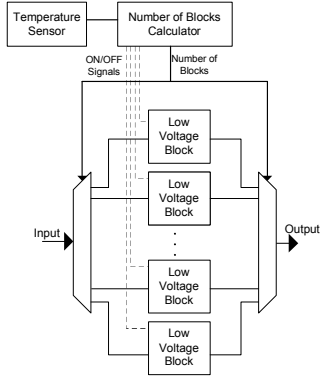


Figure 2: Organization of Temperature Dependent Deactivation

However, if we consider what happens at higher temperatures which invariably will be encountered during circuit operation, two facts are obvious: 1) the number of blocks would often be larger than needed, and 2) the circuit would be dissipating more power than needed. The first of these is obvious from the preceding discussion. As to the second, it can be explained as follows. Leakage current increases at higher temperatures, because subthreshold leakage current doubles every 7°C while gate oxide leakage is not too sensitive to temperature. Couple this with the larger-than-needed number of blocks at high-temperature, and it is clear that the power dissipation of such a system becomes an issue at higher temperatures. Thus, much of the energy benefit obtained by using a lower supply voltage and parallelism may be lost at high temperatures.

To address this problem, we propose to disable some of the blocks as temperature increases. This leads to power savings in the form of a leakage reduction, which can offset most of the increase in power as the temperature is increased. A possible implementation of this scheme is shown in Fig. 2. A temperature sensor detects the temperature that the circuit is operating at and reports it to another circuit (the “Number of Blocks Calculator”). This circuit, either through a look-up table or other means determines how many blocks have to be ON in order to obtain the required throughput. That information is fed to the multiplexor and demultiplexor, and to the blocks themselves, turning some of them ON/OFF or putting some of them into sleep mode.

This Temperature Dependent Deactivation Scheme (TDDS) allows a large energy reduction in LV systems regardless of the temperature of operation. Without it, the supply voltage of a LV system would have to be set at a higher value, where the temperature would not have a large effect on its operation, leading to lower energy savings compared to a HV system.

5. RESULTS

5.1 Low-Voltage Trends

In order to gain some insight into the effect of lower voltages on the power dissipation in a parallel system, we will first consider our generic block without considering the presence of process and other variations. For every supply voltage value, we can go through the traditional transformation shown in Fig. 1, by first finding T_o for the given voltage by simulation, then computing the required number of blocks as $m = \lceil T_{max}/T_o \rceil$, maintaining the same throughput at the different voltage settings. The results are shown in Fig. 3,




Figure 3: Effect of Supply Voltage on Dynamic Power and Leakage Power




Figure 4: Effect of Lowering the Supply Voltage on the Total Power and the Number of Blocks

where the different points at the same voltage value correspond to designs with different transistor widths in the inverter chain.

While the dynamic power consumption has almost a constant exponential drop with supply voltage, the leakage power is seen to exhibit a more interesting behavior. Initially, as the supply voltage is decreased, the total leakage power decreases, as the reduction in gate and subthreshold leakage per block outweighs the increase in leakage due to the larger number of blocks, but as more and more blocks are needed at very low voltages, the total leakage power starts to increase.

The total power of the parallel system can be computed from the dynamic and static power based on some assumed switching activity factor, α . The overhead involved in parallelizing the system, which must also be considered, consists of three components: the extra routing capacitance due to the broadcast of the input to the parallel blocks, the output routing in the multiplexor, and the multiplexor overhead and control [2]. The results of this analysis are shown in Fig. 4, based on $\alpha = 0.1$. It is found that the best operating point is at 0.3V, with $m = 18$ blocks in parallel, providing a 10.3X reduction in the power consumption of the system relative to the original HV system. The overhead is about 25% of the total power at this operating point. An important point to keep in mind is that the different points on the curves, corresponding to different supply voltages, correspond to different block counts, but the same throughput (operations completed per unit time). With regard to the chosen value of α , we will consider below the effect of variations in α , but it should be said that the observation in this section remains true: there is an optimal design point at a specific supply voltage, and the power savings can be large.

5.2 Process Variations

Suppose a design transformation as in Fig. 1 was carried out and implemented on Silicon *without* considering process variations. What then is the impact of process variations,

