

Variations-Aware Low-Power Design and Block Clustering With Voltage Scaling

Navid Azizi, *Student Member, IEEE*, Muhammad M. Khellah, *Member, IEEE*, Vivek K. De, *Senior Member, IEEE*, and Farid N. Najm, *Fellow, IEEE*

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 nonoptimal supply voltage which can underestimate the energy/operation by 8.2 \times . 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. To further limit the effect of variations, and allow for a reduced power consumption, we analyzed the effects of clustering. It was shown that providing different voltages to each cluster can provide a further 10% reduction in energy/operation to a low-voltage parallel system, and that the savings by clustering increase as technology scales.

Index Terms—Low-voltage, parallel systems, process variations.

I. INTRODUCTION

POWER consumption has become a bottleneck in microprocessor design. The core of a microprocessor, which includes the datapath, 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 way to overcome this limitation, available in some application domains, is to replicate the circuit block whose supply voltage is being reduced in order to maintain the same throughput [2], leading to what we will refer to as a “parallel system,” implementation of a logic block. It has been shown that in spite of the circuit replication this leads to large power benefits [2]. Previous studies have shown that the supply voltage can be reduced down to 0.13 V to obtain power reductions [3], after which the overhead to parallelize the system becomes larger than the energy savings obtained by lowering the supply voltage.

Manuscript received May 23, 2006; revised September 19, 2006, November 20, 2006, and January 18, 2007.

N. Azizi is with the Department of Electrical Engineering, University of Toronto, Markham, ON L3R 0C4 Canada (e-mail: nazizi@eecg.utoronto.ca).

M. M. Khellah is with the Circuits Research Laboratory, Intel Corporation, Hillsboro, OR 97124 USA (e-mail: muhammad.m.khellah@intel.com).

V. De is with the Corporate Technology Group, Intel Corporation, Hillsboro, OR 97124 USA (e-mail: vivek.de@intel.com).

F. N. Najm is with the Electrical and Computer Engineering Department, University of Toronto, Toronto, ON M5S 3G4 Canada (e-mail: f.najm@utoronto.ca).

Digital Object Identifier 10.1109/TVLSI.2007.899226

As a result of technology scaling, there are increased process variations of circuit parameters such as the transistor channel length and transistor threshold voltage [4]. The increased process variations can have a significant effect on circuit performance and power [5], variations also have an impact on how exactly a parallel system should be designed. While some studies have somewhat considered die-to-die variations [3], [6], these studies have not taken into consideration the effect of within-die (WID) variations during low-voltage operation of a parallel system. 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 [7], [8]. Local variations significantly can affect the critical path delay [7]. Given that a parallel system may have thousands of critical paths, local variations can thus have a large effect on its total throughput and power.

We present a new methodology for low-power design of parallel systems which takes into consideration the effect of WID process variations. As an expansion and extension of the work found in [9], this paper will show that the number of parallel blocks needed 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. A similar observation was seen for nonparallel subthreshold circuits where the optimal supply voltage was slightly higher when considering WID process variations [10]. We also show how correlations affect the design and the optimal choice of supply voltage.

We further show that changes in temperature can have a large effect on the power dissipation of parallel systems, and on the choice of supply voltage. Previous studies have used body bias to adjust for temperature variations [6]. 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.

To reduce power consumption even further, we use block clustering to limit the effect of the underlying variations on the performance of parallel systems, and use our methodology to determine the power savings. The method clusters parallel blocks, and then applies small differences in supply voltage to equalize the performance of each cluster. We consider different organizations of this scheme.

This paper is organized as follows. In Section II, we present some background. Section III presents the generic block that

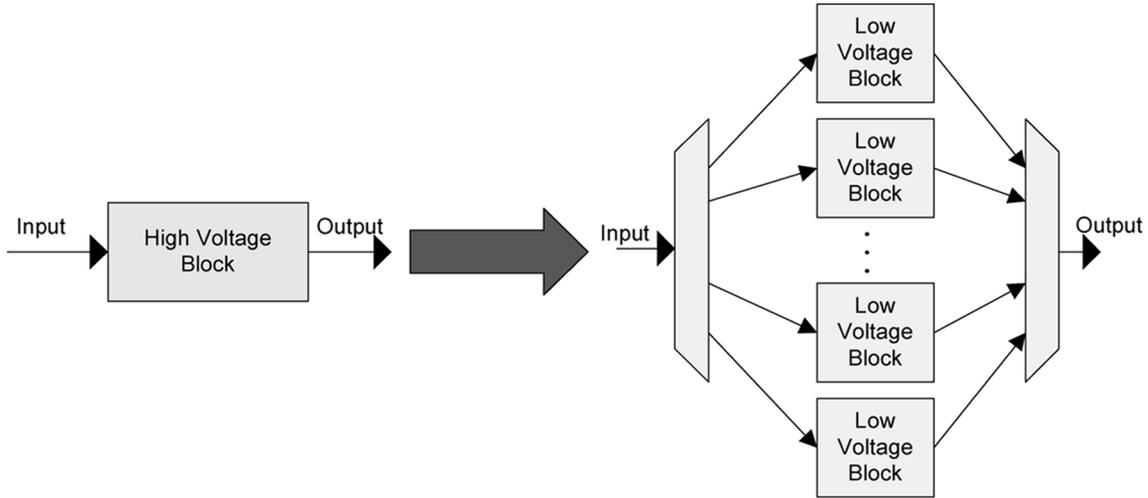


Fig. 1. Transformation into an LV parallel system.

we use as a test vehicle throughout this paper. In Section IV, we present our new methodology which takes WID process variations into consideration when designing a parallel system and results of applying this methodology are shown in Section V. We then present our technique to limit the effect of temperature variations in Section VI. Then, in Section VII, we present the clustering scheme that limits the effect of variations on a parallel system. Finally, we conclude the paper in Section VII.

II. BACKGROUND

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 digital signal processing (DSP) applications, then this allows one to maintain the same throughput (operations completed per unit time), at reduced power dissipation. We will refer to such an implementation of a logic block as the “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 required in the parallel system, 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 [9] and this paper, 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.

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). This is an important effect that has implications for the number of blocks and the supply voltage chosen.

