Ultra Low Power Control Circuits for PWM Converters

Abram P. Dancy and Anantha P. Chandrakasan
Department of Electrical Engineering and Computer Science
Massachusetts Institute of Technology
Cambridge, MA 02139 USA

Abstract— **Low voltage high-efficiency DC-DC conversion circuits are integral components of battery operated portable systems. This paper presents the design of a pulse width modulated (PWM) controller for a milliwatt level DC-DC converter. Low power control circuitry is constructed by utilizing a delay line based PWM generator and low-resolution feedback. A DC-DC converter has been fabricated which delivers 5mW at 1V, while achieving an efficiency of 88%. The PWM consumes only $10\mu W$.**

I. Introduction

In mobile 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. As the power dissipation of electronic circuits drop (exploiting low-voltage process technology and other power management techniques), there is a need for high-efficiency DC-DC conversion circuits to deliver low power levels. A high efficiency low-voltage down converter that delivered power levels of 750mW has been previously reported [1], and there are numerous commercial controllers available to create logic level outputs in the 100mW to 1W range [2], [3].

For systems that operate with a very low duty cycle (such as pagers and cellular phones), but have some circuitry which operates when the rest of the system is off, the power dissipation of the standby circuitry can have a significant effect on the battery life of the system. Systems with varying load conditions in different operating modes present the challenge of creating a voltage converter which can operate efficiently over a wide range of output powers. 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.

This work addresses high-efficiency conversion techniques at power levels on the order of hundreds of microwatts to tens of milliwatts. Although this power level is very low relative to most current system power levels, it is actually possible to implement complex DSP, such as video compression, with only a few milliwatts [4], [5]. In order to facilitate integration with other system components and the possibility of variable-voltage operation, a digital PWM controller was designed. Also, low resolution feedback of the output voltage is employed to provide regulation of the converter output with very little control overhead power. Fig. 1 shows a block diagram of our low-power integrated down converter.

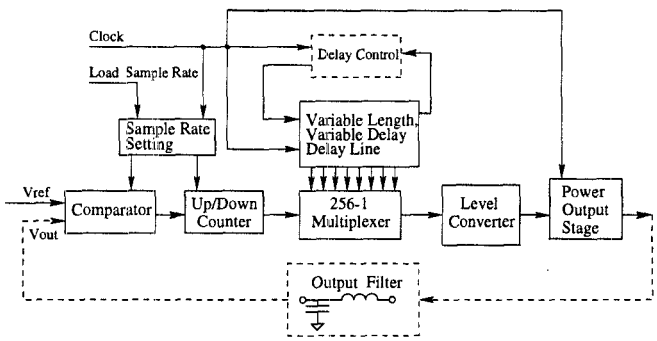

Fig. 1. Block diagram of implemented converter. Dashed lines represent blocks not integrated on chip.

II. Control Strategies for Low Power Converters

A. Requirements

Typically, power converter design has focused on output regulation, power density, and dynamic response. Low power systems present new requirements for power converters, due to the need for higher levels of integration, the optimization of the supply voltage to the load conditions, and the low output voltages of low power digital loads. The desire to increase system integration can force the power converter controller and switches to be implemented on the same die as its load. Optimization of the supply voltage can require the output supply voltage to vary dynamically, based on the computational load or environmental conditions. Intermittent operation, or multiple-mode operation requires efficient operation over a broad range of output powers. Output powers may also be reduced to the level that the power consumption of the converter control circuits is a concern.

When a power converter is integrated on the same die as a digital load, it must be designed to have some degree of immunity to the switching noise which will be generated by that load. Another challenge in such an application is the need to reduce the number of external parts and bond pads required for the operation of the power converter.

0-7803-3840-5/97/$10.00 © 1997 IEEE

The power consumption of a digital circuit increases with the square of its supply voltage, so in order to minimize power consumption of digital circuits, it is desirable to reduce the supply voltage as much as possible. For the greatest savings, this reduction can be done in a closed loop fashion, where the speed or computational load of a digital circuit is measured, and the supply voltage adjusted accordingly. This may be done to compensate for variations in circuit performance due to processing and temperature [6], or the current computational load [7]. Such measurements may be conveniently available in digital form from a frequency counter, motivating the creation of a digital power supply controller.

