A Reconfigurable Dual Output Low Power Digital PWM Power Converter

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

ABSTRACT

This versatile power converter controller provides dual outputs at a fixed switching frequency and can regulate either output voltage or target system delay (using an external L-C filter). In the voltage regulation mode, the output voltage is monitored with an A/D converter, and the feedback compensation network is implemented digitally. The generation of the PWM signal is done with a hybrid delay line/counter approach, which saves power and area relative to previous implementations. Power devices are included on chip to create the two independently regulated output PWM signals. The key features of this design are its low power dissipation, reconfigurability, use of either delay or voltage feedback, and multiple outputs.

1. Power Converter Requirements

In portable systems, electronic circuits can be designed to operate over the range of the voltages supplied by the battery over its discharge cycle. However, adding some form of power regulation can significantly increase battery life, since it allows circuitry to operate at the “optimal’’ supply voltage from a power perspective. Given the advances in power management techniques (e.g., low-voltage operation [1]), there is a need for efficient DC-DC converters at output power and voltage levels previously uncommon for such circuits. A high-efficiency low-voltage DC-DC converter has been reported that delivers 750mW [2] and several commercial controllers are currently available for the 100mW to 1W range. This paper describes techniques for high-efficiency low-voltage regulation for power levels down to 100’s μW.

Many portable systems such as cellular phones and PDAs work in an event driven fashion and have a low duty cycle. In such systems, only a small section of the chip will be turned on during the standby mode and the power dissipation of this circuitry can have a significant impact on the battery life of the system. In order to maintain efficient operation at very low output powers, the power dissipation of the control circuitry, as well as that of the power conversion circuit must be minimized. The converter must be designed to operate efficiently over wide variations in output load power.

Low power systems are being designed with multiple power supply voltages [3][4][5]. The basic approach to reduce power dissipation is to use reduced power supply voltages for modules not in the critical path of the computation. This technique requires the generation of multiple power supply voltages efficiently. A brute force approach is to use separate controllers for each output. In this paper, we describe techniques to re-use portions of the controller for multiple outputs. As an example, a dual-output supply is demonstrated.

Finally, there are many systems where the amount of processing per input sample (i.e., the computational workload) varies with time [6][7][8]. For such systems, one approach to save power is to dynamically vary the power supply voltage as the load varies. From a power supply perspective, this translates to a need to design the regulator control for a quick transient response. Even if the workload does not vary, the power supply should be dynamically adjusted to compensate for temperature and process variations [9].

2. System Architecture

Figure 1 shows a block diagram of the dual output DC-DC converter. The converter operates by creating a pulse width modulated signal of some duty cycle at node V_1 (and similarly at node V_2), whose average value is the desired output voltage. External passive filtering is used to filter the PWM signal, creating a DC voltage with some tolerable value of ripple.

In order to provide reasonable efficiencies for the low supply voltages present in low power digital systems, power converters must incorporate synchronous rectification (i.e., active power devices are used to replace diodes) [2]. A drawback of synchronous rectifi-

cation is that without explicit monitoring of the output current and control of the synchronous rectifier, the circuit will not enter discontinuous mode at light loads. The resulting ripple current in the inductor will cause resistive losses that will reduce efficiency at light loads. Hence, the ability to create a “turn-off” signal for the synchronous rectifier could be an important feature for a low power controller.

This converter has the ability to regulate either an output voltage or target system delay. That is, the input feedback signal is taken either from the A/D converter or the delay feedback input. The delay feedback input allows the controller to measure the speed of operation of a load circuit. The input is a signal from a ring oscillator formed from the critical path of the circuit to be controlled. This enables the operation of the controller in a variable supply voltage system, where the supply voltage is minimized dynamically over variations in process, temperature, and workload ([6]-[11]).

The compensation network for the output of the power converter is a variable gain integral controller. A reference value (in a digital form) is subtracted from the A/D or delay measurement, and the difference is scaled in an array multiplier stage. The product is then subtracted from the previous duty cycle command to produce the next duty cycle command. The internal representation of the duty cycle is 12 bits, and the 10 MSBs are passed to the PWM stage to create the output. The compensation sample rate is programmable; the sample rate is primarily limited by the A/D conversion time. The compensation network elements (adders and multiplier) are time multiplexed to derive duty cycle commands for both of the outputs. The reference value and gain for each of the two outputs and other configuration registers are fully programmable through a bidirectional two wire serial interface.