III. GENERIC BLOCK

To determine the effect of process variations at different supply voltages, one needs to compute the statistics of the block leakage power, as well as the statistics of the total block delay. Both these subproblems 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 the Monte Carlo (MC) analysis to estimate both delay and power distributions, for purposes of this paper. This is not the most efficient approach, but it 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

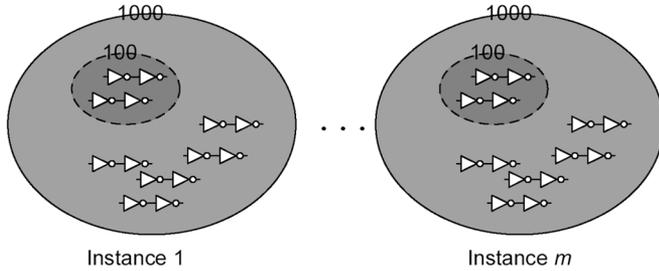


Fig. 2. Model for the parallel system.

block consisting of 1000 inverter chains, of which 100 are assumed critical, as shown in Fig. 2. 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 helps to represent the characteristics (delay, power) of a path through a typical combinational circuit. While this may appear as a simplification, simulations performed on chains of NAND and NOR gates where the supply voltage was lowered showed a similar relative increase/decrease in delay/power compared to the HV chain. Thus, using an inverter chain is warranted as it can serve to model the changes in characteristics of paths (not absolute values) as the supply voltage is lowered; using different gates in a path would not greatly affect the estimation of the required block count. Furthermore, there are previous studies that have shown that it is valid to use an inverter chain to model low-voltage operation, as can be seen in [3]. In this paper, for the critical paths, we use an inverter chain of length 14 and fan-out of 3, both of which are typical of modern circuits; for the noncritical inverter paths we use an inverter chain of length of 8 and fan-out of 3.

There are some limitations in using generic paths to determine the effects of variations on delay and power. For example, the number of critical paths in a block (in our case, we choose 100), can change the behavior of the circuit; a block with 100 critical paths that are fully correlated will show the same change in statistics as a block that has one critical path. However, a block that has 100 critical paths that are independent will exhibit different characteristics when the block count is varied. A further limitation in using a generic block is that logical dependencies between paths are not covered (for example, when two critical paths share some common subpath). However, given that we will explore a range of correlation assumptions for the delays between paths, the limitations of the generic block will not matter very much.

IV. BLOCK COUNT

To determine the number of blocks that are needed for a parallel system implementation, we follow a two-step approach, as shown in Fig. 3. First MC simulations are performed on a single generic path using HSPICE, based on a model of the variations and the correlations between the underlying variations on that

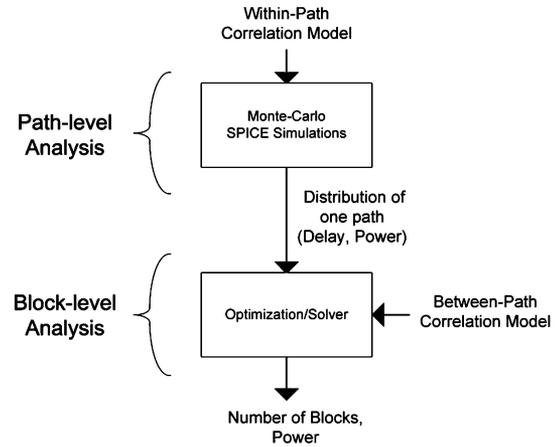


Fig. 3. Method of determining the number of blocks in parallel.

path, to find the distribution of its delay and power. These simulations are done at various voltages and temperatures, and the results are then stored and used in the second step of the process.

The second step of the process uses a fixed-point iterative algorithm to determine the number of blocks that are needed in parallel to maintain the throughput of the system as the voltage is lowered. The analysis is performed with different assumptions regarding the correlations between different paths in the system to determine the effects of different amounts of correlation on a parallel system.

The rest of this section describes the process in more detail in a top-down approach; first the fixed-point algorithm is described in Section IV-A assuming that the distribution of the single generic path delay is already known. Then, in Section IV-B, the method for obtaining the distributions of the generic path delay is discussed.

An important issue to be considered is whether the random variables representing the path delays of two disjoint paths and/or the random variables representing the underlying variations in transistors on a single path are independent or not. It simplifies the analysis to assume independence, but path delays and transistor variations may be correlated on silicon. In this paper, the effect of using different correlation assumptions, independence or correlation, on the number of blocks and consequently the power consumption of the parallel system will be explored.

In the absence of detailed information on the correlation on silicon, which is typically the case in practice, we will assume that path delay correlation and process variation correlations are nonnegative, which is a reasonable assumption in practice. For example, a physical variation that slows down one path is unlikely to speed up another. While this *nonnegativity assumption* is used at times to make the process more efficient, we also provide a slower method that can be used regardless of the correlation assumptions. Thus, our methodology is general and not really dependent on the nonnegativity assumption.

A. Block-Level Analysis

To determine the number of blocks that are needed in an LV parallel system to match the throughput of the HV system, a

fixed-point algorithm is used, which uses as input the distribution of delay of a single generic path.

The random variables representing the path delays of two or more disjoint paths in the LV parallel system can be either correlated or independent. Assuming independence between the path delays allows for an efficient analysis. This independence assumption, moreover, can be shown (if the nonnegativity assumption is used) to be the conservative assumption to use. On the other hand, if the correlations between the path delays are known, then a less efficient process has to be used.

In this section, the method for determining the number of blocks needed in parallel if the path delays are independent will be presented first. Then another method will be presented for the case of correlated paths. Finally, a proof for the statement that assuming independence between paths is the conservative approach will be given. While this proof depends on the nonnegativity assumption, the results of our methodology that will be presented in Section V do not depend on the proof, or the nonnegativity assumption.

1) *Independent Paths*: 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 MC sampling, as we will use on our generic block. If the block delay cumulative distribution function is $F_B(t | V_{dd})$, a notation which emphasizes the fact that the distribution depends on the supply voltage setting, then T_{\max} may be determined from

$$F_B(T_{\max} | V_{dd})^{\lceil \frac{T_{\max}}{T_o} \rceil} = Y \quad (2)$$

which follows from basic probability theory,¹ knowing that $m = \lceil T_{\max}/T_o \rceil$ [9]. This equation can be solved, using any method for solving nonlinear 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 IV-B), and once again assuming independence among paths, we get $F_B(t | V_{dd}) = F_I(t | V_{dd})^{100}$, resulting in

$$F_I(T_{\max})^{100 \lceil \frac{T_{\max}}{T_o} \rceil} = Y. \quad (3)$$

We solve this for T_{\max} using fixed point iteration, and then m is easily computed as $m = \lceil T_{\max}/T_o \rceil$ [9]. In this paper, a desired yield of 99.7%, which corresponds to the 3σ variation, is used.

2) *Correlated Paths*: If the path delays are not assumed independent, then the previous 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 or MC 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

¹Given a cdf $F(x)$, then the distribution of the maximum of n independent samples of $F(x)$ is $F(x)^n$.



Fig. 4. Placement of blocks in the parallel system.

system, based on a total of 100-m critical paths, and given some distribution of each path delay, $F_I(t | V_{dd})$ (see the following), to determine T_{\max} and m . For the path-to-path correlations, we used a distance-based correlation function with a quadratically decaying correlation with distance; the correlation function was obtained from industry sources. For paths within a block 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 were deemed to be near, otherwise far. For paths in different blocks, the blocks were first placed in a square fashion as shown in Fig. 4, and the distance between the paths were measured. The block width was set to the size of a specific functional unit that was a candidate to be parallelized which was $168 \mu\text{m}$ in size in a 70-nm technology.

