

# Invited Paper

## Extended Dynamic Voltage Scaling for Low Power Design

Bo Zhai, David Blaauw, Dennis Sylvester, \*Krisztian Flautner  
 {bzhai, blaauw, dennis}@umich.edu, University of Michigan, Ann Arbor, MI  
 \* krisztian.flautner@arm.com, ARM Ltd., Cambridge, UK

### Abstract

Dynamic voltage scaling (DVS) is a popular approach for energy reduction of integrated circuits. Current processors that use DVS typically have an operating voltage range from full to half of the maximum V<sub>dd</sub>. However, it is possible to construct designs that operate over a much larger voltage range: from full V<sub>dd</sub> to subthreshold voltages. This possibility raises the question of whether a larger voltage range improves the energy efficiency of DVS. First, from a theoretical point of view, we show that for subthreshold supply voltages leakage energy becomes dominant, making "just in time completion" energy inefficient. We derive an analytical model for the minimum energy optimal voltage and study its trends with technology scaling. Second, we compare several different low-power approaches including MTCMOS, standard DVS and extended DVS to subthreshold operation. Study of real applications on commercial processor shows that extended DVS has the best energy efficiency. Therefore, we conclude that extending the voltage range below V<sub>dd</sub>/2 will improve the energy efficiency for most processor designs.

### 1 Introduction

Due to technology scaling, microprocessor performance has increased tremendously albeit at the cost of higher power consumption. Energy efficient operation has therefore become a very pressing issue, particularly in mobile applications that are battery operated. Dynamic voltage scaling (DVS) was proposed as an effective approach to reduce energy use and is now used in a number of low-power processor designs [1][2][3].

Most applications do not always require the peak performance from the processor. Hence, in a system with a fixed performance level, certain tasks complete ahead of their deadline and the processor enters a low-leakage sleep mode [4] for the remainder of the time. This operation is illustrated in Figure 1(a).

In DVS systems however, the performance level is reduced during periods of low utilization such that the processor finishes each task "just in time," stretching each task to its deadline, as shown in Figure 1(b). As the processor frequency is reduced, the supply voltage can be reduced. As shown by the equations below<sup>1</sup>, the reduction in frequency [5] combined with a quadratic reduction from the supply voltage results in an approximately cubic reduction of power consumption. However, with reduced frequency the time to complete a task increases, leading to an overall quadratic reduction in the energy to complete a task.




Figure 1. Illustration of optimal task scheduling

1. The 1.3-power [5] scaling of current is only valid for high supply voltages when carrier velocity saturates. Subthreshold scaling of the supply voltage with performance for low voltage operation will be extensively discussed in Section 3.

$$Delay = \frac{1}{f} = \frac{C_s V_{dd}}{I_{dsat}} \propto \frac{V_{dd}}{(V_{dd} - V_{th})^{1.3}} \quad Power \propto f V_{dd}^2 \quad Energy \propto V_{dd}^2$$

DVS is therefore an effective method to reduce the energy consumption of a processor, especially under wide variations in workload that are increasingly common in mobile applications. Hence, extensive work has been performed on how to determine voltage schedules that maximize the energy savings obtained from DVS [4][8].

In most current DVS processor designs, the voltage range is limited from full V<sub>dd</sub> to approximately V<sub>dd</sub>/2. In Table 1, the available range of operating voltages and associated performance levels are shown for three commercial designs. The lower limit of voltage scal-

Table 1. Commercial processor designs and range of voltage scaling

|                             | Voltage Range | Frequency Range |
|-----------------------------|---------------|-----------------|
| IBM PowerPC 405LP [3]       | 1.0V-1.8V     | 153M-333M       |
| TransMeta Crusoe TM5800 [1] | 0.8V-1.3V     | 300M-1G         |
| Intel XScale 80200 [2]      | 0.95V-1.55V   | 333M-733M       |

ing is typically dictated by voltage and noise-sensitive circuits, such as pass-gates, PLLs, and sense amps and results from applying DVS to a processor "as is" without special redesign to accommodate operation over a wide range of voltage levels. However, it is well known that CMOS circuits can operate over a very large range of voltage levels down to less than 200mV. In such "subthreshold" operating regimes, the supply voltage lies below the threshold voltage and the circuit operates using leakage currents. Work has been reported on designs that operate at subthreshold voltages [6][7] and it was reported that the minimum allowable supply voltage of a functional CMOS inverter is 36mV [9]. An example of a commercial product that uses subthreshold operation for extremely low power applications is shown in [10].

With some additional design effort, it is possible to significantly extend the operating voltage range of processors. One issue that needs to be addressed is the determination of a lower limit of the voltage range for optimal energy efficiency. The optimal voltage limit depends on two factors: the power/delay trade-offs at low operating voltages and the workload characteristics of the specific processor. In this paper we address both of these issues.