For variable voltage schemes, the speed at which the output voltage is required to change can vary from system to system. For video decoders, rapid changes in the supply voltage allows the system power to be minimized depending on the content of each frame. In other cases, the voltage change can be slow, changing only as the system adjusts from full output power to a standby or "keep-alive" mode. The precise regulation of the output voltage is not necessarily critical. Most of the savings possible with variable voltage operation can be achieved by using only a few discrete output voltages [7].

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

B. Current Techniques

The most common power converter controller consists of an amplifier which creates a voltage proportional to the desired duty cycle, by comparing the output voltage to a reference, and perhaps adding some dynamics such as an integrator. Some problems with this technique for low power, integrated applications are sensitivity to noise (when integrated with a digital load circuit), difficulty interfacing to a digital frequency feedback signal (for variable voltage operation), and challenges achieving low power (microwatt level) operation.

Another control technique employed in low power DC-DC converters is pulse frequency modulation (PFM) [3]. With PFM, pulses of current are output to the load filter, and voltage regulation is achieved by varying the time delay between pulses (the power circuit is always operating in discontinuous mode). The control can be implemented without a feedback amplifier, by simply comparing the output voltage to a reference, and initiating current pulses when the output voltage is low. A slight variant of this method is a burst-mode controller (e.g. [2]), where a fixed frequency train of current pulses are applied when the output voltage is below the reference. Drawbacks of PFM operation (as described) are the use of discontinuous mode at all loads and variable frequency switching. These characteristics increase the radiated output noise (EMI), and using discontinuous operation for all loads requires oversizing the output filter inductor to accommodate the inherently large peak inductor current.

A fully digital controller can also be used to control a DC-DC power converter. A digital PID (proportional, integral, derivative) controller has been presented [8], which is suited for situations where the feedback signal is available in digital form. This type of controller has the potential to have very low power consumption, and its digital nature makes it a good choice to be integrated with other systems.

C. Low Resolution Digital Feedback

In order to achieve the low power operation of the commercial PFM controllers, while retaining fixed frequency, continuous mode operation, a very simple A/D conversion with some digital processing can accomplish the voltage feedback function. This entails far less power dissipation than that required by even the simplest analog feedback circuits, due to the static current draw of the analog circuits. Minimally, the A/D conversion can be accomplished by a comparator, which decides whether the sensed output voltage is above or below a reference voltage. If the comparator is implemented as a dynamic comparator, which evaluates only upon command, it will not dissipate any power outside of the brief evaluation period. In this single bit feedback case, a digital word representing the duty cycle is created by a counter which counts up when the output is lower than the reference and counts down when the output is higher than the reference (see Fig. 1). This digital counter accomplishes the functionality of an integrator.

There are certain disadvantages to using a single bit feedback signal. First, the output will never reach an equilibrium voltage, rather it will approach a limit cycle, since the output is always measured to be in error (either too high or too low) and the error signal is not proportional to the magnitude of the error. Second, the response of the control loop must be slow in order to limit the magnitude of the steady state limit cycle. After the duty cycle is changed, it should not be modified again until the dynamic response of the output filter has decayed. If the duty cycle is adjusted faster than the filter response time, the magnitude of the steady state

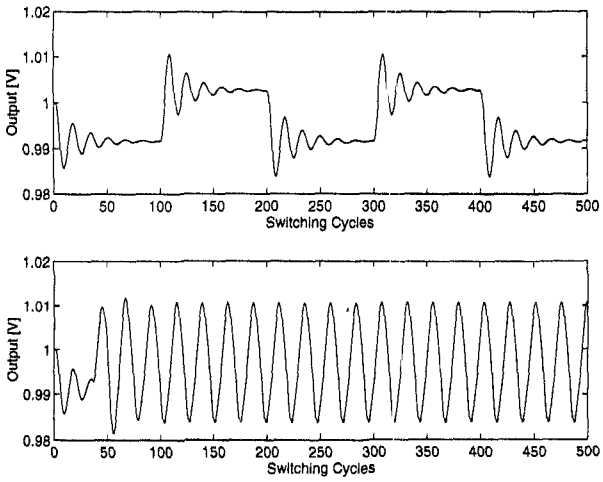

Fig. 2. Simulation of low-resolution feedback limit cycle. For both plots, the resolution on the duty cycle is 8 bits (256 different discrete duty cycles are possible). In the upper plot, output voltage is sampled once every 100 switching cycles, and the duty cycle is incremented or decremented one level per sample. In the lower plot, output voltage is sampled once every 12 switching cycles, and the duty cycle is incremented or decremented only 1/8th of a full increment in the output at each sample.

limit cycle will be large (several duty cycle increments) or unbounded.