3) *Proof That the Conservative Approach Is to Assume Independence Between Paths*: Let \mathbf{X} and \mathbf{Y} be multivariate 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 μ , 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 proven in [11] that the following relation holds:

$$P \left[\bigcap_{i=1}^n \{\mathbf{X}_i \leq \mathbf{a}_i\} \right] \geq P \left[\bigcap_{i=1}^n \{\mathbf{Y}_i \leq \mathbf{a}_i\} \right] \quad (4)$$

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 nonnegative, we have

$$P\{\mathbf{X} \leq \mathbf{a}\} \geq P\{\mathbf{Z} \leq \mathbf{a}\}. \quad (5)$$

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

$$P\{\max_i(X_i) \leq a\} \geq P\{\max_i(Z_i) \leq a\}. \quad (6)$$

In other words, if the random variables in the previous analysis are the path delays, and a is some time interval, then the independence assumption leads to the minimum timing yield, hence a conservative analysis [9].

B. Path Delay

For our generic block, the previous 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 comprehensive 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.

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 variations of each transistor in the inverter chain [12]. We estimate the distance between transistors in the path by using the amount of inverters between each pair of transistors, and then use the estimated distance to obtain a correlation between the transistors.

Assuming that the underlying variations are correlated is the conservative approach in regards to the timing yield. A proof, which uses the nonnegativity assumption, is shown as follows. While this proof depends on the nonnegativity assumption, the results of our methodology that will be presented in Section V do not depend on the proof or the nonnegativity assumption.

1) *Proof That the Conservative Approach Is to Assume Correlation Within a Path:* Let $\mathbf{W} = \mathbf{X}_1 + \mathbf{X}_2 + \dots + \mathbf{X}_n$, then the second moment of \mathbf{W} is $E((\mathbf{X}_1 + \mathbf{X}_2 + \dots + \mathbf{X}_n)^2)$ which can be expanded to

$$E[\mathbf{W}^2] = E[\mathbf{X}_1^2] + \dots + E[\mathbf{X}_n^2] + 2E[\mathbf{X}_1\mathbf{X}_2] + E[2\mathbf{X}_1\mathbf{X}_3] + \dots + E[2\mathbf{X}_{n-1}\mathbf{X}_n] \quad (7)$$

$$= \sigma_1^2 + \dots + \sigma_n^2 + 2\sigma_{12} + \dots + 2\sigma_{13} + \dots + \sigma_{n-1,n}. \quad (8)$$

If \mathbf{X} has a covariance matrix where $\sigma_{ij} \geq 0$ for all $i \neq j$ (positively correlated), then it can be seen from (8) that the case that minimizes the second moment of \mathbf{W} is when all $\sigma_{ij} = 0$ for $i \neq j$ (i.e., when the random variables \mathbf{X}_i are independent). The greater the correlation, the larger the final terms in (8), thus leading to a larger second moment. Furthermore note that the mean of W , $\mu = E[\mathbf{X}_1] + \dots + E[\mathbf{X}_n]$ is the same, regardless of whether \mathbf{X} is composed of independent or correlated random variables, and thus the variance is also minimized when there is no correlation.

Thus, if the \mathbf{X}_i 's are assumed to be the delay through each inverter in the inverter chain, and W the total delay of the chain, then it can be seen that a chain that has correlations within it will have a larger variation. Since the means are the same regardless of the correlation assumptions, then there is a larger probability that the maximum delay will be larger in a correlated set. 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.

C. 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. Furthermore, we have shown that if the nonnegativity assumption is used, then the "worst case" timing yield corresponds to a situation where there are strong correlations within a path (Section IV-B), but total independence path-to-path (Section IV-A1). Conversely, the "best case" timing yield is the reverse: strong correlation path-to-path and total independence within-path.

While the "best case" and "worst case" options use within-path and between-path correlation assumptions that are at odds with each other, they are useful to consider since they provide bounds on the effects of correlations on a parallel system. We use these bounds since we do not have sufficient data to pick a single correlation assumption that is appropriate for all processes.

V. RESULTS

A. Technology

All simulation results reported in this section are based on HSPICE, using Berkeley Predictive Technology Models (BPTM)² for a 70-nm technology. For large widths, the transistors in the process have threshold voltages of approximately 210 and -190 mV for nMOS and pMOS transistors, respectively. The transistor models were expanded to include gate tunnelling leakage which was modelled using a combination of four voltage-controlled current-sources, as in [13]. The resulting transistor macromodel was fitted to industrial data found in [14].

B. 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 any 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 of this operation are shown in Fig. 5; note that the V_t of the transistors is not changed as the supply voltage is lowered.

The dynamic power consumption of the parallel system is decreased with a reduced supply voltage with nearly a constant slope, and is due to the quadratic decrease in dynamic power with voltage, which is countered by the linear increase in the number of blocks. The leakage power exhibits 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 total leakage due to the larger number of blocks; but as 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 for all nodes in the circuit, α . While the dynamic power decreases with a reduced supply voltage, the total power

²[Online]. Available: <http://www-device.eecs.berkeley.edu/~ptm/>

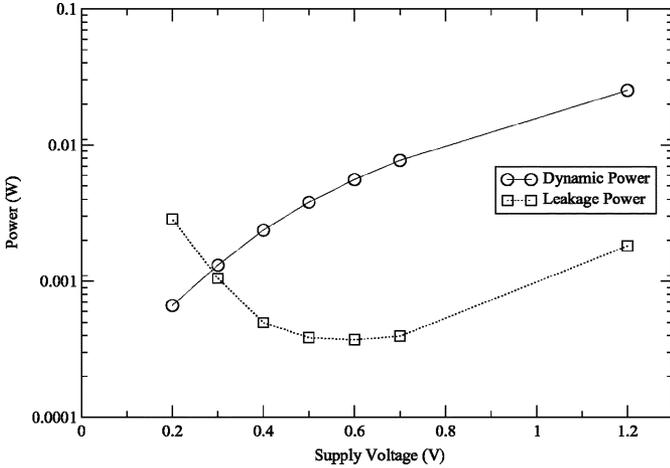


Fig. 5. Effect of supply voltage on dynamic and leakage power. Note that m is increasing as the supply voltage is reduced to maintain throughput (activity factor, $\alpha = 0.1$).