First, we show that the quadratic relationship between energy and V<sub>dd</sub> deviates as V<sub>dd</sub> is scaled down into the subthreshold region of MOSFETs. In subthreshold operation the "on-current" takes the form of subthreshold current, which is exponential with V<sub>dd</sub>, causing the delay to increase exponentially with voltage scaling. Since leakage energy is linear with the circuit delay, the fraction of leakage energy increases with supply voltage reduction in the subthreshold regime. Although dynamic energy reduces quadratically, at very low voltages where dynamic and leakage energy become comparable, the total energy can increase with voltage scaling due to the increased circuit delay. In this paper, we derive an analytical model for the voltage that minimizes energy and we show that it lies well above the previously reported [9] minimal allowable operating voltage of 36mV. We verify our model using SPICE and also study its trends as a function of different design and process parameters. As one of the results, our work shows that operation at voltages well below threshold is never energy-efficient.

A second issue that determines whether to apply such aggressive, or extended, DVS (which we refer to as INSOMNIAC) is based on the workload characteristics of the processor. Clearly it is not necessary to extend the voltage range below that which is needed based on

the expected workload of the processor. To analyze the energy efficiency of different low-power schemes, including MTCMOS and INSOMNIAC, we compare the energy consumption for a number of workload traces obtained from a processor running a wide range of applications. Our results show that most applications benefit significantly from an operating voltage range that is wider than what is available in most current DVS processors. Only very low activity applications with long idle times benefit greatly from pure MTCMOS.

The remainder of this paper is organized as follows. Section 2 provides an overview of the voltage limit for functionally correct CMOS logic. Section 3 presents our analysis of the minimum voltage scaling limit for optimal energy efficiency. Section 4 presents our energy modeling for different low-power approaches. Section 5 presents the computation and analysis based on our workload energy modeling in Section 4. Finally, Section 6 contains our conclusions.

## 2 Circuit Behavior at Ultra Low Voltages

Before we derive the energy optimal operating voltage in Section 3, we first briefly review the minimum operating voltage that is required for functional correctness of CMOS logic. The minimum operating voltage was first derived by Swanson and Meindl in [9] and is given as follows:

$$V_{dd, limit} = 2 \cdot \frac{kT}{q} \cdot \left[ 1 + \frac{C_{fs}}{C_{ox} + C_d} \right] \cdot \ln \left( 2 + \frac{C_d}{C_{ox}} \right) \quad (\text{EQ 1})$$

$$\approx 2 \cdot V_T \cdot \ln \left( 2 + \frac{C_d}{C_{ox}} \right)$$

where  $C_{fs}$  is the fast surface state capacitance per unit area,  $C_{ox}$  is the gate-oxide capacitance per unit area, and  $C_d$  is the channel depletion region capacitance per unit area. For bulk CMOS technology, we know that subthreshold swing can be expressed as follows:

$$S_s = \ln 10 \cdot V_T \cdot \left( 1 + \frac{C_d}{C_{ox}} \right) \quad (\text{EQ 2})$$

From this, we can rewrite EQ1 as follows:

$$V_{dd, limit} = 2 \cdot V_T \cdot \ln \left( 1 + \frac{S_s}{\ln 10 \cdot V_T} \right) \quad (\text{EQ 3})$$

$$\approx 52mV \cdot \ln \left( 1 + \frac{S_s}{59.87mV} \right) \text{ at } 300K$$

For 0.18um technology  $S_s$  is typically in the range of 90mV/decade, and therefore

$$V_{dd, limit} = 48mV. \quad (\text{EQ 4})$$

Hence, it is theoretically possible to operate circuits deep into the subthreshold regime given that typical threshold voltages are much larger than 48mV. In fact, SPICE simulation confirms that it is possible to construct an inverter chain that works properly at 48mV, although at this point the internal signal swing is reduced to less than 30mV. Based on SPICE simulations, we find that it is possible to operate a wide range of standard library gates at similar operating voltages and that their delays track relatively closely to that of the inverter. However it is clear that there are practical reasons why operating circuits at the minimum voltage is not desirable, such as susceptibility to noise and process variations [13]. We show in the next section that the minimum operating voltage for functionally correctness also does not provide energy optimal operation.

## 3 Minimum Energy Analysis

We first illustrate the energy dependence on supply voltage using a simple inverter chain consisting of 50 inverters. A single transition is used as a stimulus and energy is measured over the time period necessary to propagate the transition through the chain. The energy-Vdd relation is plotted in Figure 2. It is seen that the dynamic energy component  $E_{active}$  reduces quadratically while the leakage energy,  $E_{leak}$ , increases with voltage scaling. The reason for the increase in leakage energy in the subthreshold operating regime is that as the voltage is scaled below the threshold voltage, the on-current (and hence the circuit delay) increases exponentially with voltage scaling while the off-current is reduced less strongly. Hence, the leakage energy  $E_{leak}$  will rise and supersede the dynamic energy  $E_{active}$  at 180mV. This effect

creates a minimum energy point in the inverter circuit that lies at 200mV, as shown in Figure 2.