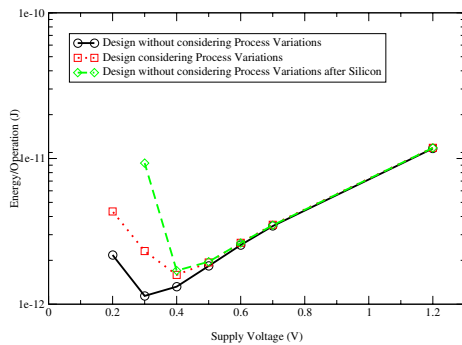


Figure 5: Effect of WID process variations on Energy/Operation

which are inevitable, on the performance of that chip? Although we did not actually measure any data on real hardware, we try to illustrate what the answer would be in Fig. 5, which assumes independence between path delays and independence between V_{IS} at $30^\circ C$ with $\alpha = 0.1$. The figure shows three curves. The bottom (solid) curve shows the *expected* performance of that design, without considering process variations based on an analysis such as in section 5.1. Recall that each point on this curve corresponds to a different block count. Now, if for each of these points, with that specific block count, we consider what happens after process variations are taken into account, we get the top (dashed) curve in the figure, marked “after Silicon.” There is a *significant* increase in the energy/operation at low voltages. When not considering process variations, the supply voltage that minimizes the energy/operation during the design phase is 0.3V, but if that design is implemented on Silicon, the energy/operation would be 8.2X times higher than expected. The resulting reduction in energy/operation compared to the original HV system is minor and not worth the trouble. In contrast, if process variations are taken into account up-front, and the block count chosen accordingly as proposed in section 3, one obtains the results shown in the middle curve in that figure. The energy/operation is *much* improved at lower voltages, showing conclusively that process variations must be taken into account, as we have described. As a result of our analysis, one sets the supply voltage to 0.4V, leading to an energy/operation of the system that is 7.4X lower than the original HV system.

5.3 Effect of Correlation

Fig. 6 shows the number of blocks that are needed under different correlation assumptions to maintain throughput. As explained in Section 3.4, the “worst case” plot corresponds to where there are strong correlations within a path, but total independence path-to-path. The “best case” is the reverse: strong correlation path-to-path and total independence within-path. It can be seen that when no variations were assumed, only 18 blocks were needed at 0.3V, but when variations are included in the analysis, the number of blocks at 0.3V needed ranges from 41 blocks to 100 blocks. Thus, the traditional approach, which did not take into consideration the effect of WID variations, considerably underestimates the number of blocks needed to obtain the required throughput and yield.

Fig. 7 shows the power consumption of the parallel system assuming different types of correlations, at $30^\circ C$ with $\alpha = 0.1$. As before, each point on the plot represents a possibly different number of blocks in parallel, as determined




Figure 6: The Effect of Process Variations on the Number of Blocks

by the procedure described in section 3.2. The first thing to observe is that irrespective of the assumed correlation structure, the curves when considering variations are all higher than when no variations are considered. This is not to say that it is better to ignore variations, because, as we saw in section 5.2, the power dissipation on Silicon would be much higher, due to the unavoidable presence of variations in practice. Thus, the curve for the “no variations” case is given only for reference and comparison and does not represent a design which is actually realizable.

The case of “no variations” would suggest that a supply voltage of 0.3V is optimal. However, with variations considered up-front, the best case curve gives an optimal V_{dd} of 0.4V, and the worst case gives a $V_{dd} = 0.5V$ (leading to a power reduction of 7.6X and 4X, respectively, compared to the original system). The number of blocks at the optimal supply voltage at the different correlation assumptions ranges from six to eleven.

Since the variations and correlations are usually not known early in the design process [6], it becomes interesting to consider the impact of designing with one set of assumptions. If, for example, the “best case” assumptions were used and then found to be incorrect, the throughput or yield of the system would be lower than anticipated. Conversely, if “worst case” assumptions were used and then found to be incorrect the supply voltage could have been lowered further in the design phase to further reduce the power. Thus, if the primary requirement of a design is performance, and power is of secondary concern, the conservative assumption would be to use the “worst case” assumptions. If, however, power is the primary concern, and performance secondary, the conservative assumption would be to use the “best case” assumptions.

5.4 Effect of Activity Factor

As α is varied, the number of blocks needed to obtain the required throughput and yield at different supply voltages does not change, because m is not a function of α . But α does affect the power consumption and, therefore, the supply voltage that minimizes the power consumption.

The solid line in Fig. 8 shows the optimal supply voltage at different α 's. Observe that as α tends toward 1, the optimal supply voltage is reduced, because the static power becomes less important to the total power consumption and thus the increased parallelism at low voltages is not a concern. As α becomes very small, the optimal supply voltage becomes larger so as to reduce the parallelism and consequently the leakage.

Also on Fig. 8 is a comparison of the power reduction that

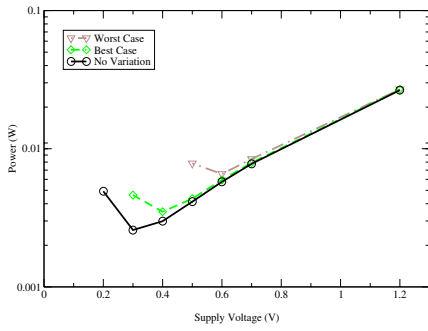


Figure 7: The effect of correlation on total power

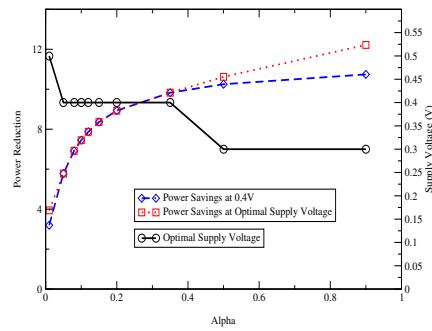


Figure 8: Best Voltage and Power Reduction at Different α 's




Figure 9: Effect of Temperature Variations on Power

is possible when using the optimal supply voltage to the power reduction when using a supply voltage of 0.4V (for our circuit, the supply voltage that maximizes the power reduction at an activity factor of 0.1). Since the two curves run along each other for most of their length, only differing slightly at their extremities, our previous conclusions about the optimal supply voltage hold true regardless of the activity factor.