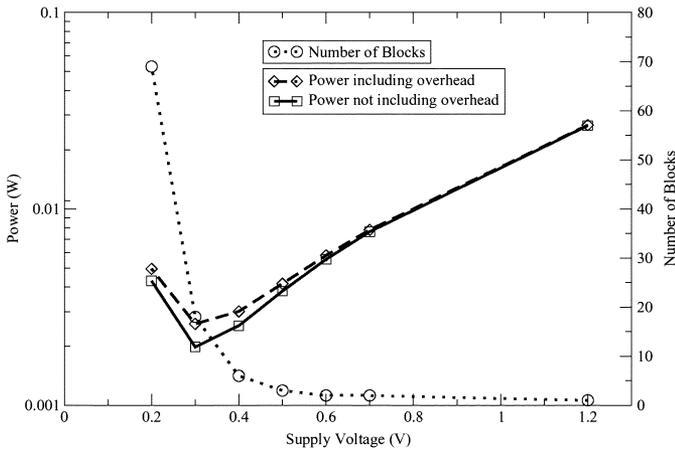


Fig. 6. Effect of lowering the supply voltage on the total power and the number of blocks (activity factor, $\alpha = 0.1$).

of the system increases rapidly near V_t since the delay starts to increase exponentially causing a rapid increase in the number of blocks m which causes a large increase in leakage [2]. These two trends result in a minimum energy point [2] as is shown in Fig. 6, based on $\alpha = 0.1$. It is found that the best operating point is at 0.3 V, with $m = 18$ blocks in parallel, providing a $10.3\times$ reduction in the power consumption of the system relative to the original HV system. 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 irrespective of the value of α : there is an optimal design point at a specific supply voltage and the power savings can be large.

In parallelizing a system, there is also some overhead involved which must also be considered. The overhead 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

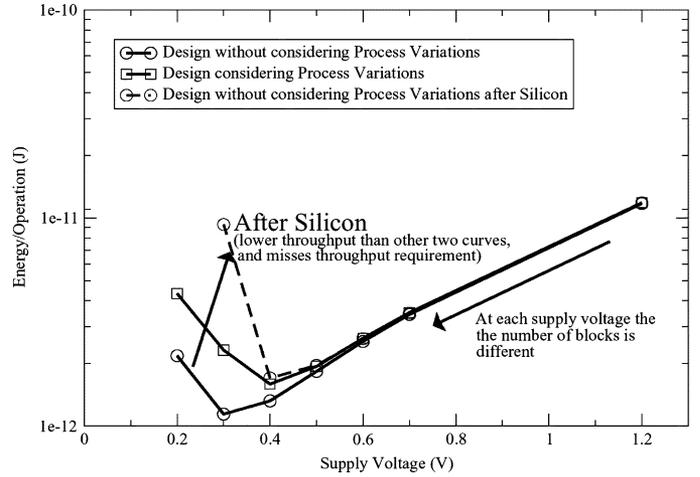


Fig. 7. Expected energy/operation by considering or not considering variations. As the supply voltage is lowered the number of blocks increases to maintain the throughput requirement. When variations are not considered in the analysis but exist on Silicon, the already-set number of blocks slow down, and thus cannot meet the throughput requirement.

[2]. We assume that there are 32 signals at the input and output of each block which results in the overhead composing around 25% of the total power at the minimum energy point.

C. 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, which are inevitable, on the performance of that chip? Although we did not actually measure any data on real hardware, we illustrate what the answer would be in Fig. 7, which assumes independence between path delays and independence between V_t 's at 30°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 V-B. 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.” Since the throughput of such a system, when variations are considered, will not be able to meet the throughput requirement and will have a lower throughput than the bottom (solid) curve, the plot uses *energy/operation* as the metric instead of *power*. 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 0.3 V, but if that design is implemented on Silicon, the energy/operation would be $8.2\times$ higher than expected [9]. 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 IV, one obtains the results shown in the middle curve in Fig. 7. Not only does the middle curve show a system where the throughput requirements are met, but the energy/operation is *much* improved

at lower voltages, showing conclusively that process variations must be taken into account, as we have described.

At the optimal supply voltage, determined when considering process variations, the energy/operation of the top and middle curves obtained by not considering and considering process variations, respectively, are close to each other. However, this does not mean that one is able to not consider process variations and use the lower curve; since the throughput at the two points are not equal, the top curve would give lower and unacceptable system throughput while the middle curve guarantees the required throughput (in this case the difference in throughput in the curves is around 30%). Furthermore, if one does not consider process variations one would not know that the optimal supply voltage is at 0.4 V to design the system with the required number of blocks at that supply voltage; instead one would have to pick the number of blocks needed at 0.3 V.

Notice also that simply sweeping the supply voltage does not mean that one is traversing the top curve, because each point on the curve corresponds to a different number of blocks. If one did not consider process variations and designed the system with the number of blocks needed at 0.3 V and then increased the supply voltage to find the optimal energy/operation, one would have a much larger number of blocks than necessary and incur large power increases of up to $5\times$ that are avoidable if one considers process variations up front.

The minimum energy point occurs at a higher supply voltage when considering variations because WID process variations cause an increase in the number of blocks that are needed to maintain throughput since the system speed is set by the slowest path. Thus, to reduce the number of blocks and limit the effects of leakage the optimal supply voltage increases; effectively the trends that formed the minimum energy point when not considering process variations are forced to occur at higher voltages. This increase in the optimal supply voltage in the parallel system due to WID variations is similar to the same effect seen in non-parallel subthreshold circuits [10]. As a result of our analysis, one sets the supply voltage to 0.4 V, leading to an energy/operation of the system that is $7.4\times$ lower than the original HV system [9].

D. Effect of Correlation

Fig. 8 shows the number of blocks that are needed in order to maintain throughput under different correlation assumptions: 1) strong correlations with a path and total independence path-to-path and 2) independence with a path and total correlation path-to-path. As explained in Section IV, if we use the nonnegativity assumption these cases lead to “worst case” and “best case” timing yield which is how we label them on the figure; “worst case” timing yield in this case, corresponds to where there are strong correlations within a path, but total independence path-to-path. The “best case” timing yield is the reverse: strong correlation path-to-path and total independence within-path. The results shown in Fig. 8 are true irrespective of the nonnegativity assumption, and only the labelling of the curves as “best case” and “worst case” depend on the nonnegativity assumption.

It can be seen that when no variations were assumed, only 18 blocks were needed at 0.3 V, but when variations are included in

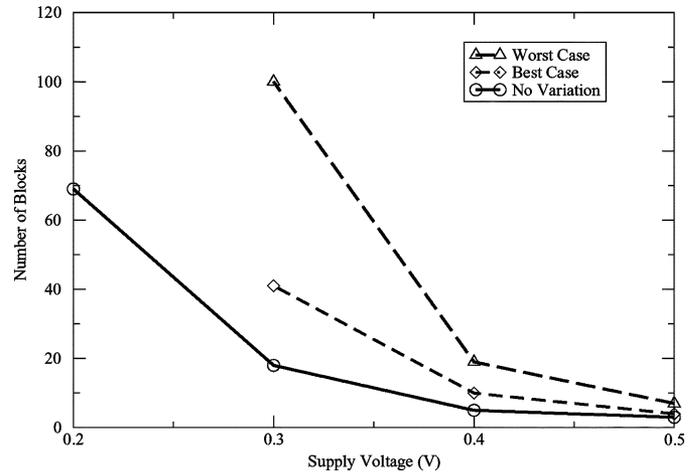


Fig. 8. Effect of variations on the number of blocks needed to maintain the throughput at each supply voltage.