There are two outputs; the first is optimized for a 20mA, 2V load, and the second is optimized for a 1mA, 1V load. Guard rings help to isolate the power output stages from the core digital logic. Additional guard rings separate the A/D capacitor array and low current bias reference from the power stages and digital core.

The switching frequency of the 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 trade-offs between efficiency, power density, and transient performance.

3. A/D Converter

The A/D converter is a seven bit, standard charge redistribution converter. The advantage of a charge redistribution converter for low power applications is that it can be implemented without amplifiers, which would typically cause significant static currents to be dissipated. A dynamic comparator was utilized to compare the capacitor array voltage to an external analog reference at each stage of the conversion.

The capacitor array utilizes common centroid layout, and there are two rows and columns of dummy devices on the perimeter of the array to enhance matching. Due to the relatively low resolution of the converter, unit capacitor sizing was rather aggressive; a 10μm by 10μm poly-poly capacitor giving 47fF of capacitance.

A schematic of the dynamic comparator used is shown in Figure 2. This design is relatively common [12]. In this comparator, the offset voltage is a function of the parameter matching (dimensions and thresh-

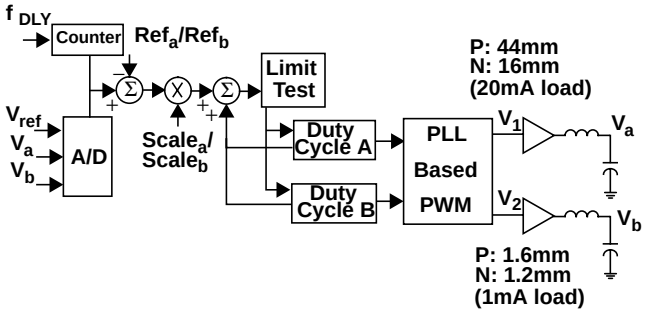

Figure 1. Block diagram of the dual output DC-DC Converter.

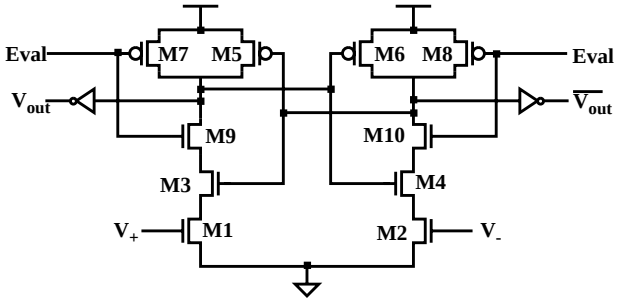

Figure 2. Schematic of the dynamic comparator used in the A/D converter.

olds) between the pairs of devices M1, M2; M3, M4; and M5, M6. When a comparison is initiated by a rising Eval signal, M1 and M2 begin discharging the nodes Out and $\overline{\text{Out}}$. The cross-coupled feedback causes whichever node is falling more slowly to become latched high.

4. Pulse Width Modulation

After a digital word representing the desired duty cycle has been created, the actual switching waveform must be generated. When using analog circuits, a PWM signal is typically created by comparing a ramp signal to a reference value with a static comparator (this requires DC current flow). Digital PWM circuits can avoid the problem of static power dissipation.

In digital systems, PWM signals are typically 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. Figure 3 shows a block diagram of the counter based PWM approach. Unfortunately, ultra-fast-clocked counters are not particularly well suited for low power operation. The counter clock frequency is chosen to be 2^N times the switching frequency of the converter, where N is the number of bits in the digital command word. The clock is used to divide the switching period into 2^N increments. For example, 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. A digital

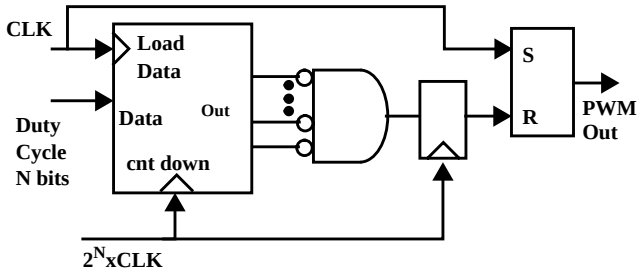

Figure 3. Fast counter based PWM generator.