The statement that increasing the sample rate can cause instability runs counter to the conventional thinking that higher sample rates improve performance. For the system described, increasing the sample rate also increases gain, since the up/down counter is incremented or decremented by one for every sample. If the amount the counter is changed at each sample is reduced, the rate at which the output voltage is sampled can be increased. The simulation results shown in Fig. 2 illustrate this tradeoff. In the upper plot, the comparator evaluates once every 100 switching cycles. On the lower plot, the comparator evaluates 8 times as often, but at each sample, the duty cycle is incremented only 1/8th of the minimum duty cycle increment. In both cases the resolution of the duty cycle is 256 levels, and the output limit cycle is the result of the commanded duty cycle switching between two consecutive values. Notice that the magnitude of the limit cycle remains approximately constant, but the frequency content changes between the two cases. Also note that the ripple is not at the switching frequency; the ripple in the output voltage is due to the excitation of the output filter resonance. The response will change significantly if the filter damping increases, but the limit cycle behavior is inherent to this control mechanism.

The existence of a steady state limit cycle is not in itself a severe problem. If the duty cycle increments are small enough, the magnitude of the limit cycle can be on the order of the switching frequency ripple. The limit cycle could also be eliminated by generating a three state feedback representing too high, too low, and acceptable. In some cases, the slow response time of the control loop may be an important consideration. This can be addressed by performing comparisons between the output voltage and other reference values, creating an error signal with more than just a single bit of resolution. Additional comparisons could eliminate the limit cycle behavior and allow faster response time.

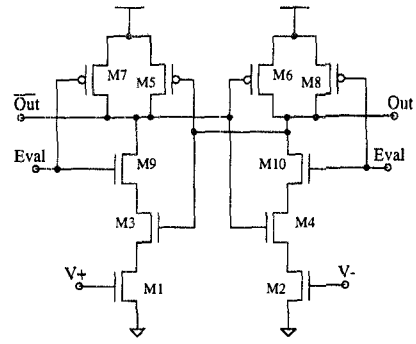

Fig. 3. A CMOS dynamic comparator. $V+$ and $V-$ are the inputs, Out and $\overline{Out}$ are the outputs. Outputs are valid some time after $Eval$ has gone high.

A schematic of the dynamic comparator used is shown in Fig. 3. This design is relatively common, for example, see [9]. In this comparator, the offset voltage is a function of the parameter matching (dimensions and thresholds) 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 $\overline{Out}$. The cross-coupled feedback causes whichever node is falling more slowly to become latched high.

III. Delay Line Based PWM

A. PWM Signal Generation

When using analog circuits, a PWM signal is typically created by comparing a ramp signal to a reference value with a static comparator (a static comparator is one which requires DC current draw). Digital PWM circuits can avoid the problem of static power dissipation. A previously reported method of creating a PWM signal from a digital command is to use fast-clocked counters [8], but the power of the reported controller alone is on the order of milliwatts. A clock frequency f_{clk} is chosen to be 2^N times the switching frequency, f_{SW}, 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. This method requires the use of very fast clocks for large values of N, and the need to run an N bit

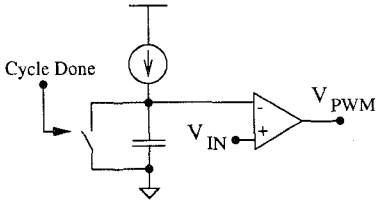

(a)

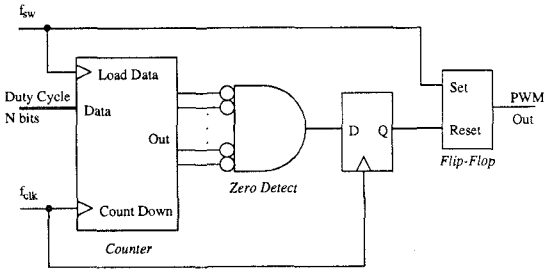

(b)

Fig. 4. Standard analog and digital PWM signal generation. In (a), V_{IN} is the analog duty cycle command, and *Cycle Done* is a narrow-pulse clock signal.