In the previous example, if the inverter chain is pipelined logic between two registers we are implicitly assuming that there is always one input transition per clock cycle. However, the switching activity varies across circuits so we include the input activity factor  $\alpha$ , which is the average number of times a node makes a power consuming transition in one clock period. We now derive an analytical expression for the energy of an inverter chain as a function of the supply voltage. Suppose we have an n-stage inverter chain with activity factor of  $\alpha$ . The standard expression for subthreshold current is given by [11]:

$$I_{sub} = \mu_{eff} C_{ox} L_{eff} \frac{W}{m} (m-1) V_T^2 e^{\frac{V_{gs} - V_{th} - V_{off}}{m V_T}} \left( 1 - e^{-\frac{V_{ds}}{V_T}} \right) \quad (\text{EQ 5})$$

where,

$$m = \frac{S_s}{\ln 10 \cdot V_T} = \frac{90}{\ln 10 \times 26} = 1.51 \quad (\text{EQ 6})$$

In EQ6 we again assume  $S_s$  is 90mV/decade. We now express the total energy  $E$  per clock cycle as the sum of dynamic, leakage energy<sup>1</sup>:

$$E = E_{active} + E_{leak}$$

$$= \alpha \cdot n \cdot E_{switch, inv} + P_{leak} \cdot t_d \quad (\text{EQ 7})$$

$$= \alpha \cdot n \cdot \left( \frac{1}{2} \cdot C_s \cdot V_{dd}^2 \right) + (n \cdot V_{dd} \cdot I_{leak}) \cdot (n \cdot t_p)$$

where

- $\alpha$  - circuit switching factor
- $n$  - number of stages
- $E_{switch, inv}$  - switching energy of a single inverter
- $P_{leak}$  - total leakage power of the entire inverter chain
- $t_d$  - delay of the entire inverter chain
- $C_s$  - total switched capacitance of a single inverter
- $I_{leak}$  - leakage current of a single inverter
- $t_p$  - delay of a single inverter

First, we focus on finding an accurate estimate of  $t_p$ . Let  $t_{p, step}$  denote the ideal inverter-delay with a step input and  $t_{p, actual}$  denote the actual inverter delay with an input rising time of  $t_r$ . We can compute  $t_{p, step}$  based on a simple charge-based expression:

$$t_{p, step} = \frac{\frac{1}{2} \cdot C_s \cdot V_{dd}}{I_{on}} \quad (\text{EQ 8})$$

where  $I_{on}$  is the average on-current of a inverter. Furthermore, for normal operating voltages, the step delay can be extended to the actual delay as follows [15],

$$t_{pHL, actual} = \sqrt{2 \cdot t_{pHL, step}^2 + \left( \frac{t_r}{2} \right)^2} \quad (\text{EQ 9})$$




Figure 2. Energy as a function of supply voltage for an inverter chain.

1. Note that we assume that short circuit power is negligible and can be ignored. This assumption is known to hold for well-designed circuits in normal (super-threshold) operation [12]. Using SPICE simulations we have found that this assumption holds in subthreshold operation as well.




Figure 3. The ratio  $\eta$  in EQ11 with Vdd (SPICE)

It is shown in [12] that if  $t_r > t_{pHL,actual}$  (which is satisfied when an inverter drives another one of the same size, as in our modelling),

$$t_{pHL,actual} = 0.84t_r \quad (\text{EQ 10})$$

Substituting EQ10 into EQ9 gives,

$$t_{pHL,actual} = 1.2445 \cdot t_{pHL,step} \quad (\text{EQ 11})$$

Similar results hold for  $t_{pLH}$  [12]. We then can estimate the average  $t_{p,actual}$  as:

$$\begin{aligned} t_{p,actual} &= 1.2445 \cdot t_{p,step} \\ &= \eta \cdot t_{p,step} \end{aligned} \quad (\text{EQ 12})$$

However, we need to test if this linear model is valid for subthreshold operation. To justify the linear modelling of  $t_{p,actual}$  with  $t_{p,step}$  at such a wide supply voltage range, we plot the calculated  $\eta$  as a function of Vdd, based on SPICE simulation in Figure 3.

From Figure 3, it is clear that the coefficient  $\eta$  increases as the supply voltage is reduced to the subthreshold regime. Other factors affecting the accuracy are that EQ5 does not perfectly model  $I_{sub}$  in subthreshold operation<sup>1</sup> and that voltage swing degrades at ultra low supply voltages. Taking these factors into account, for the target technology we set an effective  $\eta=2.1$  for subthreshold operation.

