Theoretical and Practical Limits of Dynamic Voltage Scaling

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 Vdd. However, it is possible to construct designs that operate over a much larger voltage range: from full Vdd 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 use the proposed model to study the workload activity of an actual processor and analyze the energy efficiency as a function of the lower limit of voltage scaling. Based on this study, we show that extending the voltage range below 1/2 Vdd will improve the energy efficiency for most processor designs, while extending this range to subthreshold operation is beneficial only for very specific applications. Finally, we show that operation deep in the subthreshold voltage range is never energy-efficient.

Categories and Subject Descriptors

B.8.2 [Performance and Reliability]: Performance analysis

General Terms performance, design, reliability

Keywords dynamic voltage scaling, minimum energy point

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 which are battery operated. Dynamic voltage scaling (DVS) was proposed as an effective approach to reduce energy use and is now utilized 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).

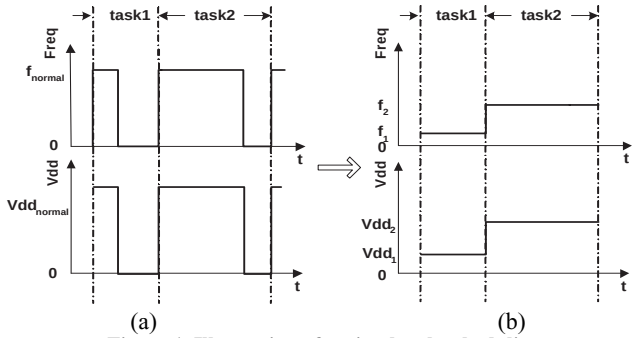

Figure 1. Illustration of optimal task scheduling

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
DAC'04, June 7-11, 2004, San Diego, California, USA
Copyright 2004 ACM 1-58113-828-8/04/0006...$5.00.

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[1], 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.

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

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 Vdd to approximately half Vdd at most. In Table 1, the available range of operating voltages and associated performance levels are shown for four commercial designs. The lower limit of voltage scaling 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 then two hundred mV. 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 ideal minimum allowable supply voltage of a functional CMOS inverter is 36mV [9]. A number of commercial products have also used subthreshold operation for extremely low power applications [10].

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

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 Vdd deviates as Vdd 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 Vdd, causing the delay to increase exponentially with voltage scaling. Since leakage

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.

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 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 the lower limit of voltage scaling is 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. Moreover, the energy/voltage relationship flattens out as the operating voltage approaches the theoretical lower limit of voltage scaling. Therefore, if the applications use low performance levels only infrequently, the gain in energy savings from extending the operating voltage range is limited. To analyze this trade-off, we study a number of workload traces obtained from a processor running a wide range of applications. Using our energy model, we investigate the trade-off between the energy efficiency of the processor and the lower limit of voltage scaling. Our results show that most applications benefit significantly from an operating voltage range that is greater than what is available in most current DVS processors, but true subthreshold operation is not required. On the other hand, applications that spend extensive time in near idle mode will benefit significantly from a voltage scaling ability from full to subthreshold voltages.

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 and discusses extensions of our model to larger circuits. Section 4 present our analysis of workload data and the practical trade-off between the minimum scaling voltage and energy efficiency. Finally, Section 5 contains our conclusions.

2 Circuit Behavior at Ultra Low Voltages

Before we derive the energy optimal operating voltage in Section 3, in this section 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:

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

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:

$$V_{dd,limit} = 2 \cdot V_T \cdot \ln\left(1 + \frac{S_s}{\ln 10 \cdot V_T}\right) \cong 52mV \cdot \ln\left(1 + \frac{S_s}{59.87mV}\right) \text{ at 300K} \quad \text{(EQ 3)}$$

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. In Figure 2, we also show that it is possible to operate a wide range of standard library gates at similar operating voltages and that their delay tracks relatively well to that of the inverter. It is, however, clear that there are practical reasons why operating circuits at the minimum voltage is not desirable, such as susceptibility to noise and process variations[15]. More importantly, we show in the next section that from an energy efficiency point of view, the minimum operating voltage for functionally correct operation does not provide the best results.