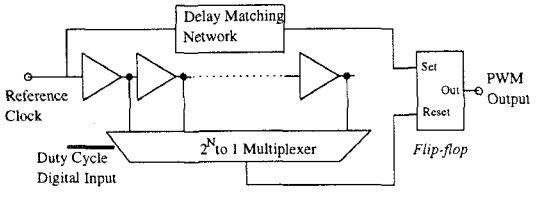

Fig. 5. Proposed PWM circuit created with a tapped delay line and a multiplexer.

counter at f_{clk} precludes the use of low supply voltages, which would otherwise save power. Fig. 4 illustrates typical analog PWM generation, as well as the fast-counter digital PWM generation scheme.

Another way to create a pulse width modulated signal from an N-bit digital value is to use a tapped delay line. Since this approach uses the switching frequency clock, the power is significantly reduced relative to the fast-clocked counter approach. Fig. 5 shows a schematic for the proposed 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

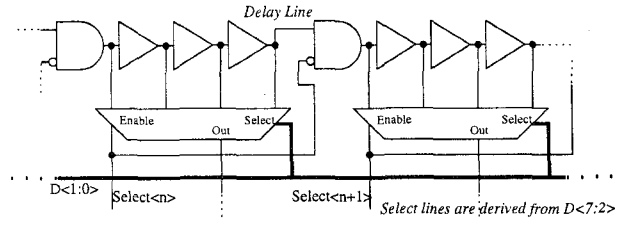

Fig. 6. Selection logic used to prevent propagation of input pulse. Outputs of 4x1 muxes are ORed together. Schematic is for a 256 stage delay line.

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. Multiple PWM signals can be generated by adding multiplexers to a single delay line, a feature which might be used to create a timing signal to turn off a synchronous rectifier.

In order to avoid superfluous switching activity (and the resulting power dissipation), the multiplexer is flat (as opposed to a tree structure), so that only the selected input propagates (causes switching transitions) on the outputs of the first level of gates. The flat multiplexer increases the energy dissipated when the selection signals change (relative to a more hierarchical, tree structured multiplexer), but due to the slow rate at which the duty cycle will change, this increase will not affect the power of the PWM circuit significantly. It is also possible to avoid switching transitions on the delay line buffers past the selected tap. As a consequence of the design and layout of the delay line and multiplexer, selection signals are available which indicate which block of four buffers contains the tap which the multiplexer is configured to select. These selection signals take the form of 64 (for a 256 stage delay line) "one-hot" signals, each available in the proximity of their respective set of buffers. These selection signals can be used to gate the pulse propagating down the delay line to save power. The pulse is prevented from propagating past the set of four buffers which is selected by the multiplexer (see Fig. 6), eliminating unnecessary switching events in the nodes downstream from the selected taps. Another signal is available to override this feature so that all pulses propagate to the end of the delay line. This is necessary to measure and control the total delay of the delay line.

B. Delay Line Length Control

The total delay of the delay line can be controlled if the buffers in the delay line are "starved," that is, the current available to switch their outputs is limited by a series MOS device in the subthreshold or linear region. A schematic of a starved buffer and measured data concerning the aggregate delay of a chain of such elements

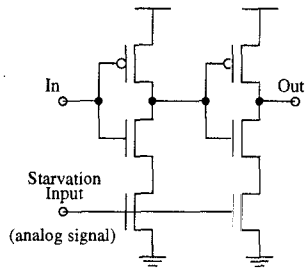

(a) Schematic of starved buffer.

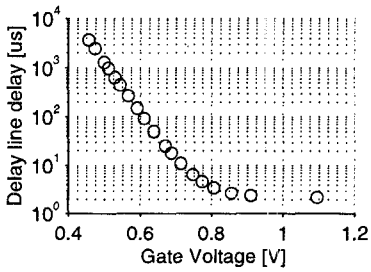

(b) Aggregate delay of 256 starved buffers vs. starvation input voltage.

Fig. 7. Circuit for variable delay buffer using starved inverters. The delay is controlled by the analog level on the gates of the lowest two NMOS devices. Measured delay of 256 stage chain shown in (b), where the supply voltage is $1.1V$.

are shown in Fig. 7. The delay decreases exponentially with the starvation gate input voltage for voltages below threshold (for the data shown, $V_{T_n} \approx 0.8V$), indicating that in this region, the starvation NMOS devices are operating in subthreshold conduction.