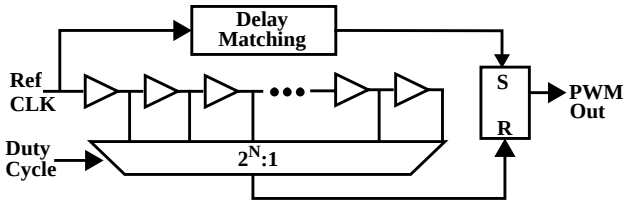

Figure 4. Delay line based PWM generator.

controller has been reported that uses the counter based approach, but the power of the controller alone is on the order of milliwatts [11]. This is acceptable for DC-DC converters that deliver power in the Watt range, but not for systems in the milliwatt range.

Another way to create a PWM signal from an N-bit digital value is to use a tapped delay line [13]. Since this approach uses the switching frequency clock, the power is significantly reduced relative to the fast-clocked counter approach. Figure 4 shows a schematic for the delay line based digital word-to-PWM circuit. The essential components of a tapped delay line PWM circuit are the delay line and a multiplexer. A pulse from a reference clock starts a cycle, and sets the PWM output to go high (after a delay designed to match the propagation delay experienced through the multiplexer). The reference pulse propagates down the delay line, and when it reaches the output selected by the multiplexer, it is used to set the PWM output low. The total delay of the delay line is adjusted so that the total delay is equal to the reference clock period. That is, feedback is used to turn the delay line into a delay-locked-loop (DLL), which locks to the period of the input clock. This approach is very power efficient, however, can require significant implementation area. If multiple PWM signals are needed, it requires the addition of multiplexers to a single delay line.

A hybrid scheme (Figure 5) is described here that provides considerable advantages over either of these

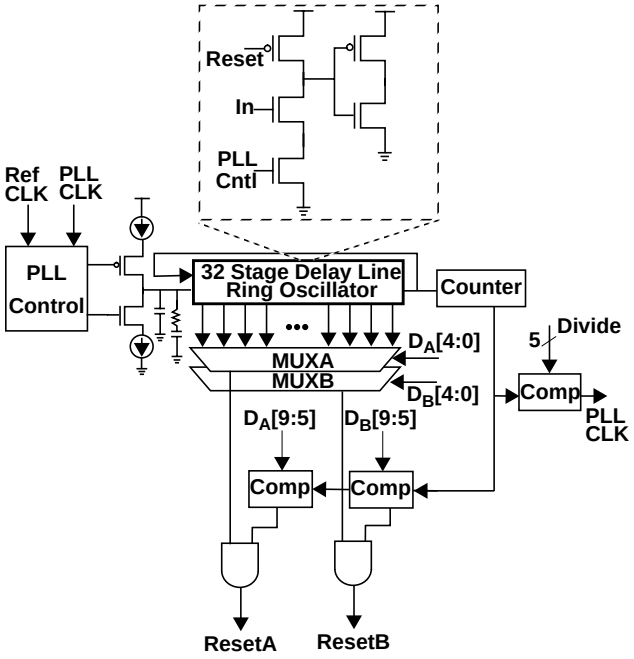

Figure 5. PWM generation block, showing PLL charge pump and dual output hybrid delay line/ counter PWM approach.

approaches. A 32 stage delay line forms the basis for the pulse width modulation stage. The delay line is configured as a ring oscillator, which is phase locked to a reference clock. A divider allows the ring oscillator frequency to be set between 2 and 32 times faster than the reference frequency. The taps of the delay line then divide the input clock period into between 64 and 1024 equal increments. The taps of the delay line are sensed by two 32 to 1 multiplexers, one for each of the output PWM signals. The rising edge of the reference clock sets the PWM signals high. A PWM signal is set low when a pulse arrives at the tap of the delay line selected by its multiplexer for the Nth time, where N represents the 5 MSBs of the 10 bit duty cycle command.