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 3. 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 3.

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. But the switching activity varies in different circuits, so we include the input activity factor α, which is the average number of times the 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 α. The standard expression for subthreshold current is given by[11]:

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

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 which is a typical value. We now express the total energy E per clock cycle as the sum of dynamic, leakage energy[1]:

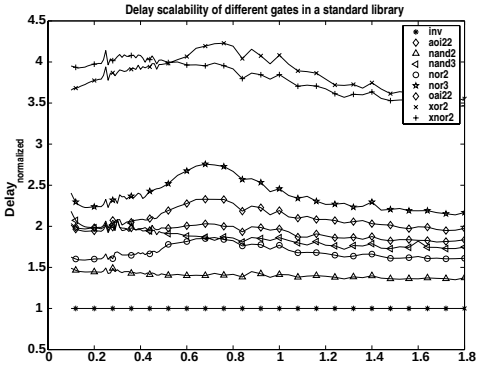

Figure 2. Delay of typical library gates over a wide voltage range, normalized to inverter delay

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 [13]. Using SPICE simulations we have found that this assumption holds in subthreshold operation as well.

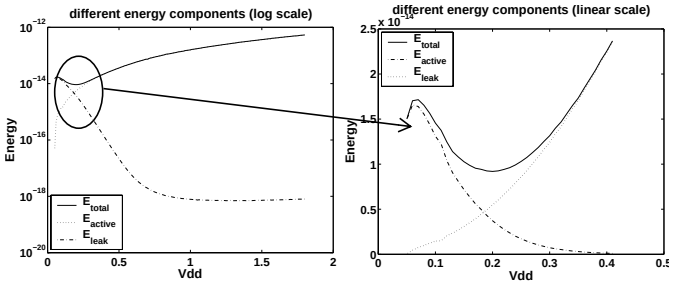

Figure 3. energy as a function of supply voltage.

$$E \cong E_{active} + E_{leak} = \alpha \cdot n \cdot E_{switch,inv} + P_{leak} \cdot t_d = \alpha \cdot n \cdot \left(\frac{1}{2} \cdot C_s \cdot V_{dd}^2\right) + (n \cdot V_{dd} \cdot I_{leak}) \cdot (n \cdot t_p) \quad \text{(EQ 7)}$$

where

- a - activity 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 [18],

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

It is shown in [13] that if $t_r > t_{pHL,actual}$ (which is satisfied when an inverter drives another one of the same size, as in our modelling),

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

Substituting EQ10 into EQ9 gives,

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

Similar results hold for t_{pLH} [13]. We then can estimate the average $t_{p,actual}$ as:

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

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 η as a function of Vdd, based on SPICE simulation in Figure 4.

From Figure 4, it is clear that the coefficient η 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[1] and that voltage swing degrades at ultra low supply voltages. Taking these factors into account, we set for this technology an effective η=2.1 for subthreshold operation.

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 μ as constant.

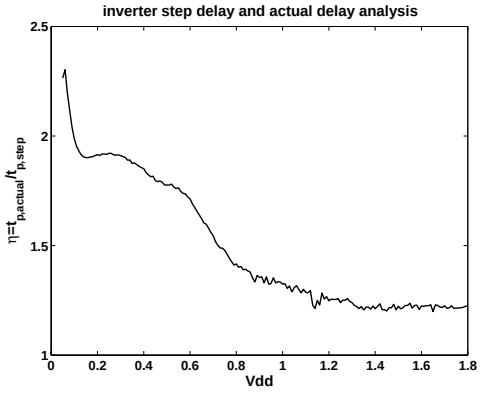

Figure 4. The ratio η in EQ11 with Vdd (SPICE)

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:

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

Note that I_{on} here is subthreshold "on" current because we are focusing on subthreshold region where V_{min} occurs. By substituting EQ5 into EQ13, we now arrive at our final expression for the total energy as a function of supply voltage for subthreshold operation:

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

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 $n \cong 4$ if η=2.1, α=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:

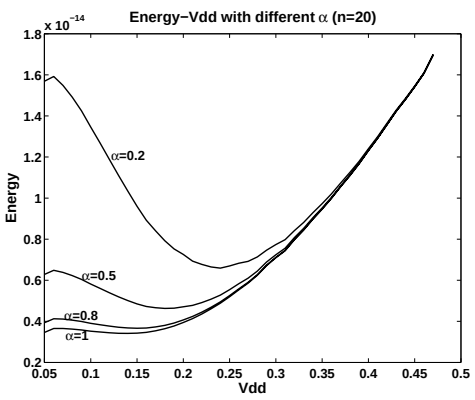

Figure 5. Energy-Vdd for an inverter chain(n=20)

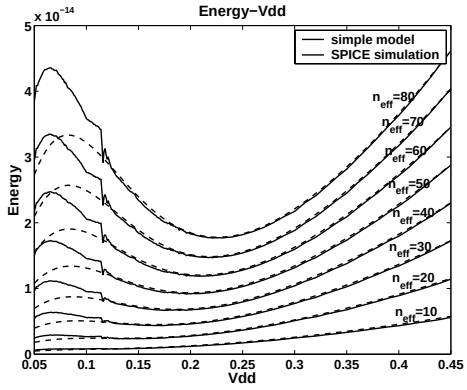

Figure 6. Inverter chain Energy-Vdd (analytical model vs. SPICE)

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

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

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

Note that in the presented model, the only parameters that are technology-dependent are η and m. Hence, when we switch from one technology to another, it is only required to determine these two parameters which can be easily 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 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 α. In a circuit with a lower α, V_{min} occurs at a larger voltage than in a circuit with higher α, because a lower α gives the circuit more time to leak and effectively increases the stage number, as shown in Figure 5. We therefore introduce the notation of effective stage number as $n_{eff} = \frac{n}{\alpha}$ to be used in the following analysis.

4 Model Verification and Extension to Circuit Blocks

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 6, we compare the energy-Vdd relationship predicted by the proposed analytical model in the subthreshold region with SPICE simulation results for an industrial 0.18um process. The plot shows a range of effective inverter chain lengths (n_{eff}). As shown in Figure 6, the analytical model matches SPICE well, except at voltages less than 100mV. In this region, the

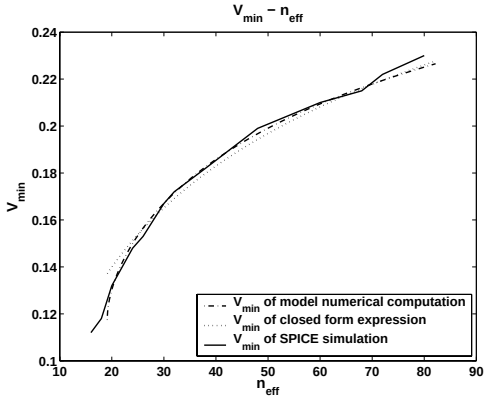

Figure 7. Minimal energy V_{min} with inverter effective stage number n_{eff}

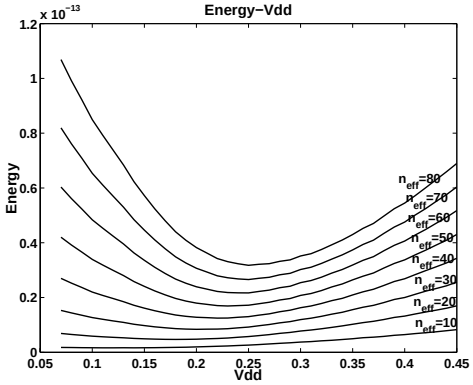

Figure 8. NAND2 chain Energy-Vdd (SPICE)

model tends to underestimate the rise in energy consumption due to the dramatic increase of η from Figure 4, 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 7, 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 NAND and NOR, as well as larger circuit blocks. Figure 8 shows results of SPICE simulations for a NAND2. As can be seen, the minimum voltage V_{min} shifts right compared with the inverter chain which means that the energy optimal voltage occurs at a higher voltage. 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. 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:

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

Using this $n'_{eff,inv}$, we obtain an accurate match between the modeled V_{min} and SPICE simulation as shown in Figure 9. Other complex gates can be modeled in a similar way by contributing to each an

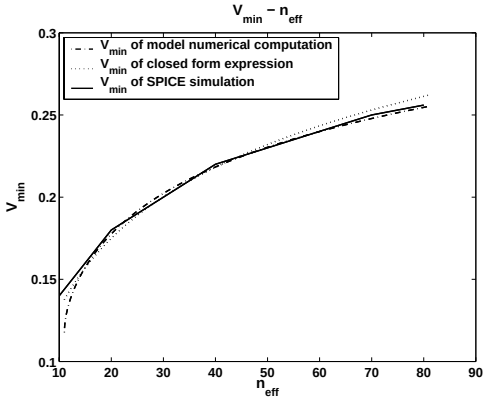

Figure 9. Minimal energy V_{min} with NAND2 effective stage number n_{eff}

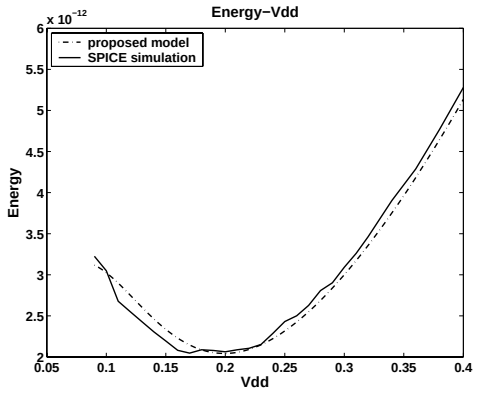

Figure 10. Energy - Vdd for 16X16 multiplier circuit

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

This approach can be extended to larger circuit blocks as well. In Figure 10, we show the total energy as a function of supply voltage obtained using SPICE for 16 x 16 multiplier when activity factor α=0.5. We estimate the total power consumption for large circuit blocks such as this by extending the expression in EQ14 as follows:

$$E_{total} = E_{active} + E_{leak} \quad \text{(EQ 20)}$$

$$E_{act} = \alpha \cdot S_{HD} \cdot C_{w0} \cdot W_{total} \cdot V_{dd}^{2} \quad \text{(EQ 21)}$$

where S_{HD} is the switching factor to model the hamming distance of the inputs[21], W_{total} is the total width of all the transistors in the circuit, C_{w0} is the capacitance of a unit width transistor. We compute the total leakage energy as follows:

$$\begin{aligned} E_{leak} &= I_{leak,total} \cdot V_{dd} \cdot t_d \\ &= (\gamma_{leak} \cdot W_{total} \cdot I_{leak0}) \cdot V_{dd} \cdot (n_{depth} \cdot t_{p,FO4}) \end{aligned} \quad \text{(EQ 22)}$$

where γ_{leak} is the leaking factor used to model the leakage stack effect and input pattern dependency, I_{leak0} is the leak current of a unit width transistor, n_{depth} is the logic depth in terms of fanout-of-four (FO4) inverter delay $t_{p,FO4}$, which is expressed as follows:

$$t_{p,FO4} = \frac{\frac{1}{2} \cdot (4 W_{inv} \cdot C_{w0}) \cdot V_{dd}}{W_{inv} \cdot I_{on0}} \quad \text{(EQ 23)}$$

where I_{on0} is the on-current a unit width inverter. Note that S_{switch} may change with supply voltage as glitches are sensitive to circuit delay although for simplicity we treat S_{HD} as a constant. Substituting EQ21 and EQ22 into EQ20, we can derive the following expression for total energy of a circuit block as a function of supply voltage in a manner similar to EQ14:

$$E_{total} = C_{w0} W_{total} V_{dd}^{2} \left(\alpha S_{HD} + 2\gamma_{leak} n_{depth} e^{\left(-\frac{V_{dd}}{mV_T}\right)} \right) \quad \text{(EQ 24)}$$

For the test circuit in Figure 10, the following parameters for the model were found using SPICE simulation: $S_{HD} \cong 0.55$, $\gamma_{leak} \cong 0.5$, $n_{depth} \cong 65$. The total energy predicted by EQ24 with above parameters is shown in Figure 10 for the 16x16 multiplier block together with SPICE simulation results.

