Techniques for Aggressive Supply Voltage Scaling and Efficient Regulation

Abram Dancy Anantha Chandrakasan
Department of EECS,
Massachusetts Institute of Technology, Cambridge

ABSTRACT

Aggressive voltage scaling to 1V and below through technology, circuit, and architecture optimization is the key to low-power design. Threshold voltage scaling enables aggressive supply scaling but increases leakage power. Technology and circuit trends to control idle leakage power are presented including MTCMOS, variable V_T bulk-CMOS, and variable V_T SOI. Power can also be reduced by adaptively varying the supply voltage in applications where the computational workload varies with time. Aggressive voltage and power level scaling requires efficient DC-DC conversion circuitry and in some cases, it is necessary to embed this function in the processor.

1. Analysis of Power Components

There are three main sources of power consumption in digital CMOS circuits: switching power, short circuit power, and leakage power. In conventional process technology, the switching component dominates and is given by $\alpha \cdot C_L \cdot V_{DD} \cdot \Delta V \cdot f_{clk}$ where α is the node transition activity factor, C_L is load capacitance, V_{DD} is the supply voltage, ΔV is the voltage swing, and f_{clk} is the clock frequency.

The activity factor (α) is a strong function of signal statistics and circuit topology. A variety of techniques have been proposed to reduce switching activity ranging from low level logic restructuring and power down techniques, to architectural restructuring and selection of data representation [1]. It is important to consider transitions that are fundamentally required to perform a given operation as well as spurious transitions arising from imbalances in the signal path. Various approaches have been proposed to reduce glitching transitions in arithmetic structures by balancing signal paths (e.g., [2]).

In conventional CMOS logic with rail-to-rail swing, the most efficient approach to lower energy consumption is to operate at the lowest possible power supply voltage. The individual circuit elements, however, run slower at lower supply voltages and circuit performance degrades. One approach to maintain throughput at reduced voltages is to use parallel architectures to compensate for increased gate delays [1]. Significant power reduction over conventional approaches is possible at the cost of increased silicon area. A variety of voltage scaling strategies are described here which reduce power without significantly increasing silicon area. Associated with aggressive voltage and power scaling is the need for high-efficiency regulation techniques. In many cases, embedding the power converter control in the processor can significantly reduce power dissipation. The switching power can also be reduced at a fixed supply voltage by using low-swing signalling. One approach involves reducing the swing on high capacitance nodes using reference voltages generated on-chip [3]. Charge sharing techniques can also be used significantly reduce the voltage swing on data busses [4].

The short-circuit component arises when both the NMOS and PMOS transistors are "ON" simultaneously, providing a direct path from V_{DD} to ground. By sizing transistors such that the input and output rise-times are approximately equal, the short circuit component can be kept to less than 10% of the total power [5].

Leakage power results from reverse biased diode conduction and subthreshold operation. The sub-threshold leakage occurs due to carrier diffusion when the gate-source voltage, V_{GS}, has exceeded the weak inversion point, but is still below the threshold voltage V_T, where carrier drift is dominant. In this regime, the current is exponentially dependent on the gate-source voltage V_{GS}. Scaling supply voltages below 1V requires the scaling of the threshold voltage, which unfortunately comes at the cost of increased leakage (Figure 1). While leakage is typically negligible when circuits are active, it can be significant during idle mode. A variety of technology and circuit solutions are discussed here which address the conflicting requirements of high-performance during

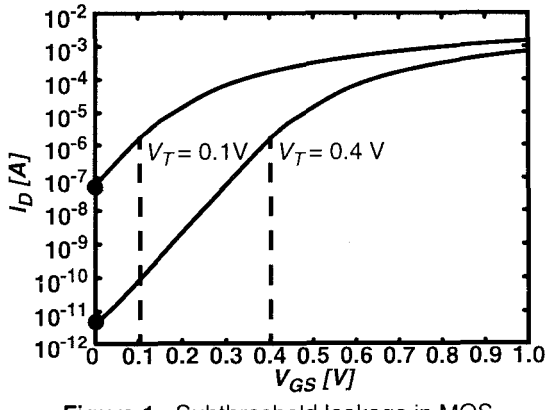

Figure 1. Subthreshold leakage in MOS.

0-7803-3669-0 $5.00 © 1997 IEEE

active periods (favoring low thresholds) and low leakage during standby (favoring high thresholds).

2. Adaptive Supply Voltage Scaling

There are many systems where the amount of processing per input sample (henceforth called computational workload or simply workload) varies with time. For example, in differential video compression algorithms, the amount of computation performed depends highly on the frame to frame image correlation. Typically each frame is divided into small blocks and only those blocks that change more than some threshold are processed. A scene change requires the entire frame to be coded while a highly correlated sequence requires only a few blocks to be coded. Figure 2 shows a histogram of the number of blocks that changed in a sequence of MPEG2 video. The wide variation in workload can be exploited to save power by shutting down the processor using gated clocks when the workload is less than the peak; this results in a linear reduction in power dissipation.