The delay of the delay line is controlled by adjusting the gate signals on starvation-type NMOS devices. The gate control signal controls the speed of the positive going edge at the output of each buffer. Figure 6 shows how post-charge logic is used to ensure that the

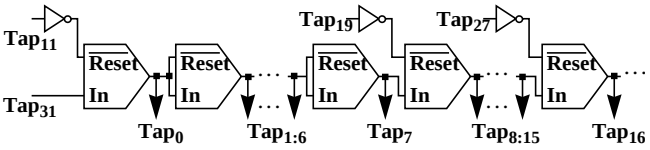

Figure 6. Post-charge logic in delay line matches leading edge and falling edge propagation times, and allows a ring oscillator to be created with even number of stages.

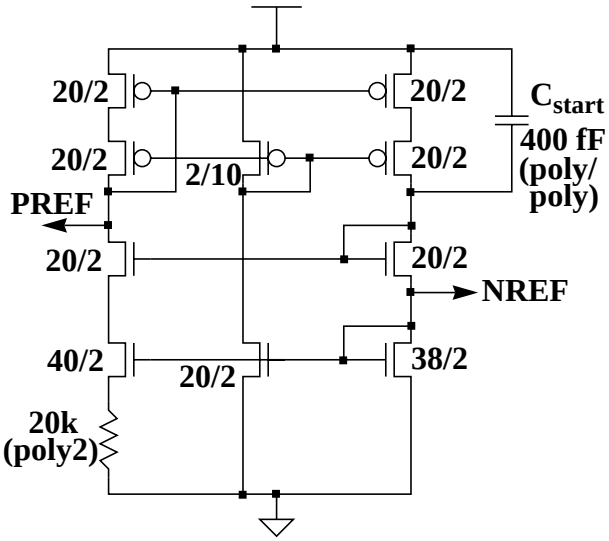

Figure 7. Low voltage modified Widlar 100nA current source, using MOS devices in subthreshold. This circuit generates bias voltages for the PLL charge pump. All dimensions in µm.

negative edge of the outputs travel at the same speed as the positive edge. The control node is charged up and down using a current source. The biasing for the current source is generated on chip with a MOS Widlar current source (Figure 7). The compensation network for the PLL control node is also implemented on chip with poly-poly capacitors and a poly-2 resistor.

The hybrid delay line/counter circuit reduces power dissipation relative to the fast-counter approach, by a 32X reduction in counter clock frequency (in this implementation). Compared to the delay line based PWM circuit, the hybrid approach gives a 9 times reduction in area; when leveraged to provide multiple outputs as done here, the effective area reduction is a factor of 12.

5. Experimental Results

Figure 8 shows a die photo of the dual output DC-DC converter. The A/D, control circuitry and power switches are integrated on the same die. The power switches are sized to trade-off the switching and conduction losses.

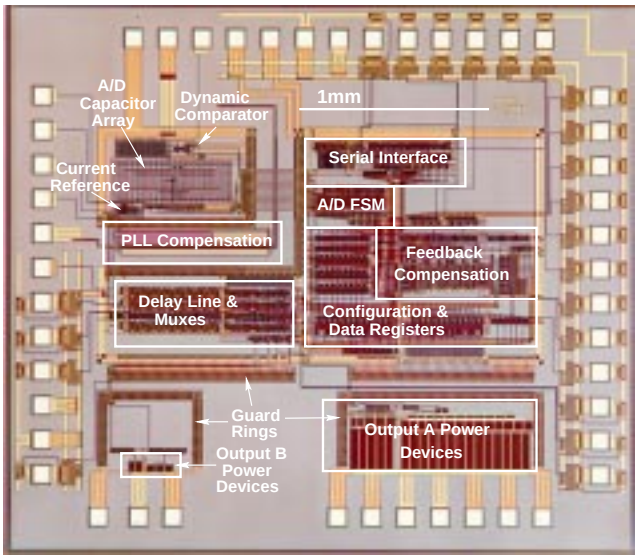

Figure 8. Die photo of dual output DC-DC converter.

Table 1 gives the details of the power converter chip including the chosen filter values.