It is important to note that for a generic circuit block n_{eff} is defined as $n_{eff,block} \equiv \frac{n_{depth}}{\alpha \cdot S_{HD}}$. Therefore when the activity factor α and switching factor S_{HD} are very low, based on circuit structure or the input data stream, the $n_{eff,block}$ is actually much larger than the real logic depth n_{depth}. In a real processor, the activity factor varies across the chip because not all the circuit blocks are working intensively at all times. Therefore, in order to gain energy efficiency, designers must take into account the α difference before estimating the average V_{min}. In other words, for the purposes of optimizing DVS, low activity and large logic depths are interchangeable as they both lead more quickly to leakage dominated designs.

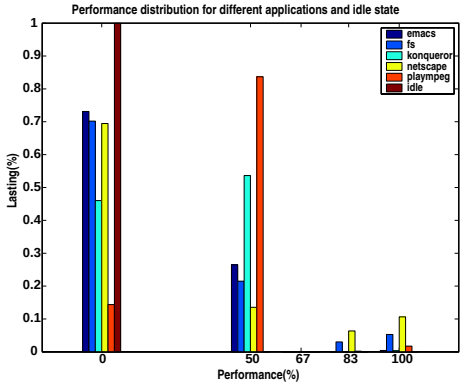

Figure 11. Performance distribution of different applications

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. Hence, we studied a number of different applications running on Linux using an ARM926 and Transmeta Crusoe TM5600 processors with dynamic 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
- *idle* traces the activity when the system has no dominant workloads and as a result contains very little activity and mostly operating system housekeeping tasks.

The dynamic performance management policy is based on Vertigo [8] and ARM's Intelligent Energy Manager. The distribution of the four available performance levels (with a highest frequency of 600MHz) among the executed tasks is shown in Figure 11 for each application. As the bar graph shows, the processor spends significant time in sleep mode, meaning that the processor completes many tasks well ahead of schedule. Most importantly, we observed that during the execution of all tasks a run-then-idle pattern was seen 50% of the time. This implies that many tasks could run at a frequency less than the minimum (50%) available on the processor if it was able to do so.

By extending the lower limit of voltage scaling, the amount of idle time can be reduced leading to more energy-efficient operation. Based on the previous analysis, energy efficiency can increase until it reaches the energy optimal voltage V_{min}. In addition, by eliminating the need to enter a sleep state, any energy overhead due to switching to and from sleep mode is also avoided, further increasing the energy efficiency.

We therefore study the total energy consumption of the processor as a function of the lower limit of the performance that the processor provides, denoted by f_{limit}. Assuming that we have an ideal performance scheduler that is able to set the performance exactly sufficient to just complete every task, we can compute the optimal energy consumption with different f_{limit} values. The total energy is based on the proposed energy model of Section 3 for subthreshold voltage opera-