The voltage level on the starvation gates is controlled so that the total delay of the chain of buffers matches the period of the input reference clock (at the switching frequency). In order to match the delay line delay and the clock period, the time of arrival of a pulse to the last tap of the delay line is compared to the time of the next clock pulse. The control circuitry determines which pulse arrived first, then charges or discharges the starvation gate voltage until the other pulse arrives (see Fig. 8).

The control signals can modify the gate control voltage with the circuit of Figure 9. The gain of this control loop is determined by the charging resistance, R_p, the total capacitance on the signal node, $C_{gate} + C_B$, the supply voltage, and the operating point DC voltage on the output. C_{gate} is the distributed gate capacitance of the delay line, and C_B is a capacitor added for loop stabilization. Because the charging or discharging current depends on the voltage drop across R_p, the gain differs for positive and negative errors. The power dissipation

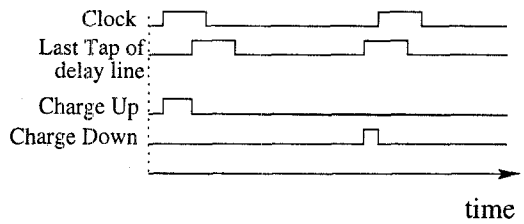

Fig. 8. Sample timing diagram showing generation of control signals.

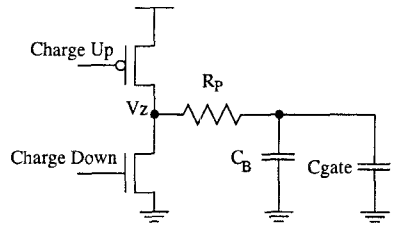

Fig. 9. Circuitry to charge rate control gate of delay line. C_{gate} is the distributed gate capacitance of the delay line, R_p, and C_B are added for stabilization.

of this feedback loop is small because only a few gates switch at each evaluation, and operation is at a fraction of the switching frequency.

For a fully integrated controller, it is prudent to have the "charge up/down" control signals gate a current source to increase or decrease the voltage on the output node. This configuration avoids the dependence of the gain on the supply voltage and the bias voltage, and avoids the problem of building either a large capacitor or a large resistor on chip, saving die area. For a practically sized integrated capacitor, the current source current will be very small (on the order of $0.1\mu A$ to $1\mu A$) so the additional power dissipation (from the current source) will be minimal.

Gardner [10] discusses the stability of phase locked loops controlled by a charge pump (the mechanism described above). Controlling a delay locked loop as proposed here is fundamentally less difficult, since a DLL does not integrate phase errors as a PLL. (A DLL has a lower order than a PLL.) In order to generate a model for the DLL, we consider it to be a discrete time system, and linearize about the operating point of the sub-threshold starvation devices:

$$D[n] = T_0 - Kv[n] \tag{1}$$

$$v[n+1] = v[n] - \frac{I_c}{C_B + C_{gate}}(T_0 - D[n]) \tag{2}$$

Here, $D[n]$ is the delay of the delay line in a given cycle, K is the incremental gain of the delay line in units *seconds/volt*, $v[n]$ is the small-signal value of the incremental rate control voltage, and I_c is the charging current. For the case shown in Figure 9, I_c will depend on

the power supply, the bias point on the rate control gate, the resistance R_p, and whether the delay error is positive or negative. The $\mathcal{Z}$-transform of this first order system yields the characteristic equation,

$$z = 1 - \frac{I_c}{C_B + C_{gate}} K, \tag{3}$$

and the system is stable if $-1 < z < 1$.

In the discussion of saving power by gating the propagation of the delay line pulse, it was assumed that the length comparison is not done every cycle. This property actually helps to avoid the problem that the delay of the delay line may be a multiple of the clock period. If the delay of the delay line is twice the clock period, pulses at the last tap may align perfectly with input clock pulses, and the control circuitry will not converge correctly. When the length comparison is not done every cycle, pulses are only allowed to propagate to the last tap of the delay line when a comparison is desired, and the threat of converging to the wrong delay length is reduced. The sample rate for the length comparison is determined by the leakage on the control gate, the tolerable ripple of the output duty cycle (caused by delay length jitter), and the necessary transient response for the delay length.

IV. Experimental Implementation

A. Design

To test and evaluate the low resolution feedback and digital PWM circuit discussed above, the circuits were fabricated on an integrated circuit, along with output switches to create a down converter. The fabrication was done on a 0.6μm CMOS process.