Table 1. DC-DC converter chip summary

Parameter	Value
Die Size	3.2mm x 2.8mm (test chip is pad limited)
Technology	0.6µm DPDM

Table 1. DC-DC converter chip summary

Parameter	Value
A/D INL	± 0.5 LSB
A/D DNL	+0.3, -0.4 LSB
Inductor Value	220µH
Capacitor Value	0.22µF
Filter Area	0.024in^2
Output Ripple	< 40mV for switching frequency (f_{SW}) > 500kHz
Operating Frequency	< 2.5MHz

Figure 9 shows the transient response of the output voltage to step changes in commanded output. The switching frequency is 1MHz, and the feedback sampling period is 25µs. The output settles to 90% of the desired value in 100µs. The switching speed is limited by the conversion time of the A/D converter (since the A/D takes multiple switching periods for data conversion). The ability the supply voltage on demand allows the minimization of energy dissipation in variable load systems. The DC-DC converter, as mentioned earlier, can also be configured in a performance feedback mode. The performance feedback has been tested and is functional. The performance feedback requires the DSP load circuit to provide a clock signal derived from a ring oscillator matched to the critical path circuitry.

The jitter of the PLL is 5.5ns (Figure 10). The effect of this jitter on the output voltage is a slight broadening of the spectrum of the switching frequency

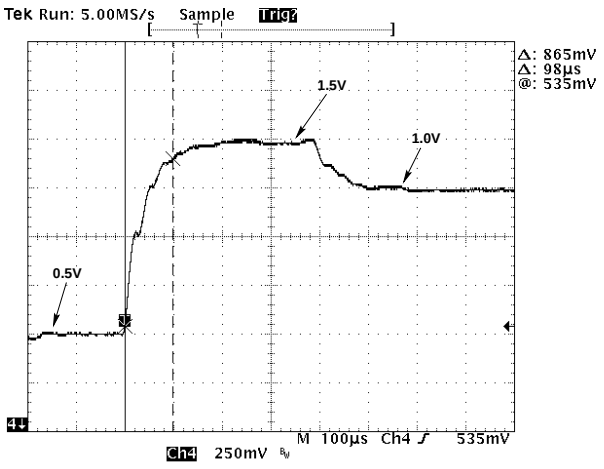

Figure 9. Transient response of filtered output voltage to changing digital reference commands (f_{SW}=1MHz).

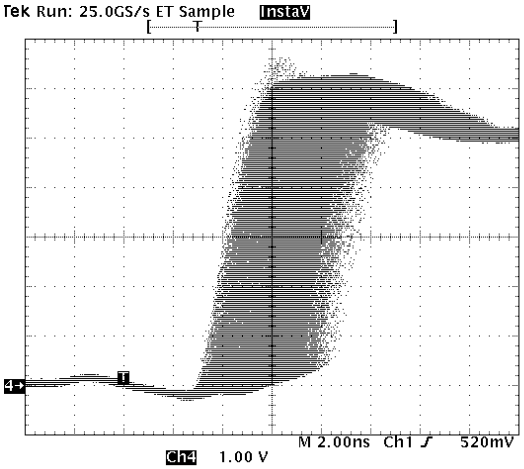

Figure 10. PLL jitter at last tap of delay line.

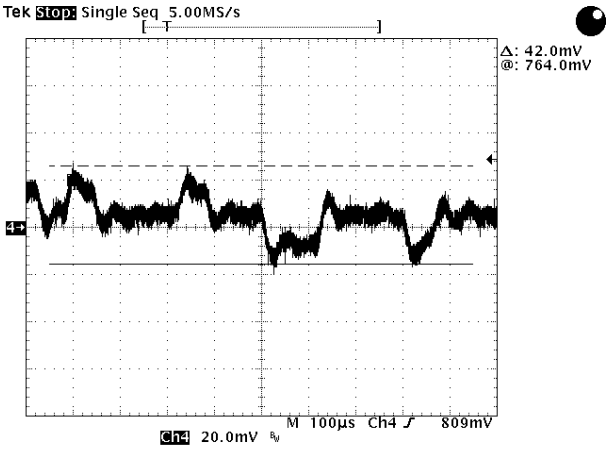

Figure 11. Regulated output voltage and ripple.

ripple. Figure 11 shows the regulated output voltage and the ripple.