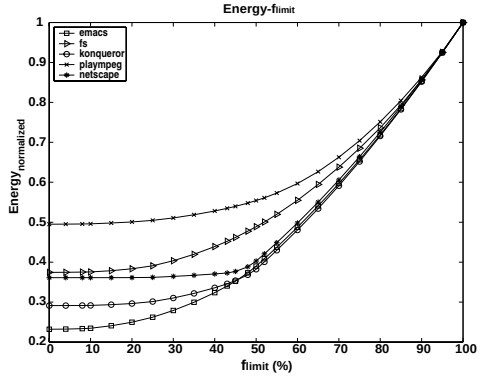

Figure 12. Energy - f_{limit} for different applications

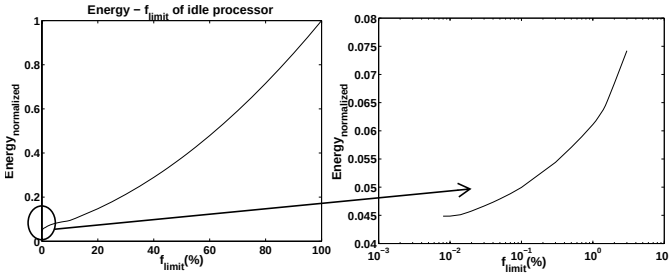

Figure 13. Energy - f_{limit} for an idle processor

tion, combined with a simple fitted model for energy and delay at super-threshold operating voltages. Note that we do not consider the sleep-wakeup energy overhead although this could be easily incorporated in our analysis. We show the energy / f_{limit} trade-off for the first five applications in Figure 12. As can be seen, the energy efficiency improves as the f_{limit} is reduced and levels off for most applications below 10%, which corresponds to a V_{dd}/V_{dd0} of 30.7% (553mV for a V_{dd0} of 1.8V).

Finally, we also analyze the energy / f_{limit} trade-off for the idle-mode trace, in which the processor is mostly in sleep mode, waking up only to do regular "housekeeping" chores for the operating system. Note that this state can be quite common on a processor. The results are given in Figure 13, and show that the energy continues to reduce down to a performance level of 0.02%, corresponding to a V_{dd}/V_{dd0} of 13% (234mV for a V_{dd0} of 1.8V). Note that in such low activity situations the *practical* V_{min} value approaches the theoretical V_{min} levels of Section 3. The energy savings of a more scalable processor over the traditional one are summarized in Table 2, and how that substantial energy savings can be obtained by extending the voltage range appropriately.

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 on V_{th}, 3) the logic depth and switching factor of the circuit 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.

We then used these models along with workload traces for an existing DVS processor to show how the *practical* minimum energy voltage compares to the theoretical V_{min} value. We find that under typical workload requirements, the operating voltage (frequency) should be scaled to approximately 30% (10%) of the maximum. Since in current DVS-based processors V_{min} is commonly 50% of the maximum, this implies that there is room for improvement in the energy efficiency of these systems.

Table 2. Energy consumption comparison between aggressive DVS and traditional DVS approaches

Application	Normalized Energy		Energy Savings
	Aggressive DVS	Traditional DVS	
emacs	0.235	0.324	21.7%
fs	0.376	0.439	10.5%
konqueror	0.292	0.336	9.32%
netscape	0.361	0.370	1.62%
mpeg	0.496	0.528	4.38%
idle state	1.76E-4	9.63E-4	81.7%

note: In aggressive DVS, V_{dd}/V_{dd0} is 30.7% for general applications, 13% for idle state; in traditional DVS, V_{dd}/V_{dd0} is assumed as 50%.

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] H. Soeleman, K. Roy and B. Paul, "*Robust ultra-low power subthreshold DTMOS logic*," in ISLPED, pp. 25-30, 2000.

[13] J. Rabaey, "*Digital Integrated Circuits: A Design Perspective*", Prentice Hall, 1996.

[14] H. Soeleman and K. Roy, "*Ultra-low Power Digital Subthreshold Logic Circuits*", in ISLPED, pp. 94-96, 1999.

[15] H. Soeleman, K. Roy, "*Digital CMOS logic operation in the sub-threshold region*", in GVLSI, pp. 107-112, March 2000.

[16] A. Forestier and M.R. Stan, "*Limits to voltage scaling from the low power perspective*," in 13th Symposium on Integrated Circuits and Systems Design, pp. 365-370, Sept. 2000.

[17] 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

[18] D. Hodges and H. Jackson, *Analysis and Design of Digital Integrated Circuits*, McGraw-Hill, 1988

[19] F. Brglez and H. Fujiwara, "*A Neutral Netlist of 10 Combinational Circuits and a Target Translator in Fortran*", Proc. IEEE ISCAS, pp. 663-698, June 85.

[20] T.D. Burd and R.W. Brodersen, "*Design issues for Dynamic Voltage Scaling*", ISLPED, pp. 9-14, 2000

[21] *"Power Aware Computing"*, edited by R. Graybill and R. Melhem, Kluwer Academic/Plenum Publishers, May 2002