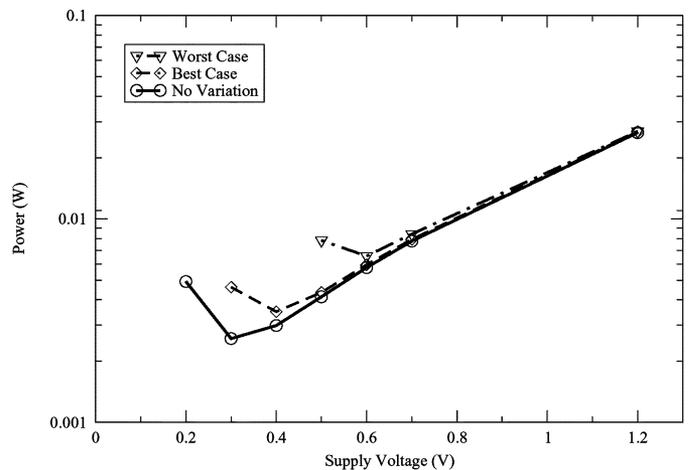


Fig. 9. Effect of correlation on total power.

the analysis, the number of blocks needed at 0.3 V ranges from 41 to 100 blocks. The increase in the number of blocks is due to the WID variations, which causes some blocks to be slowed and lowering the throughput of the system, thus necessitating an increase in the number of blocks to regain the lost throughput. 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. 9 shows the power consumption of the parallel system under different levels of correlation, at 30°C with $\alpha = 0.1$. As before, each point on the plot represents a possibly different number of blocks in parallel, as determined by the procedure described in Section IV-A. 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 since variations cause an increase in the number of block which move the minimum energy point to a higher supply voltage as explained in Section V-C. This is not to say that it is better to ignore variations, because, as we saw in Section V-C, 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 0.3 V is optimal. However, with variations considered up-front, the best case curve gives an optimal V_{dd} of 0.4 V, and the worst case gives a $V_{dd} = 0.6$ V (leading to a power reduction of $7.6\times$ and $4\times$, respectively, compared to the original system) [9]. 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 [8], it becomes interesting to consider the impact of designing with one set of assumptions. If, for example, the “best case” assumptions were used (strong path-to-path correlation and within-path independence) and then found to be incorrect, then Fig. 9 shows that 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. This final conclusion depends on the nonnegativity assumption since it depends on the true power curve lying in between the “best case” and “worst case” curves. If the nonnegativity assumption is not used, the results would still remain true and using the “best case” assumption would still be the more conservative in terms of power than the “worst case” assumption, and the “worst case” assumptions would be more conservative in terms of performance than the “best case” assumptions, but they would not be the **most** conservative options.

E. Effect of Activity Factor

The previous results and conclusions used an activity factor α of 0.1; in this section, we will show that in general our conclusions about the optimal supply voltage hold true regardless of the 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, can affect the supply voltage that minimizes the power consumption.

With larger α 's, larger power reductions are possible (even though the absolute power would still be higher), and the supply voltage that provides the largest power reduction becomes lower. The reason for a lower optimal supply voltage when activity factors are higher is that, with larger α 's, the static power consumption is less important, and thus the increased parallelism that is needed at very low voltages is not at issue. An opposite argument holds when the activity factor is lower, and thus the supply voltage has to be increased to limit the parallelism.

The solid plot in Fig. 10 shows the optimal supply voltage at different α 's. Observe that, as α tends toward 1, the optimal

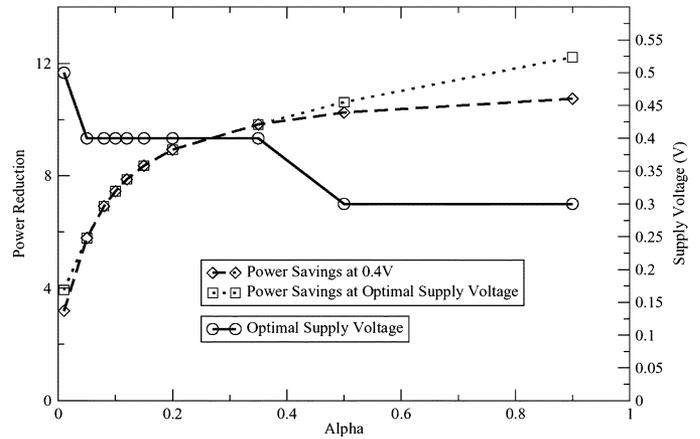


Fig. 10. Optimal voltage and power reduction at different α 's.

supply voltage is reduced, because the static power becomes less important compared 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 [9].

Also in Fig. 10 is a comparison of the power reduction that is possible when using the optimal supply voltage to the power reduction when using a supply voltage of 0.4 V (for our circuit, the supply voltage that maximizes the power reduction at an activity factor of 0.1). Since the two curves are close 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.

F. Changing Transistor Characteristics

As the voltage is being lowered we also optimize the circuit by changing the transistor characteristics in the critical paths.³ The width of the transistors in the critical paths are changed from being minimum width to up to $4\times$ minimum width for nMOS transistors and up to $8\times$ minimum width for pMOS transistors. 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. While at high supply voltages, small transistors are preferable, as increases in performance by using wider transistors are offset by increased power consumption, wider transistors are necessary at low voltages. At lower supply voltages, where process variations have a large effect, using wider transistors increases performance *and* decreases the variation in V_t . The variation in V_t due to random dopant fluctuations has an inverse relationship to the transistor area (specifically $1/\sqrt{WL}$) and thus larger transistors have a smaller variation in V_t [15] causing a smaller variation in the transistor performance at low voltages [10]. 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 similar to what is seen for nonparallel subthreshold circuits in [10], but opposite to what is seen in [16] where WID process variations were not considered, and minimum width transistors were found to be optimal at low

³The width of the transistors in the noncritical paths remain at minimum width

voltages. In the results presented thus far, we have used the transistor width that minimizes the power at each data point; at the minimum energy points (for both worst case and best case correlation assumptions) the width that minimizes the energy/operation is still the minimum size; only at supply voltages which are lower than the supply voltage that minimizes the power consumption are larger transistors beneficial.

VI. 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. When considering the worst-case operating conditions the benefits of the LV system are reduced, and thus we introduce the TDDS which allows large energy reductions in parallel systems regardless of the operating conditions [9]. The TDDS turns off a portion of the parallel blocks as operating conditions improve to maintain throughput but reduce energy consumption.

At traditional supply and threshold voltages, the speed of circuits depends on the on-current through the transistor. It is well known that both the mobility of charge carriers (which effects the on-current) and the threshold voltage decrease with increased temperature; these two effects usually lead to a net performance decrease with increased temperature at high voltages. This behavior, however, is no longer true when supply voltages become lower, and the decrease in the threshold voltage can have a larger impact on the total performance of the transistor allowing circuits to speed up as the temperature is increased. Consequently, at low voltages the performance of blocks will be lower at low temperatures and thus the number of blocks needed must be chosen at low temperature.