However, in such variable workload systems, power can be further reduced if a variable power supply is used in conjunction with a variable switching speed processor. Circuits with a fixed supply voltage work at a fixed speed and idle if the data sample requires less than the peak computation. Less power would be consumed if supply voltage were lowered when the workload decreased. Figure 3 shows a plot of power vs. normalized workload ($V_{DD(max)}$=2V and V_T=0.4) [6]. The straight line represents a fixed power supply system. If the workload is less than peak workload (i.e., *workload=1*), then the processor would compute as fast as possible at a fixed voltage and idle for the remaining fraction *(1-workload)* of the sample period. Power is reduced in a linear fashion since the energy per operation is fixed, and only the number of operations changes.

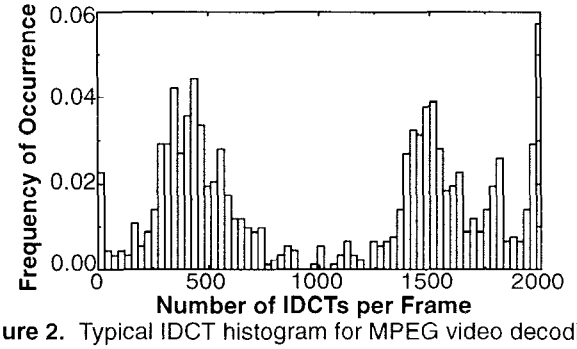

Figure 2. Typical IDCT histogram for MPEG video decoding

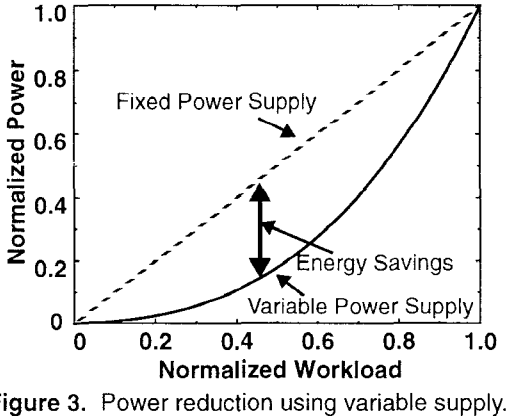

Figure 3. Power reduction using variable supply.

The lower curve of Figure 3 represents the case for a variable power supply system. If the workload for a given sample period is less than peak, then the delay of the processing element can be increased by a factor of *1/workload* without loss in throughput, allowing the processor to operate at a lower power supply voltage. The energy per sample therefore varies not only because fewer operations are performed, but also because each operation consumes less energy.

Self-timed circuits have been proposed to exploit variable processing times to dynamically vary the supply voltage [7]. Self-timed circuits implicitly account for variations in bit-level data dependent delay, algorithmic workload, temperature, and process parameters. Unfortunately, implementation requires the use of dual-rail coding, which comes at the cost of high switching activity. An alternative approach for implementing a variable supply system is to use standard *synchronous* design techniques with a ring-oscillator based clock and switching supply [6]. While this involves giving up the ability to exploit bit-level data dependent delay variations (e.g., data dependent variable ripple length of an adder), algorithmic workload variations can be exploited. For example, the supply voltage can be adapted based on computation requirements of different applications running on a general purpose processor [8].

The supply voltage is adjusted by changing the duty cycle of the switching supply based on the workload derived from the input signal (the details of efficient switching regulator design are described later). The clock frequency is controlled by a digital phase-locked loop with the variable V_{DD} as the control signal for the ring-oscillator. A lookup table (in the rate controller) records the voltage needed for each processing rate, and the feedback

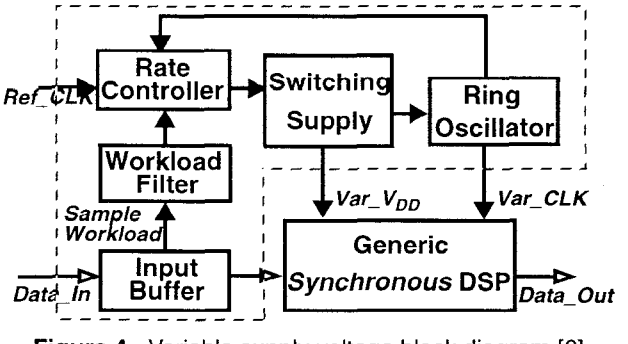

Figure 4. Variable supply voltage block diagram [6].

slowly adjusts the entries in the table as needed. Thus the loop tracks process and temperature, but computation rate is changed open-loop by reading table entries.Even when workload is fixed, the supply can be minimized by tracking process and temperature variations [9],[10].