As the supply voltage reduces, the total energy consumption reaches a minimum at some supply voltage (referred to as  $V_{min}$ ) since the delay of the circuit increases and the circuit now leaks over a larger amount of time. Substituting the equation for circuit delay EQ12 into EQ7, we obtain the following expression for total energy:

$$\begin{aligned} E &= \frac{1}{2} \cdot \alpha \cdot n \cdot C_s \cdot V_{dd}^2 + n \cdot V_{dd} \cdot I_{leak} \cdot n \cdot \frac{\eta C_s V_{dd}}{2I_{on}} \\ &= \frac{1}{2} n C_s V_{dd}^2 \cdot \left( \alpha + \eta \cdot n \cdot \frac{I_{leak}}{I_{on}} \right) \end{aligned} \quad (\text{EQ 13})$$

Note that  $I_{on}$  here is subthreshold "on" current because we are focusing on the subthreshold region where  $V_{min}$  occurs. By substituting




Figure 4. Energy-Vdd for an inverter chain ( $n=20$ )

1. We find that over the entire subthreshold region ( $0 < V_{dd} < V_{th}$ ),  $I_{sub}$  deviates from the simple exponential equation (EQ5) by at most 20% if we treat mobility  $\mu$  as constant.

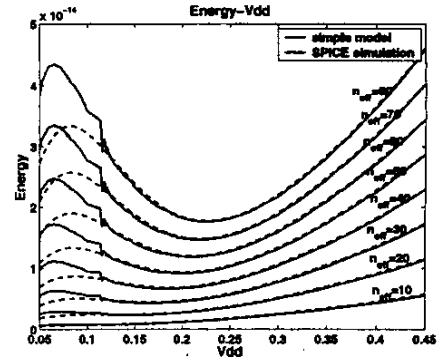


Figure 5. Inverter chain Energy-Vdd (analytical model vs. SPICE) EQ5 into EQ13, we now arrive at our final expression for the total energy as a function of supply voltage for subthreshold operation:

$$E = \frac{1}{2} n C_s V_{dd}^2 \cdot \left( \alpha + \eta \cdot n \cdot e^{\left( \frac{V_{dd}}{mV_T} \right)} \right) \quad (\text{EQ 14})$$

Based on this simple expression of total energy, we can find the optimal minimum energy voltage  $V_{min}$  by setting  $\partial E / \partial V_{dd} = 0$ . Let  $u = \eta \cdot n / \alpha$  and  $t = V_{dd} / mV_T$ , we obtain:

$$e^t = \frac{u}{2} \cdot t - u \quad (\text{EQ 15})$$

We rewrite the above equation as:

$$u = \frac{e^t}{\frac{t}{2} - 1} \quad (\text{EQ 16})$$

By doing this, we can easily find that only if  $u \geq 2e^3$  ( $t=3$ ) can  $E$  have a minimum, which means the lowest  $V_{min}$  is  $3mV_T$ . This corresponds to  $n \geq 4$  if  $\eta=2.1$ ,  $\alpha=0.2$ .

Since EQ15 is a non-linear equation, it is impossible to solve it analytically. Hence, we use curve-fitting to arrive at the following closed-form expression:

$$t = 1.587 \ln u - 2.355 \quad (\text{EQ 17})$$

Substituting the original variables gives the following final expression for the energy optimal voltage:

$$V_{min} = \left[ 1.587 \ln \left( \eta \cdot \frac{n}{\alpha} \right) - 2.355 \right] \cdot mV_T \quad (\text{EQ 18})$$

Note that in the presented model, the only parameters that are technology-dependent are  $\eta$  and  $m$ . Hence, when we switch from one technology to another, it is only required to determine these two parameters which is readily accomplished. Interestingly, the total energy in EQ14 and the optimal energy voltage  $V_{min}$  do not depend on the threshold voltage  $V_{th}$ , as verified using SPICE. This independence is caused by the fact that in subthreshold operation both leakage and delay have similar dependencies on  $V_{th}$ , and hence the effect




Figure 6. Minimal energy  $V_{min}$  with inverter effective stage number  $n_{eff}$




Figure 7. Minimal energy  $V_{min}$  with NAND2 effective stage number  $n_{eff}$  of  $V_{th}$  on the total energy cancels out. Also, we find that the minimum energy voltage is strongly dependent on the number of stages in the inverter chain. This is due to the fact that in a longer inverter chain more gates are leaking relative to the dynamic energy component, causing  $V_{min}$  to occur at a higher voltage. Finally, we point out that  $V_{min}$  is strongly related to the activity factor  $\alpha$ . In a circuit with a lower  $\alpha$ ,  $V_{min}$  occurs at a larger voltage than in a circuit with higher  $\alpha$ , because a lower  $\alpha$  gives the circuit more time to leak and effectively increases the stage number, as shown in Figure 4. We therefore introduce the notation of effective stage number as  $n_{eff} = \frac{n}{\alpha}$ .

In order to verify the accuracy of the proposed model, we compared the results from EQ14 with SPICE simulations for inverter chains of different lengths. In Figure 5, we compare the energy-Vdd relationship predicted by the proposed analytical model in the sub-threshold region with SPICE simulation results for an industrial 0.18 $\mu$ m process. The plot shows a range of effective inverter chain lengths ( $n_{eff}$ ). As shown in Figure 5, the analytical model matches SPICE well, except at voltages less than 100mV. In this region, the model tends to underestimate the rise in energy consumption due to the dramatic increase of  $\eta$  from Figure 3, resulting in a delay that is greater than expected. However, this is not a severe problem since the important region of modeling is around  $V_{min}$ , where the proposed model shows good accuracy.