As a parallel system is designed for the worst-case (lowest) temperatures, we must consider what happens at higher temperatures which invariably will be encountered during circuit operation. At high temperatures the number of blocks would often be larger than needed and the circuit would be dissipating more power than needed. 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 would hopefully offset most of the increase in power as the temperature is increased. A possible implementation of this scheme is shown in Fig. 11. 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. To turn off the different blocks, many different techniques have been presented in the literature [17] such as

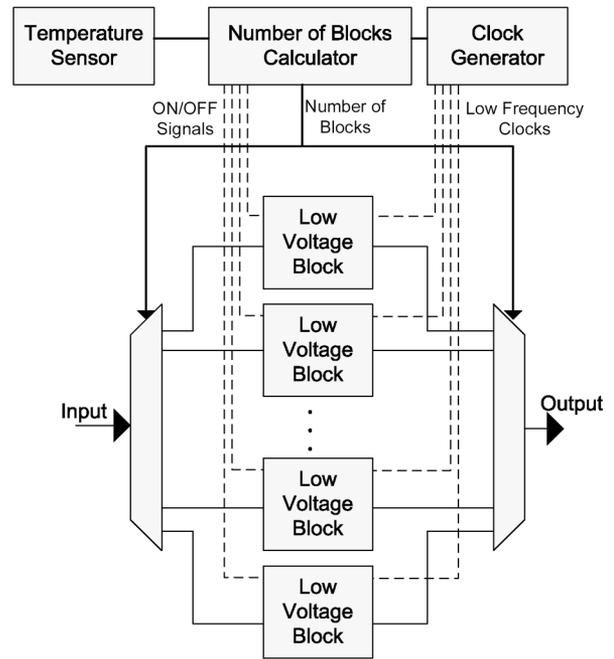


Fig. 11. Organization of temperature dependent deactivation.

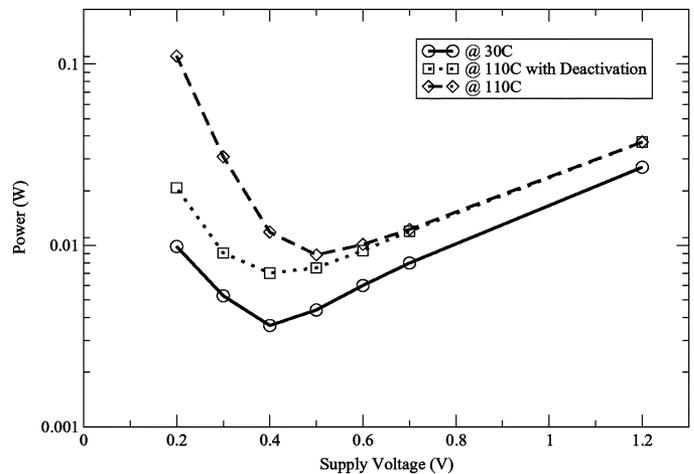


Fig. 12. Effect of temperature variations on power.

sleep transistors, which have limited effects on the performance of the block they are controlling when the block is turned on.

This 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 an HV system.

Fig. 12 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 0.3 V is larger than that of the original system. When TDDS is used, the increase in the power consumption is limited as blocks that are

unused are turned off at high temperatures [9]. At 0.4 V, when using TDDS, there is a power increase of $1.9\times$ when the temperature changes from 30 °C to 110 °C instead of an increase of $3.3\times$ when not using TDDS.

VII. CLUSTERING

Now that a methodology exists for determining the effects of process variations on a parallel system, we can look at different architectures and circuit schemes that try to limit the effect of the underlying variations on the performance of the parallel system to further reduce the power consumption. The main reason that process variations affect the throughput and consequently the power consumption of a parallel system is that all blocks that are in parallel have to operate at the speed of the slowest block. If this were not the case, then the blocks which are faster could operate at a faster speed thus allowing for increased throughput, which would allow the number of blocks that are required in parallel to be reduced, further reducing power.

One way of reducing the dependence of the system throughput on the speed of the slowest block is to have different clusters of parallel blocks, and then apply small differences in the supply voltage to each cluster to equalize the performance of the clusters. For example, we could speed up all clusters to the speed of the fastest cluster by increasing the supply voltage of the slower clusters. In this way, blocks that are extremely slow or extremely fast only affect the power consumption of the blocks in their cluster. It will be seen that clustering with voltage differences can reduce the power consumption and provide significant benefit for multicore microprocessor systems and for further scaled technology.

A. Results of Clustering

By using clustering, the energy/operation was reduced by 3.7% when using two clusters and by 4.7% when using three clusters.⁴ The average voltage difference between each cluster is 20 mV and the maximum voltage difference from the nominal is 40 mV. All simulation results reported in this and subsequent sections are based on a commercial 65-nm technology, where the optimal supply voltage is at 0.5 V.

Given the limited benefit of clustering around 5% seen before for functional blocks, we have looked at the parallelization of much larger blocks such as complete microprocessor cores. Multicore microprocessors have already started to appear to increase the performance of systems [18]–[21], and in this case, we are looking at multicore microprocessors to instead lower the power consumption at the same throughput.

Thus we have reperformed our analysis with the core block size being increased from 168 to 2500 μm which can represent a small microprocessor core [19]. For a system that uses these larger parallel blocks, the reduction in energy/operation by clustering is slightly more beneficial, leading to a reduction of over 5% when three clusters are used.

1) *Unlimited Clustering*: Given the size of each of the blocks and the limited size of a number of blocks in parallel it becomes

⁴To remain focused on the effects of clustering the assumption that correlation exists within a path, and between paths will be used throughout this section. Using other correlation assumptions show very similar relative decreases in energy due to clustering.

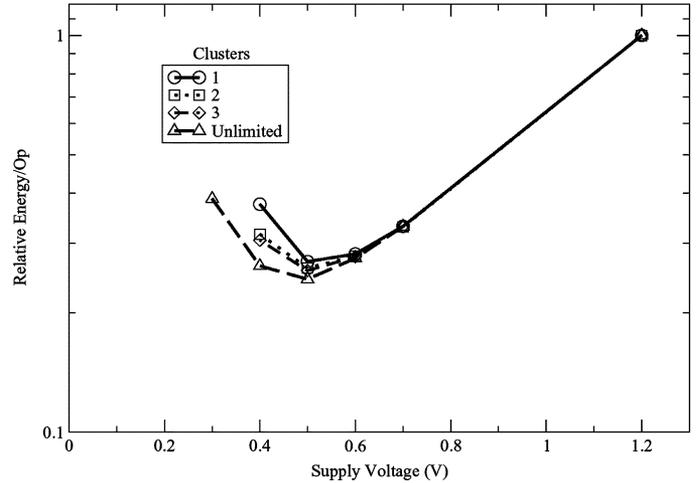


Fig. 13. Energy/operation by using unlimited clustering.

possible to provide each block its own voltage regulator that changes the supply voltage slightly relative to the global low supply voltage. Thus, each block is in its own cluster, achieving an *unlimited clustering* scheme.