3. Threshold Voltage Scaling

Reducing the threshold voltage (V_T) of the devices allows the supply voltage to be scaled down without loss in performance. Figure 5 shows an experimentally obtained plot of V_{DD} vs. V_T while keeping performance (gate delay) constant. These experimental plots were obtained from ring oscillator structures by adjusting the V_T (using the variable V_T SOI device described later) and V_{DD} for a fixed delay. Figure 6 shows a plot of energy vs. V_T for two different speeds of a 101 stage ring oscillator (i.e., all points on each curve have a fixed delay). Here, the power supply voltage is allowed to vary to keep the performance fixed.

For a fixed delay, the supply voltage and therefore the switching component of power can be reduced while reducing the threshold voltage. However, at some point, the threshold voltage and supply reduction is offset by an increase in the leakage current, resulting in an optimum threshold voltage and power supply voltage. That is, the optimum threshold voltage must compromise between lower switching energy and higher sub-threshold leakage

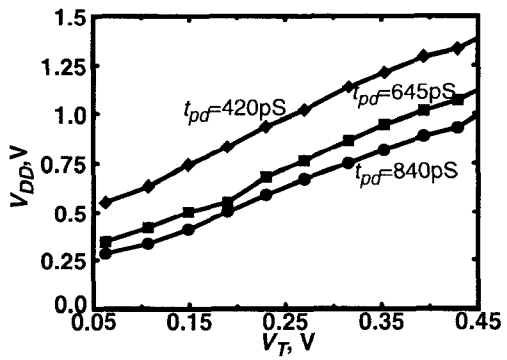

Figure 5: Experimental V_{DD} vs. V_T for a fixed delay.

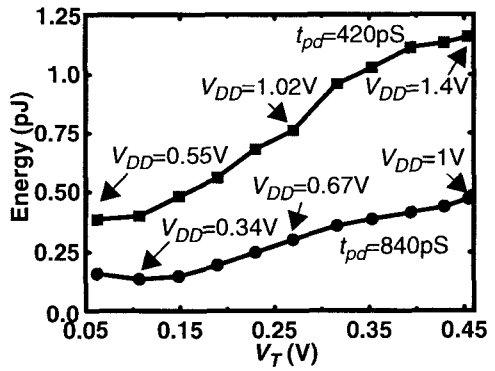

Figure 6. Experimentally derived optimum V_{DD}/V_T point using the dual-gate SOI technology described in [12].

energy at low supply voltages. It is interesting to note that the optimum supply voltage is significantly lower than 1V, even with a very low activity factor. Circuits with supply voltages as low as 200mV have been demonstrated [11].

4. Leakage Control in Burst Mode Applications

The previous section addressed supply voltage scaling via of threshold voltage scaling for continuous mode circuits. However, not all computations are continuously operational. An important class of high-performance computation is "event-driven" computation in which intermittent computation activity triggered by external events is separated by long periods of inactivity - examples include X-servers, communication interfaces, etc. An obvious mechanism for saving energy is to shut down parts of the system hardware that are idle because they are waiting for I/O from outside the system or from other parts of the system. For example, analyzing several traces obtained from real X-sessions indicates that the processor spends more than 95% of its time in the *off* state suggesting large energy reductions are achievable under ideal shutdown conditions [13].

While the processor is in its shutdown state, the system should ideally consume nearly zero power. This is only possible if the device leakage current is low - i.e., the devices have a high threshold voltage. However, for low voltage high-performance operation, reduced threshold devices are required. To satisfy the contradicting requirements of high-performance during active periods and low-standby leakage, several device technologies and circuit techniques have recently been introduced.

One approach to dynamically control leakage currents is to use a multiple threshold CMOS process (Figure 7). The logic circuits are implemented using low threshold devices and the low-V_T transistors are "gated" using series connected high threshold switches. During idle periods, the high threshold devices are cut-off, significantly reducing the subthreshold leakage. During active periods, the high threshold switches are turned on and circuits resume normal low threshold high speed opera-

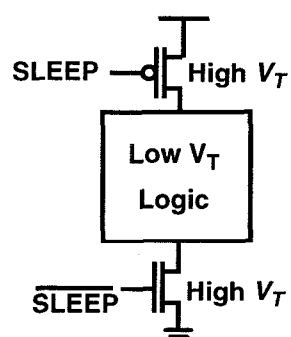

Figure 7. Multiple Threshold CMOS ([14], [15]).