In Figure 6, we compare the predicted minimum energy voltage  $V_{min}$  based on our model with that measured by SPICE simulation. In the plot, the results using the fitted closed-form expression of EQ18 are shown as well as the numerical solution of the non-linear equation in EQ15. As can be seen, both match SPICE with a high degree of accuracy for a wide range of effective inverter chain lengths  $n_{eff}$ .

We now consider the energy optimal voltage for more complex gates, such as NANDs and NORs. Based on SPICE simulations for a 2-input NAND (NAND2), we find that the minimum voltage  $V_{min}$  shifts upward compared with the inverter chain. This is caused by the fact that for a chain of NAND2 gates, the number of leaking PMOS transistors is doubled in every other gate and NMOS transistors are twice the size compared to the inverter. The capacitance increase does not affect the  $V_{min}$  because the delay and the switching energy are proportional to the loading  $C_s$ . Now we introduce  $n'_{eff,inv}$  as the equivalent stage number of a inverter chain that gives the same  $V_{min}$  as a NAND2 chain with  $n_{eff,nand2}$ . The  $n'_{eff,inv}$  proves a little smaller than twice  $n_{eff,nand2}$  due to the stack affect in the nmos transistors and a slightly larger driving ability of the pull-down nmos. We therefore compute  $n'_{eff,inv}$  value for the NAND2 chain:

$$\begin{aligned} \frac{n'_{eff,inv}}{n_{eff,nand2}} &= \frac{I_{leak,nand2}}{I_{leak,inv}} \cdot \frac{I_{on,inv}}{I_{on,nand2}} \\ &= \frac{1.91}{1.1} \\ &\approx 1.74 \end{aligned} \quad (\text{EQ 19})$$

Using this  $n'_{eff,inv}$  we obtain an accurate match between the modeled  $V_{min}$  and SPICE simulation as shown in Figure 7. Other complex gates can be modeled in a similar way by calculating an appropriate

$n'_{eff,inv}$  value.

## 4 Energy Modelling for Different Low-Power Approaches

In this section, we compare the energy efficiency of MTCMOS-based power gating and DVS. First, we define five different systems:

- $S_{basic}$ , a basic system with clock gating but without power-gating or DVS.
- $S_{mtcmos}$ , a system with the ability of power-gating during idle mode (clock gating is implied).
- $S_{dvspg}$ , a partial DVS system with power-gating ability where the minimum scalable voltage is  $V_{limit}$ , set to be  $V_{dd}/2$  in the next section.
- $S_{dvsonly}$ , a system similar to  $S_{dvspg}$  but without power-gating.
- $S_{insomniac}$ , an INSOMNIAC system with unlimited voltage scaling ability, down to the energy optimal voltage.

For  $S_{basic}$ , the energy consumption during idle mode is leakage and we can model the totally energy for workload  $E_{basic}$  as:

$$E_{basic} = P_{act,vdd} \cdot t_{on} + P_{leak} \cdot (t_{on} + t_{off}) \quad (\text{EQ 20})$$

where  $t_{on}$  is the time the processor stays busy and  $t_{off}$  is the idle time.

For  $S_{mtcmos}$ , the energy consumption is modeled as:

$$\begin{aligned} E_{mtcmos} &= P_{act,vdd} \cdot t_{on} + P_{leak} \cdot t_{on} + E_{overhead} \\ E_{overhead} &= \frac{1}{2} \cdot C_{powerrail} \cdot V_{dd}^2 + \frac{1}{2} \cdot C_{internal} \cdot V_{dd}^2 + C_{sleep} \cdot V_{dd}^2 \end{aligned} \quad (\text{EQ 21})$$

where  $E_{overhead}$  is the overhead energy when gating the power supply, the  $C_{powerrail}$  is the virtual supply rail capacitance,  $C_{internal}$  is the total internal node cap,  $C_{sleep}$  is the gate capacitance of the sleep transistors.  $E_{overhead}$  results from three sources: the power rail charge loss, the circuit internal node charge loss and sleep transistor gate charge needed to conduct power gating. Depending on whether a header or footer transistor is used in  $S_{mtcmos}$ , the processor will lose the voltage level on virtual  $V_{dd}$  or virtual ground respectively, shortly after it enters sleep mode [18]. Without loss of generality, we use the average between virtual ground and virtual  $V_{dd}$  capacitance. To consider the internal node capacitances, we assume half of internal nodes are in at 1/0. Hence, half of the internal nodes are charged high or discharged low.

According to recent research in [18][19], the virtual power rail can be restored in several cycles if the circuits is carefully designed and thus the wakeup process can be treated as instantaneous compared to the thousands/millions of cycles of useful operation.