With unlimited clustering the optimal supply voltage stays at 0.5 V, but there is an additional reduction in the energy/operation since the effect of the variations can be compensated for a thus a smaller number of blocks is needed. In Fig. 13, it can be seen that there is almost a 10% reduction in energy/operation at the optimal supply voltage, more than any of the other clustering options. In this analysis, we assume that the extra supply voltages can be delivered with 100% efficiency and 0% overhead.

B. Future Trends

As process technologies continue to scale, there will be increased variation [22], [23]. Furthermore, HV multicore processor systems will become more prevalent [18]–[21]. In this section, the effects of increasing variation and of multicore HV reference systems on the benefit of a clustered LV parallel system will be explored.

1) *Increased Variation*: As the variation increases the optimal supply voltage in an unclustered LV parallel system increases to limit the parallelism to reduce the effect of process variations at LVs. This necessitated increase in supply voltage reduces the power and energy savings that were possible by further reducing the supply voltage, and thus limiting the benefits of using a parallel system.

By clustering the different blocks, however, an LV parallel system can compensate for these variations and keep the optimal supply voltage low. For example, if the standard deviation of the V_t variation is increased by 50%, the optimal supply voltage of an unclustered system increases to 0.6 V. However, for the unlimited clustered system, regardless of the variation, the optimal supply voltage stays at 0.5 V. Fig. 14 shows the reduction in the energy/operation as the standard deviation of the V_t variation is increased by $1.5\times$, $2\times$, and $3\times$. It can be seen that for a $1.5\times$ increase in the standard deviation of the V_t variation, clustering is able to help reduce the energy/operation even further; for example, three clusters can reduce the energy/operation

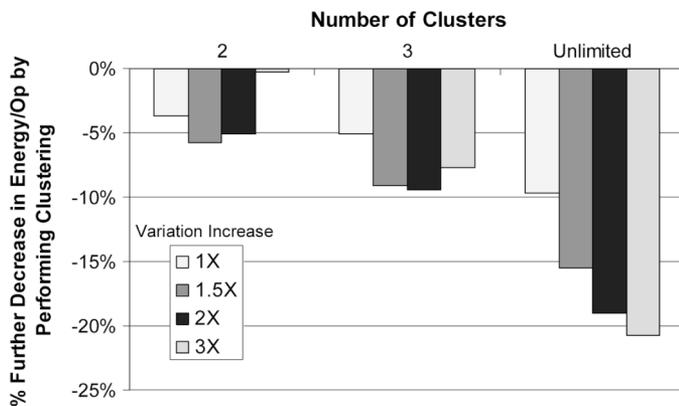


Fig. 14. Reduction in energy/operation with increased variation.

tion by more than 9% rather than the 5% when the variation was not increased. This further reduction in the energy/operation is due to the unclustered system's higher energy/operation in the presence of larger variation, rather than a lower energy/operation in the clustered systems.

For larger increases of variation, however, a small amount of clusters is not able to decrease the energy/operation considerably because the likelihood of a very slow block appearing in each is quite high. For example, using two clusters in the presence of a $3\times$ increase in variation is not able to reduce the energy/operation at all. The energy/operation benefit of using three clusters also sees a drop as the variation is increased to $3\times$. This behavior is due to the increasing effect of the variation and for which clustering, in limited amounts, cannot overcome. The benefits of unlimited clustering, however, continue to increase as the variation increases since the technique can always compensate for the variations, showing a 15% and 19% energy/operation reduction with a $1.5\times$ and $2\times$ increase in V_t , respectively.

Considering that there are only around ten blocks in parallel and that scaling will continue to decrease the size of logic blocks, and increase the underlying variation, unlimited clustering can provide a significant benefit in lowering the energy/operation.

VIII. CONCLUSION

Power consumption is increasingly becoming the barrier in submicrometer integrated circuit design. An 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 larger than what was expected. We have presented a new methodology that takes WID variations into consideration when designing a parallel system and showed that the supply voltage that minimizes 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 $7.6\times$.

We further 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.

To further limit the effect of variations, and allow for a reduced power consumption, we analyzed the effects of clustering. It was shown that providing different voltages to each cluster can provide a further 10% reduction in energy/operation to a LV parallel system, and that the savings by clustering increase as technology scales.

ACKNOWLEDGMENT

This work was done in part at Intel Corporation. The authors would like to thank K. Heloue for providing assistance.

REFERENCES

- [1] J. Schutz and C. Webb, "A scalable X86 CPU design for 90 nm process," in *Proc. IEEE Int. Solid-State Circuits Conf.*, 2004, pp. 62–513.
- [2] A. P. Chandrakasan and R. W. Brodersen, *Low Power Digital CMOS Design*. Norwell, MA: Kluwer Academic, 1995.
- [3] D. Liu and C. Svensson, "Trading speed for low power by choice of supply and threshold voltages," *IEEE J. Solid-State Circuits*, vol. 28, no. 1, pp. 10–17, Jan. 1993.
- [4] S. Nassif, "Delay variability: Sources, impacts, trends," in *Proc. ISSCC*, 2000, pp. 368–369.
- [5] S. Borkar, T. Karnik, S. Narendra, T. Tschanz, A. Keshavarzi, and V. De, "Parameter variations and impact on circuits and microarchitecture," in *Proc. Design Autom. Conf.*, 2003, pp. 338–342.
- [6] C. Kim, H. Soeleman, and K. Roy, "Ultra-low-power DLMS adaptive filter for hearing aid applications," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 11, no. 6, pp. 1058–1067, Dec. 2003.
- [7] M. Eisele, "The impact of intra-die device parameter variations on path delays and on the design for yield of low voltage digital circuits," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 5, no. 4, pp. 360–368, Dec. 1997.
- [8] D. Boning and S. Nassif, "Models of process variations in device and interconnect," in *Design High-Performance Microprocessor Circuits*. Piscataway, NJ: IEEE Press, 2001.
- [9] N. Azizi, M. M. Khellah, V. De, and F. N. Najm, "Variations-aware low-power design with voltage scaling," in *Proc. ACM/IEEE Design Autom. Conf.*, 2005, pp. 529–534.
- [10] B. Zhai, S. Hanson, D. Blaauw, and D. Sylvester, "Analysis and mitigation of variability in subthreshold design," in *Proc. ACM/IEEE Int. Symp. Low-Power Electron. Design*, 2005, pp. 20–25.
- [11] Y. L. Tong, *The Multivariate Normal Distribution*. New York: Springer-Verlag, 1990.
- [12] J. Luo, S. Sinha, Q. Su, J. Kawa, and C. Chiang, "An IC manufacturing yield model considering intra-die variations," in *Proc. ACM/IEEE Design Autom. Conf.*, 2006, pp. 749–754.
- [13] D. Lee, W. Kwong, D. Blaauw, and D. Sylvester, "Simultaneous subthreshold and gate-oxide tunneling leakage current analysis in nanometer CMOS design," in *Proc. ISQED*, 2003, pp. 287–292.
- [14] W. K. Henson, "Analysis of leakage currents and impact on off-state power consumption for CMOS technology in the 100-nm regime," *IEEE Trans. Electron Devices*, vol. 47, no. 2, pp. 440–447, Feb. 2000.
- [15] D. J. Frank, R. H. Dennard, E. Nowak, P. M. Solomon, Y. Taur, and H. S. P. Wong, "Device scaling limits of Si MOSFETs and their application dependencies," *Proc. IEEE*, vol. 89, no. 3, pp. 259–288, Mar. 2001.
- [16] B. H. Calhoun, A. Wang, and A. Chandrakasan, "Device sizing for minimum energy operation in subthreshold circuits," in *Proc. Custom Integr. Circuits Conf.*, 2003, pp. 95–98.
- [17] S. Mutoh, T. Douseki, Y. Matsuya, T. Aoki, S. Shigematsu, and J. Yamada, "1-V power supply high-speed digital circuit technology with multithreshold-voltage CMOS," *IEEE J. Solid-State Circuits*, vol. 30, no. 8, pp. 847–854, Aug. 1995.
- [18] S. Naffziger, E. Stackhouse, and T. Grutkowski, "The implementation of a 2-core multi-threaded Itanium-family processor," in *Proc. IEEE Int. Solid-State Circuits Conf.*, 2005, pp. 182–183.