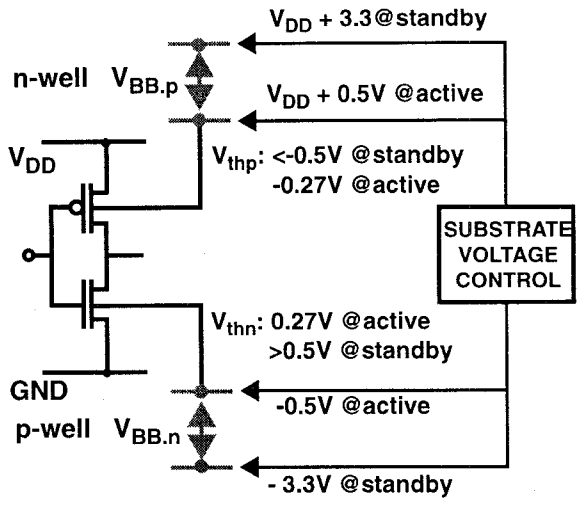

Figure 8. Adaptive V_T control using triple-well CMOS [16].

tion, provided that the devices are sized properly. Multiple threshold technology has been applied to both logic and memory. A key challenge in designing with MTCMOS is transistor sizing of high V_T devices which is very strongly dependent on data dependent current profiles. Also, preserving state requires additional overhead circuitry [15].

Another proposal for idle mode leakage control involves changing the threshold voltage dynamically by changing the substrate bias - that is, use a lower threshold voltage to operate at low power supply voltages during active periods and raise the threshold voltage during idle periods to lower leakage power [16]. One potential problem with this approach is that the threshold voltage changes in a square root fashion with respect to *source to bulk voltage* and therefore a large voltage may be required to change V_T by a few hundred mV. A high body factor is required which comes at the cost of increased parasitic capacitance.

A third approach is to use a silicon-on-insulator-with-active-substrate (SOIAS), to achieve dynamically variable threshold voltages [12]. Figure 9 shows the cross section of the SOIAS device. The back gate controls the threshold voltage (V_T) of the front gate device since the surface potentials at the front and back interfaces are coupled in fully depleted SOI devices. Leakage power and on-currents are controlled by lowering V_T when a circuit is active and raising V_T when the circuit is idle. This addresses the opposing requirements of high performance and low power, particularly at low power supply voltages. A 250mV change in threshold voltage results in a 3.5-4 decade reduction in off current and an 80% switching current increase at 1V operation [12].

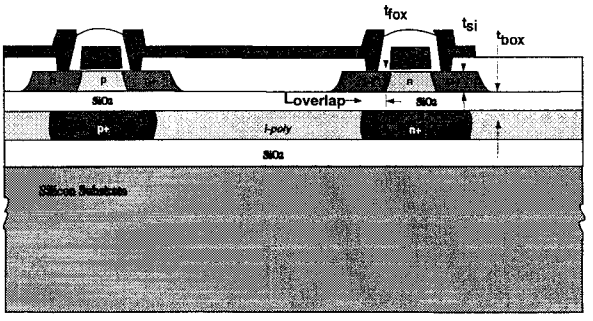

Figure 9: Silicon-on-Insulator Active Substrate [12].

5. Low Voltage Low Power Switching Regulator Design

When voltage scaling is used to lower the power consumption of a system, the generation of the low voltage supply becomes a consideration. Even if a battery cell voltage is near the desired output voltage, droop in the battery voltage over the battery lifetime may preclude its use without regulation. Thus, low power, low voltage systems create a need for efficient DC-DC converters at output power and voltage levels previously uncommon for such circuits. The most common application of a DC-DC converter for a voltage scaled system is a "down converter" (Figure 10). The converter operates by creating a pulse width modulated signal with duty cycle *D* at the node labeled V_x, whose average value is equal to the desired output voltage. Lossless passive filtering is used to filter the PWM signal, creating a DC voltage with some tolerable value of ripple.

In its simplest incarnation, a down converter requires a semiconductor switch for S1, and a diode for S2. Low voltage applications necessitate the use of a synchronous rectifier for switch S2 in order to maintain reasonable efficiencies. (Synchronous rectification is the use of an active switch in place of a rectifier.) Typical voltage and current waveforms are shown for such a converter in Figure 11. There are two fundamental modes of opera-

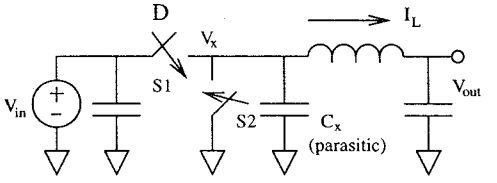

Figure 10: Standard Down Converter Topology.