In order to model  $S_{dvspg}$  and  $S_{dvsonly}$  we must know how a DVS system implements the voltage scaling process. There are several ways to design a DVS system [16][17]. In this paper, we assume a method similar to the IBM405LP [17], which is illustrated in Figure 8. First, we define some system constants:

- $\tau_f = 1\mu s$ , the constant time it takes to scale the frequency
- $k_v = 2mV/\mu s$ , the speed of the regulator to scale the voltage

Suppose the system scales from state ( $V_1, f_1$ ) to ( $V_2, f_2$ ), then there are two possible cases:

1. Scaling voltage down --  $V_2 \leq V_1$

For  $S_{dvsonly}$  the energy consumption for a certain application is:




Figure 8. Illustration of scaling process for a DVS system

$$E_{dvs\text{only},sd} = (P_{act,v_1} + P_{leak,v_1}) \cdot \tau_f + E_{vscale,f2,v1 \rightarrow v2} + (P_{act,v_2} + P_{leak,v_2}) \cdot t_{rest} + P_{leak,v_2} \cdot t_{idle} \quad (\text{EQ 22})$$

For  $S_{dvs\text{spg}}$ , the energy consumption is

$$E_{dvs\text{spg},sd} = (P_{act,v_1} + P_{leak,v_1}) \cdot \tau_f + E_{vscale,f2,v1 \rightarrow v2} + (P_{act,v_2} + P_{leak,v_2}) \cdot t_{rest} + \delta \cdot \left( \frac{1}{2} \cdot C_{internal} \cdot V_2^2 + C_{sleep} \cdot V_2^2 + \frac{1}{2} \cdot C_{powerrail} \cdot V_2^2 \right) \quad (\text{EQ 23})$$

In both cases,

$$\begin{aligned} \text{if } (V_{2,request} \geq V_{limit}) \quad & V_2 = V_{2,request}, t_{idle} = 0, \delta = 0 \\ \text{if } (V_{2,request} < V_{limit}) \quad & V_2 = V_{limit}, t_{idle} > 0, \delta = 1 \end{aligned}$$

where  $V_{2,request}$  is the target voltage request by the application.

2. Scaling voltage up --  $V_2 > V_1$ , for both  $S_{dvs\text{only}}$  and  $S_{dvs\text{spg}}$

$$E_{dvs,su} = (P_{act,v_2} + P_{leak,v_2}) \cdot \tau_f + E_{vscale,f1,v1 \rightarrow v2} + (P_{act,v_2} + P_{leak,v_2}) \cdot t_{rest} + \frac{1}{2} \cdot C_{powerrail} \cdot (V_2^2 - V_1^2) + \frac{1}{2} \cdot C_{internal} \cdot (V_2^2 - V_1^2) \quad (\text{EQ 24})$$

For an optimal voltage up-scaling process,  $t_{idle}$  is always zero implying that  $S_{dvs\text{only}}$  and  $S_{dvs\text{spg}}$  have the same expression of energy consumption.

A key difference between scaling up and scaling down is that scaling up involves extra energy consumption caused by voltage level change as it draws current from the power supply to charge up the level difference.

In order to make a practical comparison, we extract detailed physical parameters from an existing Alpha processor from its layout as the basis for the following discussion (shown in Table 2).

**Table 2. Physical parameters from an Alpha processor in a 0.18um technology (without caches considered)**

|   |             |
|---|-------------|
| voltage   | 1.8 volts   |
| total # of transistors                                  | 554,052     |
| total # of gates  | 39,703      |
| power rail capacitance ( $V_{dd}$ & ground)             | 447.34 pF   |
| internal node capacitance (interconnect included)       | 3519.669 pF |
| total gate capacitance                                  | 948.079 pF  |
| sleep transistor gate cap. (assumed to be 10% gate cap) | 95 pF       |


In leading-edge processors leakage power is much more substantial than that shown in the baseline 0.18um Alpha processor. Therefore, we intentionally reduce the  $V_{th}$  in this process to make the leakage power around 11% of active power, which is reasonable for modern processors.

Due to the fact that a DVS system involves a wide range of operating voltages, the regulator efficiency (usually a buck dc-dc converter) will degrade if the loading power is small due to the internal loss of regulator becoming relatively large. We developed a regulator efficiency ( $\eta_{regulator}$ ) model based on [20][21][22] to consider this effect.


## 5 Energy Optimality for Different Work Loads

As discussed earlier, the energy optimal voltage depends on both circuit and technology characteristics. At the same time, the best choice for the minimum allowed voltage for a processor depends on its workload distribution. If the workload of a processor is such that low performance levels are never or rarely required, the minimum operating voltage for energy-efficient operating will be larger than the minimum voltage  $V_{min}$  computed in Section 3. Furthermore, for MTCMOS, the length of idle periods determine whether the switching overhead can be amortized.