The intended application is a 5mW, 3V in, 1V out DC-DC converter, and the output switches were sized for this load. In order to give about 10mV resolution for the output voltage, a 256 tap delay line was used. (The output resolution is about $V_{IN}/\#$ Taps.) The dynamic comparator and digital up/down counter were also integrated on the chip, but the delay line length control and output L–C filter were implemented externally. The comparator compares the filtered output voltage to an external reference.

The output switches are a 2mm wide NMOS device and a 3.8mm wide PMOS device, constructed with many parallel fingers (see [1] for a discussion of the layout of CMOS power devices). The gates of these devices were driven by the same signal — there was no attempt to achieve zero voltage switching by careful timing of the gate signals.

The output filter for the down converter consisted of a 220μH inductor and a 0.22μF ceramic capacitor. The total filter board area is $0.024in^2$. This filter was sized to allow operation at a switching frequency around 330kHz with an output power of 5mW.

TABLE I
Summary of PWM and Power Converter Performance.

Minimum PWM V_{DD}:	1.1V
PWM Current with ($V_{DD} = 1.2V$, $f_{SW} = 330kHz$)	
D = 12%:	I=7.1μA
D = 50%:	I=8.8μA
Delay Range:	$< 1kHz$ to $1MHz$ ($V_{DD} = 2V$)
Efficiency with 5mW load:	88%
($V_{IN} = 3V$, $V_{OUT} = 1V$, $f_{SW} = 330kHz$, $L = 220\mu H$, $C = 0.22\mu F$, Filter area is $0.024in^2$)	

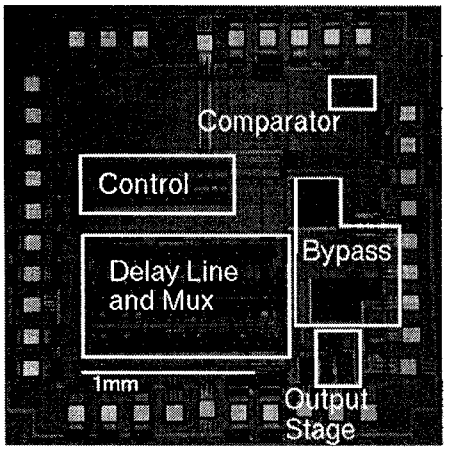

Fig. 10. Photograph of fabricated chip. Chip area is 2.3mm x 2.4mm, technology is 0.6μm CMOS.

B. Results

As shown in Table I, the power dissipation of the control circuitry is on the order of a mere 10μW at low supply voltages. This power varies depending on the switching frequency and the control circuit supply voltage. The power dissipation also depends on the duty cycle selected, since nodes downstream from the selected tap of the delay line are not activated, except during cycles where the delay line total delay is adjusted. For this data, half of all input pulses traverse the total length of the delay line, which explains why the power dissipation at 12% duty cycle is not one-quarter the dissipation at 50% duty cycle. With a $2V$ supply, the controller is suitable for switching frequencies as high as 1MHz, and the lower limit on the switching frequency is less than 1kHz (limited by jitter on the delay line length).

Fig. 11 shows the action of the delay line length control. For a clock frequency of 250kHz (shown in the figure), the delay length jitter is 40ns, or about 1% of the switching period. The magnitude of this jitter will increase if the

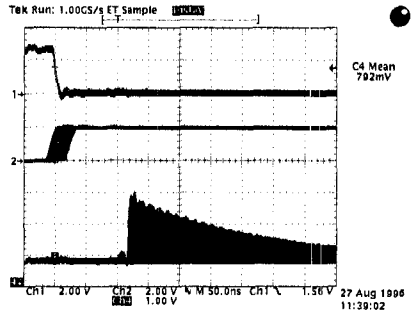

Fig. 11. Jitter of delay line delay. The top signal is an inverted version of the reference clock, and the second signal is the output of the last tap of the delay line (oscilloscope is in a long persistence mode). The bottom signal is node V_z of Fig. 9. There is some leakage from the control node to ground, so the delay control circuit generates occasional pulses to compensate.