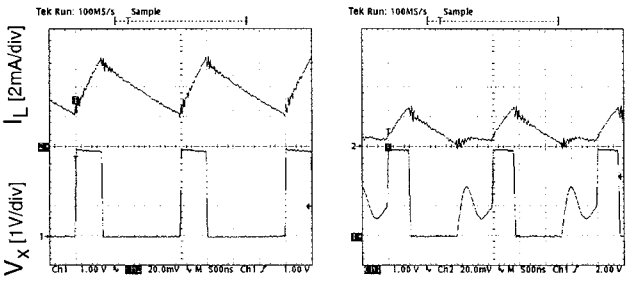

Figure 11: Voltage and current scope waveforms.

tion. In the continuous mode, either S1 or S2 is conducting at all times, and the output voltage is related to the input voltage by the expression $V_{out} = DV_{in}$. In the discontinuous mode, the inductor current goes to zero at some point, and S2 turns off. A voltage ring results on node V_x, the frequency of which is given by the inductance and the parasitic capacitance at that node. The output voltage is no longer a simple function of the input voltage and duty cycle; in discontinuous mode, the relationship depends on the load, as shown in Equation 1,

$$V_{out} = \frac{V_{in}^2}{\frac{2L}{D^2T}I_{out} + V_{in}} \quad (1)$$

where I_{out} is the average output current, L is the inductance in the filter, and T is the switching period. Both continuous and discontinuous mode operation can occur in the same system if the load current varies enough.

5.1 Frequency Selection

The switching frequency of a converter, the physical size of the output filter, and the efficiency of the converter are inextricably linked. The volume of the output filter is roughly proportional to the energy which it must store over a single cycle, which in turn is proportional to the power being processed times the period of a single cycle. The relationship between the cut-off frequency of the output filter and the switching frequency determines the size of the ripple on the output voltage. The power dissipation in a switching converter will always increase with increasing switching frequency. Choosing the switching frequency requires making tradeoffs between efficiency, power density, and transient performance.

When designing for maximum efficiency, the general statement can be made that for a fixed available filter volume, the filter should be designed so as to accommodate the lowest possible switching frequency. If the filter is designed in this way, the task of trading off the efficiency and power density is transformed into a problem of determining the maximum available filter volume, or making a tradeoff between the filter volume and the battery volume.

For a battery powered system, tradeoffs can be made between the volume of the converter and the volume of the battery. Reducing the volume of the battery in order to accommodate a power converter can increase the total system run time due to the advantages offered by voltage regulation. Since the efficiency of a DC-DC converter improves with increasing volume, battery volume and converter volume can be balanced to provide the longest possible run time, per Equation 2 [17],

$$\frac{P_{system} + P_{diss}(V_{supply})}{V_{battery}} = -\left(\frac{dP_{diss}}{dV_{supply}}\right), \quad (2)$$

where P_{system} is the dissipation in the load, P_{diss} is the loss in the converter (and a function of the volume of the power converter), $V_{battery}$ is the volume of the battery, and V_{supply} is the volume of the power converter. The run time is maximized when the incremental reduction in power dissipation caused by increasing the size of the power converter equals the power density delivered by the battery.

5.2 Improving Efficiency of the Output Stage

Although an ideal down converter is lossless, any physical realization of such a converter will be subject to resistive losses and switching losses. For low output voltage systems, synchronous rectification is practically required for tolerable efficiency. In order to achieve synchronous rectification, an NMOS device is used as the low side switch. If the high side switch is implemented with a PMOS device, the output stage of the converter resembles a standard digital CMOS buffer, and could be driven as such (with the gates tied together). However, there are important differences between the down converter switches and a CMOS buffer, and such a drive scheme should be avoided to improve efficiency.

5.2.1. Short Circuit Dissipation

The load on a CMOS buffer is capacitive, and the sizing of such buffers is best chosen so that the ratio of the load capacitance to the input capacitance of the buffer is related by a factor of about 3. Under this condition, the rise time on the input signal will match the fall time on the output. Short circuit power is typically limited to around 10% of the losses incurred charging the output capacitance [5]. For the down converter switches, the device sizing is chosen not to drive capacitive loads, but to handle the DC current present when the switches are on. As a result, the capacitance on the output terminals (drains) of the switches is significantly less than 3 times the input capacitance on the device gates. The output transition is much faster than the input transition, and as a result, the short circuit power can increase dramatically (over the CMOS buffer case).

Short circuit dissipation can be reduced by introducing a fixed delay between the signal which turns off one switch and the signal which turns on the other switch. This delay must be long enough to ensure that the two devices to not conduct simultaneously. The maximum tolerable fixed delay is set by the desire to avoid body diode conduction in the synchronous rectifier. When the upper switch turns off, the current in the inductor discharges the node V_x. If the lower switch is not turned on in time, the body diode of the lower switch will become forward biased. This situation results in undesired losses. The maximum delay is approximately given by

$t_{d-max} = (C_x V_{in})/I_{peak-max}$, where C_x is the parasitic capacitance at node V_x.

5.2.2. Zero Voltage Switching (ZVS)

A much discussed method of improving efficiency is zero voltage switching (ZVS) [18]. ZVS requires the precise timing of the gate drives of the output switches to avoid dissipating the energy stored in parasitic capacitances at node V_x. During switch transitions, the parasitic capacitance C_x is shorted to either ground or the input voltage through one of the power switches. This causes the dissipation of an energy $(C_x \Delta V^2)/2$, where ΔV is the voltage across the capacitor when it is shorted. ZVS is the technique of timing the switch transitions so that V_{DS} is zero when a switch is turned on (the inductor fully charges or discharges C_x). Figure 12 illustrates the timing of the switches and the V_x waveform.