Hence, we studied a number of different applications running on Linux with a Transmeta Crusoe TM5600 processor with dynamic



**Figure 9. Comparison of energy consumption under different low-power schemes**



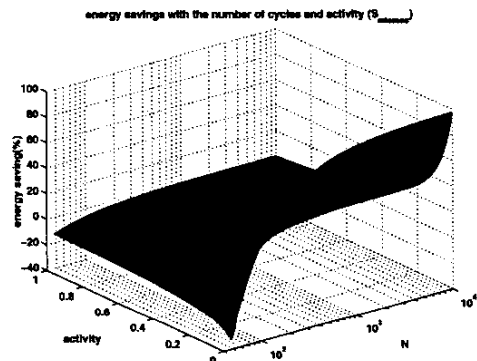
**Figure 10. Illustration of cycle number  $N$  and activity**

voltage scaling and recorded traces of the minimum necessary performance levels for each application. The applications comprise both multimedia and interactive applications:

- *emacs* is a trace of user activity using the editor performing light text editing tasks
- *konqueror* and *netscape* are traces of web browsing sessions using the two browsers
- *fs* contains a record of filesystem-intensive operations
- *mpeg* is a trace using MPEG2 video playback

To make a fair comparison, we convert these traces to fit into the Alpha processor mentioned earlier. By applying four different low-power schemes to these workloads, we can compute the energy savings relative to  $S_{basic}$  based on the models in Section 4. The results are shown in Figure 9. As the bar graph shows, the largest savings for all five applications are seen on  $S_{insom}$ . Note that for a less leaky process,  $S_{insom}$  will be even more efficient. Also, the graph shows that it is usually helpful if we apply power gating when the processor is not able to scale as needed, because for all these applications the processor runs long enough to amortize the switching overhead and thus can save significant leakage energy.

To obtain a more general evaluation for these approaches, we analyze the energy savings under artificial workload. We characterize a workload by two parameters:  $N$ , activity, where the  $N$  is the total number of cycles that the processor will run before the deadline at normal operation mode, and the activity is the ratio of the number of working cycles to  $N$ , as shown in Figure 10. Both of these two factors



**Figure 11. Energy savings with  $S_{mtcmos}$  for general workloads.**

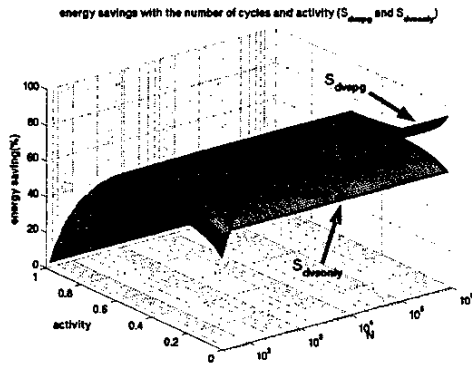


Figure 12. Energy savings with  $S_{dvspg}$  and  $S_{dvsonly}$  for general workload influence the design choice of low-power systems.

The energy savings for  $S_{mtcmos}$  with general workloads are shown in Figure 11.  $S_{mtcmos}$  is useful when activity is very low and the number of cycles is large, which gives savings close to 100% and even better savings than  $S_{dvspg}$ . This is because energy is almost completely due to leakage in the  $S_{basic}$  system. The reason that we find negative savings for  $S_{mtcmos}$  is because the extra power gating energy outweighs the potential leakage energy in  $S_{basic}$ .

Figure 12 contains the results for both  $S_{dvsonly}$  and  $S_{dvspg}$ . A system without power gating is independent of the total number of cycles, which can be easily derived from the models in Section 4. When activity is high and the application never requests a voltage less than half  $V_{dd}$ , there will be no idle period if an optimal scheduling is used. Therefore no difference exists between  $S_{dvsonly}$  and  $S_{dvspg}$  at this activity range. When activity drops below a certain value and leads to requesting voltages lower than  $V_{limin}$ ,  $S_{dvspg}$  becomes better than  $S_{dvsonly}$  because of leakage saving. This confirms again that for modern state-of-the-art partial-DVS systems, it is helpful if the system also includes power-gating to avoid unwanted leakage at run-time.

For  $S_{insom}$ , the energy savings are independent of  $N$ , but for consistency we still plot the savings in three dimensions as shown in Figure 13. It is clearly shown that  $S_{insom}$  can easily provide energy gains for a very wide range of activity, and gives much improved savings than  $S_{dvspg}$  or  $S_{mtcmos}$ . Only when activity is very low does the energy savings become saturated, which occurs since the leakage dominates the overall system energy consumption. The system has scaled below  $V_{min}$  at this point, as described in Section 3.

## 6 Conclusions

In this paper, we developed analytical models for the most energy efficient supply voltage ( $V_{min}$ ) for CMOS circuits. A number of interesting conclusions can be drawn: 1) Energy shows clear minimum in the subthreshold region since the time over which a circuit is leaking (delay) grows exponentially in this region while leakage current itself does not drop as rapidly with reduced  $V_{dd}$ , 2)  $V_{min}$  does not depend

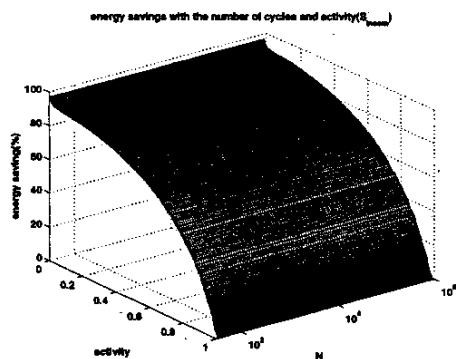


Figure 13. Energy savings with  $S_{insom}$  for general workload

on  $V_{th}$  if  $V_{min}$  is smaller than  $V_{th}$ , 3) the circuit logic depth and switching factor impacts  $V_{min}$  since it relates to the relative contributions of leakage energy and active energy and 4) the only relevant technology parameters to  $V_{min}$  are subthreshold swing and the dependency of delay on input transition time. The analytical models presented are shown to match very well with SPICE simulations.