- [19] A. S. Leon, J. L. Shin, K. W. Tam, W. Bryg, F. Schumacher, P. Konge-tira, D. Weisner, and A. Strong, "A power-efficient high-throughput 32-thread SPARC processor," in *Proc. IEEE Int. Solid-State Circuits Conf.*, 2006, pp. 98–99.
- [20] M. Golden, S. Arekapudi, G. Dabney, M. Haertel, S. Hale, L. Herlinger, Y. Kim, K. McGrath, V. Palisetti, and M. Singh, "A 2.6 GHz dual-core 64b-x86 microprocessor with DDR2 memory support," in *Proc. IEEE Int. Solid-State Circuits Conf.*, 2006, pp. 104–105.
- [21] E. B. Cohen, N. J. Rohrer, P. Sandon, M. Canada, C. Lichtenau, M. Ringle, P. Kartschoke, R. Floyd, M. Ross, T. Phueger, R. Hilgendorf, P. McCormich, G. Salem, J. Connor, S. Geissler, and D. Thygesen, "A 64B CPU pair: Dual- and single-processor chips," in *Proc. IEEE Int. Solid-State Circuits Conf.*, 2006, pp. 106–107.
- [22] ITRS, U.S., "International technology roadmap for semiconductors," 2005.
- [23] C. Visweswariah, "Death, taxes and failing chips," in *Proc. ACM/IEEE Design Autom. Conf.*, 2003, pp. 343–347.



Navid Azizi (S'03) received the B.A.Sc. degree (with honors) in computer engineering, and the M.A.Sc. and Ph.D. degrees in electrical engineering from the University of Toronto, Toronto, ON, Canada, in 2001, 2003, and 2007, respectively.

In June 2007, he joined Altera Corporation, Toronto, as a Senior Software Engineer. He has previously worked at Altera Corporation, Toronto, and with Intel Corporation, Hillsboro, OR. His research interests include low-leakage and low-variability digital integrated circuit design and analysis.

Mr. Azizi was a recipient of the Postgraduate Scholarship Award from the National Sciences and Engineering Research Council (NSERC) of Canada for 2001–2003 and 2003–2005.



Muhammad M. Khellah (M'03) received the B.Sc. degree from KFUPM, Saudi Arabia, in 1991, the M.A.Sc. from the University of Toronto, Toronto, ON, Canada, in 1994, and the Ph.D. degree from the University of Waterloo, Waterloo, ON, Canada, in 1999, all in electrical and computer engineering.

He is a Staff Research Engineer at the Circuits Research Laboratory (CRL), Intel Corporation, Hillsboro, OR. He is currently working on advanced low-power memory and interconnect circuits. He first joined Intel in 1999, where he designed embedded SRAM caches for P3 and P4 microprocessor products.

He has authored or co-authored 25 papers in refereed international conferences and journals. He has 23 patents granted and 38 more patents filed. He is on the TPC of several IEEE conferences.



Vivek K. De (SM'07) received the Bachelor's degree in electrical engineering from the Indian Institute of Technology in Madras, Madras, India, in 1985, the Master's degree in electrical engineering from Duke University, Durham, NC, in 1986, and the Ph.D. degree in electrical engineering from Rensselaer Polytechnic Institute (RPI), Troy, NY, in 1992.

He is an Intel Fellow and Director of Circuit Technology Research in the Corporate Technology Group. He joined Intel in 1996 as a Staff Engineer in the Circuits Research Laboratory (CRL), Hillsboro, OR.

Since that time he has led research teams in CRL focused on developing advanced circuits and design techniques for low-power and high-performance processors. In his current role, he provides strategic direction for future circuit technologies and is responsible for aligning CRL's circuit research with technology scaling challenges. He has published 147 technical papers in refereed conferences and journals and six book chapters on low-power circuits. He holds 117 patents, with 73 more patents filed (pending). He received an Intel Achievement Award for his contributions to a novel integrated voltage regulator technology. Prior to joining Intel, he was engaged in semiconductor devices and circuits research at Rensselaer Polytechnic Institute and Georgia Institute of Technology, Atlanta, and was a Visiting Researcher at Texas Instruments, Dallas, TX.



Farid N. Najm (F'03) received the B.E. degree in electrical engineering from the American University of Beirut (AUB), Beirut, Lebanon, in 1983, and the M.S. and Ph.D. degrees in electrical and computer engineering (ECE) from the University of Illinois at Urbana-Champaign (UIUC), Urbana, in 1986 and 1989, respectively.

In 1999, he joined the ECE Department at the University of Toronto, Toronto, ON, Canada, where he is currently a Professor and Vice-Chair of ECE. From 1989 to 1992, he worked with Texas Instruments in

Dallas, TX. He then joined the ECE Department, UIUC, as an Assistant Professor, becoming an Associate Professor in 1997. He is the co-author of *Failure Mechanisms in Semiconductor Devices*, (2nd Ed., Wiley, 1997). His research is on computer-aided design (CAD) for integrated circuits, with an emphasis on circuit level issues related to power dissipation, timing, and reliability.

Dr. Najm is an Associate Editor for the IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN ON INTEGRATED CIRCUITS AND SYSTEMS. He was a recipient of the IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN ON INTEGRATED CIRCUITS AND SYSTEMS Best Paper Award in 1992, the National Science Foundation (NSF) Research Initiation Award in 1993, the NSF CAREER Award in 1996, and was an Associate Editor for the IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION (VLSI) SYSTEMS from 1997 to 2002. He served as General Chairman for the 1999 International Symposium on Low-Power Electronics and Design (ISLPED'99), and as Technical Program Co-Chairman for ISLPED'98. He has also served on the technical committees of ICCAD, DAC, CICC, ISQED, and ISLPED.