Achieving ZVS requires adaptively setting the delay between the turn off of one switch and the turn on of the other switch, since the rise time and fall time of V_x varies with the load current. Obtaining ZVS on the turn on of the high side device requires that the inductor current be negative at the start of the switching cycle, so that the current in the inductor charges V_x to V_{in}. Stratakos [18] proposes setting $I_{pk-pk} > 2I_{avg}$ to allow ZVS to function at all loads. Note that if ZVS is practiced, there will be no short circuit dissipation.

5.2.3. Discontinuous Conduction

When the two switches of the output stage are driven with complementary signals, so that one switch is always on, the converter will always remain in continuous mode. The lower switch will not act as a true rectifier. In continuous conduction mode, the inductor current is a triangular wave with some offset, the average value being the actual output current. The resistive losses associated with such a waveform are $R(I_{avg}^2 + I_{pk-pk}^2/12)$. Because of the small constant in front of the ripple term, one may design I_{pk-pk} to be large relative to the nominal value of I_{avg} (note however that a large ripple reduces the energy storage density of the magnetic elements). In many systems, not only is the efficiency important at full load, but the efficiency at relatively light loads may also be a matter of concern. When the output current is reduced, I_{pk-pk} remains constant, and this ripple term will dominate the resistive losses. If $I_{pk-pk} = 2I_{avg}$ at maximum loading (the condition required for ZVS to function), then at ten percent loading, the dissipation due to DC current will be 1/100 the peak value, although the total dissipation only decreases to 1/4 the peak value.

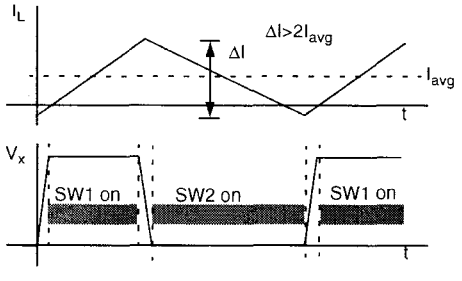

Figure 12: Illustration of switch timing for ZVS.

Entering discontinuous conduction mode, which requires additional control of the synchronous rectifier, can provide high efficiency over a broader range of loads. If discontinuous mode is used, the low side switch must be turned off when its current goes to zero and starts to become negative. This can be accomplished by turning off the switch when V_x crosses zero or by explicitly measuring the current in the switch or inductor. For the same case examined above, where at maximum load $I_{pk-pk} = 2I_{avg}$, if the converter enters discontinuous mode at lower loads, the resistive dissipation at 10% load will be $1/(10\sqrt{10})$ (about 1/32) the peak dissipation, an improvement of a factor of 8.

5.3 Reducing Control Power

A complete switching converter includes control circuitry which monitors the output voltage to set the duty cycle on the power switches, and creates the pulse width modulated signal to drive the switches. A block diagram of a converter with feedback networks is shown in Figure 13. The compensation filter, *G(s)*, shown depicts a zero frequency integrator for precise DC output regulation. A zero cancels the phase lag added by the integrator, providing extra phase margin at the unity gain frequency. A higher frequency pole accomplishes a first order roll off at crossover. Crossover (and hence the high frequency pole) is designed to be below the output filter resonance. The ultimate output of the compensation filter is a signal representing the desired duty cycle. A second block in the system takes the output of the compensator, be that a digital word or analog signal, and generates the pulse width modulated signal, at a fundamental frequency well above the filter resonance (typically 20—100). This signal is driven onto the power switches and filtered by the output filter, *F(s)*. The magnitude of the resonance of this filter is a concern; a large resonance can cause instability in the closed loop system.

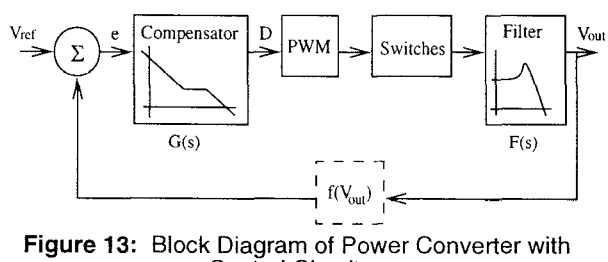

Figure 13: Block Diagram of Power Converter with Control Circuitry.

For a variable voltage output, as discussed in Section 2, the feedback error signal is not a measure of the difference between the output voltage and an external reference, but may be a metric based on the throughput or workload of the DSP load. In this case an extra block is added to the feedback path, $f(V_{out})$, which for instance, could represent the delay of the digital load as a function of the supply voltage. In such a case the dynamics of this block are extremely fast (relative to the filter frequency and crossover frequency), so no phase is added to the loop transfer function. However, $f(V_{out})$ will probably have significant non-linearities which will require special attention in the design of the feedback gain.