Measured output efficiencies were between 89% and 80% over a range of output currents, for the particular output filter selected. There is a trade off between the size and cost of the output filter and the achievable efficiency. The filter selected here represents a low cost, small area selection. The losses are dominated by a 9.5Ω resistance in the output inductor at high output powers. Table 2 shows the output efficiency for this low cost, small sized inductor. Using an optimized inductor with low series resistance, we were able to achieve a total efficiency of 95% for the 2V output at a load of 45mA.

Table 2. Output efficiency using a low cost inductor with a series resistance of 9.5Ω

Output (f_{sw}=500KHz)	**Efficiency**
Output A	
V_{out}=2V	
I_{out} = 10mA	89%
I_{out} = 45mA	80%
Output B	
V_{out}=1V	
I_{out} = 750µA	89%
I_{out} = 10mA	80%

At 500kHz, with 256 levels of duty cycle resolution, the control circuit draws less than 45µA. Table 3 shows the power dissipation for two different configurations of the power converter (with 8 bit and 10 bit resolutions on the PWM).

Table 3. Power Dissipation for 1024 and 256 levels of the PWM.

Parameter	**Value**
1024 Taps, 500kHz	
Min. Supply Voltage	2.05V
PLL & Logic Current	199.3µA
Analog Circuits Current	1.5µA
256 Taps, 500kHz	
Min. Supply Voltage	1.35V
PLL & Logic Current	42.8µA
Analog Circuits Current	1.5µA

6. Conclusion

A digital PWM power supply converter has been described here that produces dual output voltages efficiently. This converter can be configured to regulate a fixed supply voltage or a processing speed. The ability to adapt supply voltage quickly can be exploited to minimize power dissipation in applications where the workload varies rapidly. Minimizing the fixed overhead loss of the converter is the key to achieving high efficiency at light load conditions. The power supply converter features a hybrid delay line and counter based PWM generator that is area and power efficient for generating multiple outputs.

Acknowledgments

This work is sponsored by DARPA, under contract #DAAL-01-95-K3526. Abram Dancy was supported by an NSF Graduate Fellowship.

References

[1] A. P. Chandrakasan, R. Brodersen, Low-Power CMOS Design, IEEE Press, January 1998.

[2] 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.

[3] S. Raje, M. Sarrafzadeh, “Variable Voltage Scheduling,” 1995 International Symposium on Low Power Design, pp. 9-14, 1995.

[4] K. Usami, M. Igarashi, F. Minami, T. Ishikawa, M. Kanazawa, M. Ichida, and K. Nogami, “Automated Low-Power Technique Exploiting Multiple Supply Voltages Applied to a Media Processor,” IEEE Journal of Solid-State Circuits, pp. 463-472, March 1998.

[5] J.-M. Chang and M. Pedram, “Energy Minimization Using Multiple Supply Voltages,” IEEE Transactions on Very Large Scale Integration (VLSI) Systems, pp. 436-443, December 1997.

[6] V. Gutnik, A. P. Chandrakasan, “Embedded Power Supply for Low-Power DSP,” IEEE Transactions on Very Large Scale Integration (VLSI) Systems, pp. 425-435, December 1997

[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, pp. 391-397, December 1994.

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

[9] P. Maken, M. Degrauwe, M. Van Paemel, H. Oguey, “A Voltage Reduction Technique for Digital Systems,” IEEE International Solid-State Circuits Conference, pp. 238-239, February 1990.

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

[11] G. Wei and M. Horowitz, “A Low Power Switching Power Supply for Self-Clocked Systems,” IEEE/ACM 1996 International Symposium on Low Power Electronics and Design, pp. 313-318, 1996.

[12] T.B. Cho and P.R. Gray, “A 10-bit, 20MS/s, 35mW Pipeline A/D Converter,” IEEE Custom Integrated Circuits Conference, pp. 23.2.1-23.2.4, 1994.

[13] A. Dancy, A. P. Chandrakasan, “Ultra Low Power Control Circuits for PWM Converters,” IEEE Power Electronics Specialists Conference, pp. 21-27, June 1997.