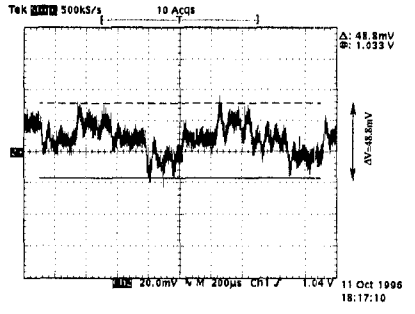

Fig. 12. Limit cycle behavior of system under closed loop feedback control, at $5mW$ output power. Arrow on left indicates $1.2V$ level.

delay length calibration is done less frequently. Here, the charge pump corrects for errors every fourth cycle, or $16\mu s$.

Fig. 12 shows the steady state limit cycle behavior of the closed loop system with low resolution feedback. In this case the switching frequency is 330kHz, and the duty cycle is incremented or decremented by one level every 75 switching cycles. The magnitude of the limit cycle ripple at steady state is 50mV, and the switching frequency ripple is 13mV. The response of the system to a step change in commanded voltage (not shown) is slew limited, since the duty cycle is only incremented by one step when it is in error, no matter how large the error. This slow behavior is not inherent to the delay line PWM, as illustrated in Fig. 13, where the selected duty cycle is being changed every cycle.

The low power of this controller enabled the construction of an 88% efficient power converter. It is interesting to note that the die area consumed by the power switches for a 5mW load is significantly smaller than the area of the circuit which controls those switches.

V. Conclusions

We have shown that it is possible to achieve stable operation with only a single bit of feedback regarding the output voltage, albeit with degradation of output response time. A circuit was presented which creates a pulse width modulated signal from digital inputs at frequencies appropriate for very small low power voltage converters, with a power dissipation on the order of 10μW.

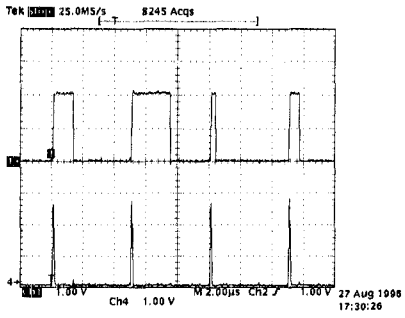

Fig. 13. Pulse width modulated signal, with clock. Duty cycle is varied between duty cycles 25%, 50%, 12%, and 37%.

Acknowledgments

This work is funded by DARPA contract #DAAL-01-95-K3526. Abram Dancy is supported by a NSF Graduate Research Fellowship. We would also like to recognize Rajeevan Amirtharajah for the dynamic comparator design.

References

[1] A. Stratakos, S. Sanders, and R. Broderson, "A low-voltage CMOS DC-DC converter for a portable battery-operated system," in *IEEE Power Electronics Specialists Conference*, 1994, pp. 619–626.

[2] Linear Technology, *LT1073 Datasheet, Micropower DC-DC Converter*, Available from `www.linear.com`.

[3] Maxim Integrated Products, *MAX 756 Datasheet, Adjustable-Output Step-Up DC-DC Converter*, January 1995, Available from `www.maxim-ic.com`.

[4] Anantha P. Chandrakasan and Robert W. Brodersen, *Low Power Digital CMOS Design*, Kluwer Academic Publishers, Norwell, Massachusetts, 1995.

[5] B. Gordon and T. Meng, "A 1.2mW video-rate 2-D color subband decoder," *IEEE Journal of Solid-State Circuits*, pp. 1510–1516, December 1995.

[6] Vincent Von Kaenel, Peter Macken, and Marc G. R. Degrauwe, "A voltage reduction technique for battery-operated digital systems," *IEEE Journal of Solid-State Circuits*, vol. 25, no. 5, pp. 1136–1140, October 1990.

[7] Vadim Gutnik and Anantha Chandrakasan, "An efficient controller for variable supply-voltage low power processing," in *Symposium on VLSI Circuits*, 1996, pp. 158–159.

[8] Gu-Yeon Wei and Mark Horowitz, "A low power switching power supply for self-clocked systems," in *1996 International Symposium on Low Power Electronics and Design*, 1996, pp. 313–318.

[9] T.B. Cho and P.R. Gray, "A 10-bit, 20-MS/s, 35-mW pipeline A/D converter," in *1994 IEEE CICC*, 1994, pp. 23.2.1–23.2.4.

[10] Floyd M. Gardner, "Charge-pump phase-lock loops," *IEEE Transactions on Communications*, vol. COM-28, no. 11, pp. 1849–1858, November 1980.