When considering the efficient design of power supplies for load powers in the milliwatt range, we must reduce the power consumed by the power supply's control circuitry. Publications to date [18], [19] regarding power supplies for low power loads have not addressed the need to push control power into the microwatt range (and for good reason; they addressed loads in the hundreds of milliwatts, where control power was not as critical). To control systems which monitor the delay of a digital circuit, it is convenient for the power supply controller to also be a digital circuit [6], [19]. Digital controllers may also help provide the control functions with very little power. The essential functions of a controller are: measuring the output voltage and comparing it to a reference, creating a signal representing the desired duty cycle, and creating a pulse width modulated signal to drive the output power switches.

5.3.1. Control Strategy

One way to reduce the control power for low power regulators is to implement control strategies which inherently require very little power dissipation. Some examples of such strategies are variable voltage control based on the frequency of a ring oscillator and single bit (comparator) feedback on the output voltage. The former method entails measuring the delay of a digital load circuit using a ring oscillator and a counter. The value which accumulates in the counter after a fixed period of time can then be compared to a reference value representing the desired delay of the digital load. The difference of the two values can be used to update the PWM duty cycle signal. The entire operation avoids the static dissipation of analog circuitry, and requires very little digital computation. Single bit feedback on the output voltage [17] utilizes a dynamic comparator to occasionally sample the output voltage and determine whether the output voltage is above or below the reference value. If the output voltage is higher than the reference, the duty cycle is reduced slightly and if the output is below the reference, the duty cycle is increased slightly. Since the comparator is dynamic, and is operated infrequently, its dissipation can be very small. Comparator signals are used for feedback in some commercial controllers with pulse-frequency control [20].

5.3.2. PWM Design

After a digital word representing the desired duty cycle has been created, the actual switching waveform must be generated. Typically, in digital systems PWM signals are created by using a clock at some multiple of the switching frequency with a counter. The PWM signal is set high at the beginning of a switching period, and then reset after the counter detects that some number of cycles of the faster clock have passed. This approach has been applied to digital power supply controllers [19].

Ultra-fast-clocked counters are not particularly well suited for low power operation. The clock frequency at which the counter operates can be extremely high (e.g., a 1MHz switching waveform with 256 discrete levels of duty cycle requires a 256Mhz clock!). As a result of the short delay requirement, such a circuit does not lend itself to voltage scaling. The power reported [19] for a controller utilizing a fast-clocked counter is on the order of milliwatts or more.

A second approach to transforming a digital word to a pulse width modulated signal is to use a tapped delay line [17]. The delay line can be constructed with current starved buffers, and controlled with a feedback loop so that the delay of the delay line matches the period of the switching clock. The number of taps on the delay line is chosen to be the number of desired discrete duty cycle levels. The taps of the delay line are then addressed with a multiplexer. The PWM signal is set high as a pulse is

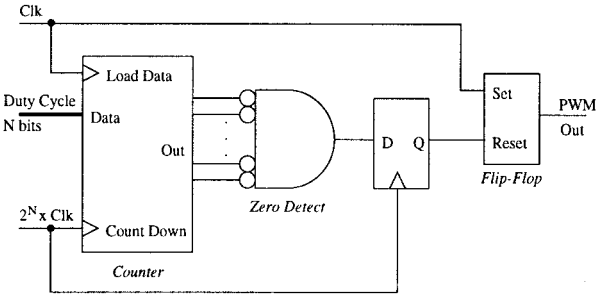

Figure 14: Fast clocked counter approach to create a PWM signal.

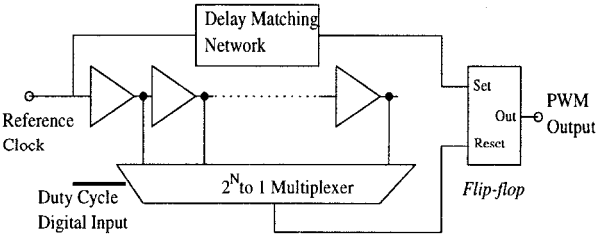

Figure 15: Tapped delay line PWM.

input to one end of the delay line, and reset when that pulse arrives at the selected tap. Lacking any need for short delays, the tapped delay line PWM can tolerate deep voltage scaling. Such a PWM circuit has been fabricated, and demonstrates a power dissipation on the order of tens of microwatts[17].

6. Conclusion

Extremely low power operation can be achieved for digital circuits by aggressively reducing the supply voltage. New technologies can be exploited to enhance the efficiency of low supply voltage circuits, such as multiple threshold CMOS, variable threshold bulk-CMOS and variable threshold SOI. Variable supply voltage circuits can be used to increase power savings for variable workload applications. Low voltage and low power systems also require designers to investigate power converter circuits to create low supply voltages and variable supply voltages efficiently. The efficiency of such converters can be improved with synchronous rectification, discontinuous conduction at light loads, and a reduction in the dissipation of the control circuitry.