5.5 Best Transistor Widths

We explored the effect of transistor widths on the number of blocks needed, and the power consumption of the system. It was found that the transistor width that minimizes the power consumption changes as the supply voltage is lowered in the presence of WID variations. At high supply voltages, small transistors are preferable, as increases in performance by using wider transistors are offset by increased power consumption. But at lower supply voltages, where process variations have a large effect, using wider transistor increases performance *and* decreases the variation in V_t . The decreased variation results in a lower number of blocks needed to obtain the same throughput, thus more than offsetting the extra leakage incurred by using wider transistors. This observation is opposite to what is seen in [11] where WID process variations are not considered and minimum width transistors are optimal at low voltages. In the results presented thus far we have used the transistor width that minimizes the power at each data point.

5.6 Operating Condition Variations

Fig. 9 shows the power consumption of the parallel system at different temperatures. At high temperatures the power consumption increases, partly due to an increase of subthreshold leakage. Observe that at low voltage, where there are many blocks in parallel, there is a large increase in power as there is a considerable increase in the leakage due to the parallelism. At 110°C, the power consumption of the parallel system at 0.3V is larger than that of the original system.

When the TDDS of Section 4 is used, the increase in the power consumption is limited as blocks that are unused are turned off at high temperatures. At 0.4V, when using TDDS, there is a power increase of 1.9X when the temperature changes from 30°C to 110°C instead of an increase of 3.3X when not using TDDS.

6. CONCLUSION

Power consumption is increasingly becoming the barrier in submicron integrated circuit design. A LV parallel system

is one possible option to reduce the power consumption of the datapath of microprocessors.

Ignoring WID variations, however, during the design process can lead to silicon which has an energy/operation many times that what was expected. We presented a new methodology that takes WID variations into consideration when designing a parallel system and showed that the supply voltage that minimized power consumption at the required throughput and yield was higher than when not considering WID variations. Even in the presence of WID variations, power can be reduced by up to 7.6X.

We also showed that parallel systems have large increases in power consumption when the temperature increases thus reducing their benefit. We introduced a novel scheme, the TDDS, which allows parallel systems to be used across a wide range of temperatures. As temperatures increased, our scheme reduced the power increase by 43% allowing the system to remain at its optimal supply voltage across different temperatures.

7. ACKNOWLEDGMENTS

This work was done in part at Intel Corporation. We would like to thank Khaled Heloue for providing assistance.

8. REFERENCES

- [1] J. Schütz and C. Webb. A scalable X86 CPU design for 90nm process. *ISSCC*, 2004.
- [2] A. P. Chandrakasan and R. W. Brodersen. *Low Power Digital CMOS Design*. Kluwer Academic Publishers, 1995.
- [3] D. Liu and C. Svensson. Trading speed for low power by choice of supply and threshold voltages. *IEEE Journal of Solid-State Circuits*, 28(1):10–17, January 1993.
- [4] C. Kim, H. Soeleman, and K. Roy. Ultra-low-power DLMS adaptive filter for hearing aid applications. *IEEE Transactions on VLSI*, 11(6):1058–1067, December 2003.
- [5] M. Eisele, et al. The impact of intra-die device parameter variations on path delays and on the design for yield of low voltage digital circuits. *IEEE Transactions on VLSI*, 5(4):360–368, December 1997.
- [6] D. Boning and S. Nassif. Models of process variations in device and interconnect. In A. Chandrakasan, W. J. Bowhill, and F. Fox, editors, *Design of High-Performance Microprocessor Circuits*. IEEE Press, New York, NY, 2001.
- [7] <http://www-device.eecs.berkeley.edu/~ptm/>.
- [8] D. Lee, W. Kwong, D. Blaauw, and D. Sylvester. Simultaneous subthreshold and gate-oxide tunneling leakage current analysis in nanometer CMOS design. *ISQED*, pages 287–292, 2003.
- [9] W.K. Henson et al. Analysis of leakage currents and impact on off-state power consumption for CMOS technology in the 100-nm regime. *IEEE Transactions on Electron Devices*, 47(2):440–447, February 2000.
- [10] Y.L. Tong. *The Multivariate Normal Distribution*. Springer-Verlag, 1990.
- [11] B.H. Calhoun, A. Wang, and A. Chandrakasan. Device sizing for minimum energy operation in subthreshold circuits. *Custom Integrated Circuits Conference*, 2003.