In the second part of the paper, we compare the energy savings of different low-power schemes, namely, pure MTCMOS, partial-DVS, partial-DVS with power gating, and INSOMNIAC (extended-DVS). The comparison for five applications traces recorded on a commercial processor shows that INSOMNIAC provides the largest savings. A comparison for arbitrary workloads shows that for the majority of different application activity ratios INSOMNIAC continues to yield the largest energy improvements.

## Acknowledgements

This research was supported by ARM, NSF, SRC, GSRC-DARPA

## References

- [1] Transmeta Crusoe. <http://www.transmeta.com/>
- [2] Intel XScale. <http://www.intel.com/design/intelxscale/>
- [3] IBM PowerPC. <http://www.chips.ibm.com/products/powerpc/>
- [4] K. Flautner, S. Reinhardt, and T. Mudge, "Automatic Performance Setting for Dynamic Voltage Scaling," In Proc. of the 7th Annual International Conference on Mobile Computing and Networking (MobiCom'01), May 2001.
- [5] T. Sakurai and A. Newton, "Alpha-Power Law MOSFET Model and Its Applications to CMOS Inverter Delay and other Formulas", IEEE JSSCC, Vol. 25, No. 2, April 1990.
- [6] M. Miyazaki, J. Kao, A. Chandrakasan, "A 175mV Multiply-Accumulate Unit using an Adaptive Supply Voltage and Body Bias (ASB) Architecture", ISSCC 2002, pp. 58-59.
- [7] A. Wang, A. Chandrakasan, "A 180mV FFT Processor Using Subthreshold Circuits Techniques", ISSCC 2004, pp. 292-294.
- [8] K. Flautner and T. Mudge, "Vertigo: automatic performance-setting for Linux," In 5th Symp. Operating Systems Design & Implementation, pp. 105-116, Dec 2002.
- [9] J. D. Meindl and J. A. Davis, "The fundamental limit on binary switching energy for terascale integration (TSI)," IEEE JSSCC, vol. 35, pp. 1515-1516, Oct. 2000.
- [10] F. Møller, "Algorithm and architecture of a 1-V low-power hearing instrument DSP", ISLPED, pp. 711, Aug. 1999
- [11] BSIM3. <http://www-device.eecs.berkeley.edu/~bsim3/get.html>
- [12] J. Rabaey, "Digital Integrated Circuits: A Design Perspective", Prentice Hall, 1996.
- [13] H. Soeleman, K. Roy, "Digital CMOS logic operation in the sub-threshold region", in GVLSI, pp. 107-112, March 2000.
- [14] A. Wang, A.P. Chandrakasan and S.V. Kosonocky, "Optimal supply and threshold scaling for subthreshold CMOS circuits," IEEE Symposium on VLSI, pp. 5-9, April 2002
- [15] D. Hodges and H. Jackson, *Analysis and Design of Digital Integrated Circuits*, McGraw-Hill, 1988
- [16] "Energy Efficient Microprocessor Design", T. Burd and R. Brodersen, Kluwer Academic Publishers
- [17] K. Nowka, G. Carpenter, et al, "A 32-bit PowerPC System-on-a-Chip With Support for Dynamic Voltage Scaling and Dynamic Frequency Scaling", JSSCC, pp. 1441-1447, vol. 37, Nov. 2002
- [18] J. Tschanz, S. Narendra, Y. Ye, et al, "Dynamic Sleep Transistor and Body Bias for Active Leakage Power Control of Microprocessors", JSSCC, pp.1838-1845, vol. 38, Nov. 2003
- [19] S. Kim, S. Kosonocky, D. Knebel, "Understanding and Minimizing Ground Bounce during Mode Transition of Power Gating Structures", ISLPED pp. 22-25, Aug. 2003.
- [20] R. Erickson and D. Maksimovic, "High Efficiency DC-DC Converters for Battery-Operated Systems with Energy Management," Worldwide Wireless Communications, Annual Reviews on Telecommunications, 1995.
- [21] A. Dancy and A. Chandrakasan, "Techniques for aggressive supply voltage scaling and efficient regulation," in Proc. IEEE Custom Integrated Circuits Conf. (CICC), 1997, pp. 579-586
- [22] J. Kim and M. A. Horowitz, "An Efficient Digital Sliding Controller for Adaptive Power-Supply Regulation," in JSSCC, May 2002, vol. 37, pp. 639-647