Acknowledgments

This work is sponsored by ARPA, under contracts #DAAL-01-95-K3526 and #DABT63-95-C-0088. Abram Dancy is also supported by an NSF Graduate Fellowship. We would like to thank V.Gutnik, I. Yang and Prof. D. Antoniadis for their contributions to this paper.

References

[1] A. Chandrakasan, R. Brodersen, *Low Power Digital CMOS Design*, Kluwer Academic Publishers, July 1995.

[2] Toshiyuki Sakuta, Wai Lee and Poras T. Balsara, "Delay Balanced Multipliers for Low Power/Low Voltage DSP Core," *1995 IEEE Symposium on Low Power Electronics*, pp. 36-37, Oct. 9-11, 1995.

[3] Y. Nakagome, "Sub-1-V Swing Internal Bus Architecture for Future Low-Power ULSI's," *IEEE Journal of Solid-State Circuits*, Vol. 28, No. 4, pp. 414-419, April 1993.

[4] M. Hiraki, H. Kojima, H. Misawa, T. Akazawa, Y. Hatano, "Data-Dependent Logic Swing Internal Bus Architecture for Ultra low-Power LSI's," *IEEE Journal of Solid-state Circuits*, pp. 397-401, April 1995.

[5] H. Veendrick, "Short-Circuit Dissipation of Static CMOS Circuitry and Its Impact on the Design of Buffer Circuits," *IEEE Journal of Solid-state Circuits,* pp. 468-473, August 1984.

[6] V. Gutnik, A. Chandrakasan, "An Efficient Controller for Variable Supply-voltage Low Power Processing," *IEEE VLSI Circuits Symposium*, June 1996.

[7] L. Nielsen, C. Niessen, J. Sparso, K. van Berkel, "Low-Power Operation Using Self-Timed Circuits and Adaptive Scaling of Supply Voltage," *IEEE Transactions on VLSI systems*, December 1994, pp 391-397.

[8] M. Weiser, B. Welch, A. Demers, S. Shenker, "Scheduling for Reduced CPU Energy,"pp. 13-22, *First Symposium on Operating Systems Design and Implemenration*, Usenix, 1994.

[9] P. Maken, M. Degrauwe, M. Van Paemel, H. Oguey, "A Voltage Reduction Technique for Digital Systems," *IEEE ISSCC*, pp. 238-239, February 1990.

[10] M. Horowitz, "Low Power Processor Design Using Self-Clocking," *Workshop on Low-power Electronics*, 1993.

[11] J. Burr, J. Shott, "A 200mV Self-Testing Encoder/Decoder Using Stanford Ultra-Low-Power CMOS," *IEEE ISSC*, pp. 84-85, 1994.

[12] I. Yang, C. Vieri, A. Chandrakasan, and D. Antoniadis, "Back Gated CMOS on SOIAS for Dynamic Threshold Control," *IEEE 1995 IEDM*, December 1995.

[13] M. Srivastava, A. P. Chandrakasan, and R.W. Brodersen, "Predictive System Shutdown and Other Architectural Techniques for Energy Efficient Programmable Computation," *IEEE Transaction on VLSI Systems*, pp. 42-55, March 1995.

[14] T. Sakata et al., "Subthreshold-Current Reduction Circuits for Multi-Gigabit DRAM's," *IEEE Journal of Solid-state Circuits*, Vol. 29, No. 7, pp. 761-769, July 1994.

[15] M. Mutoh, T. Douskei, Y. Matsuya, T. Aoki, S. Shigematsu, and J. Yamada, "1-V Power Supply High-Speed Digital Circuit Technology with Multi-threshold Voltage CMOS," *IEEE Journal of Solid-state Circuits*, pp. 847-854, August 1995.

[16] T. Kuroda et. al., "A 0.9V 150MHz 10mW 4mm^2 2-D discrete cosine transform core processor with variable threshold-voltage (V_T) scheme," *IEEE Journal of Solid-state Circuits*, pp. 1770-1779, November 1996.

[17] A. Dancy, *Power Supplies for Ultra Low Power Applications,* M.Eng. Thesis, Massachusetts Institute of Technology, August 1996.

[18] A. Stratakos, S. Sanders, R. Brodersen, "A Low-Voltage CMOS DC-DC Converter for a Portable Battery-Operated System," *IEEE Power Electronics Specialists Conference,* pp. 619-626, 1994.

[19] G. Wei and M. Horowitz, "A Low Power Switching Power Supply for Self-Clocked Systems," *1996 International Symposium on Low Power Electronics and Design,* pp. 313-318, 1996.

[20] Maxim Data Sheets, MAX756, MAX624. *http://www.maxim-ic.com/p_switch.html*