A Fully Digital, Energy-Efficient, Adaptive Power-Supply Regulator

Article *in* IEEE Journal of Solid-State Circuits · May 1999

DOI: 10.1109/4.753685 · Source: IEEE Xplore

2 authors, including:

Gu-Yeon Wei
Harvard University
234 PUBLICATIONS **7,701** CITATIONS

SEE PROFILE

A Fully Digital, Energy-Efficient, Adaptive Power-Supply Regulator

Gu-Yeon Wei and Mark Horowitz

***Abstract*—A voltage scaling technique for energy-efficient operation requires an adaptive power-supply regulator to significantly reduce dynamic power consumption in synchronous digital circuits. A digitally controlled power converter that dynamically tracks circuit performance with a ring oscillator and regulates the supply voltage to the minimum required to operate at a desired frequency is presented. This paper investigates the issues involved in designing a fully digital power converter and describes a design fabricated in a MOSIS 0.8-μm process. A variable-frequency digital controller design takes advantage of the power savings available through adaptive supply-voltage scaling and demonstrates converter efficiency greater than 90% over a dynamic range of regulated voltage levels.**

***Index Terms*—Adaptive control, dc–dc power conversion, frequency-locked loops, power supplies.**

I. Introduction

REDUCING power dissipation in digital integrated circuits is a key design constraint, and proliferation of battery-operated, portable applications further emphasizes the need to reduce power consumption [1]–[4]. Power dissipation in synchronous digital circuits is dominated by CV^2f dynamic power, so reducing the clock frequency of a synchronous digital circuit proportionally reduces the power, as shown by the dotted line in Fig. 1. Simply lowering frequency linearly reduces power but does not affect the energy consumed per cycle of operation. Since the delay of digital circuits is nominally related to the inverse of voltage, significant power and energy savings are possible if the supply voltage were to scale down with frequency [5]. The solid curve in Fig. 1 illustrates this optimal potential for reducing energy consumption. Adaptively regulating the supply voltage realizes this potential savings without sacrificing peak performance. This technique requires a feedback loop and supply-voltage regulator to dynamically set the supply voltage to the minimum required for any desired frequency of operation. This paper discusses a fully digital implementation of such an adaptive power-supply regulator.

For a digital system to operate with high energy efficiency and achieve the power savings discussed, three components are required: a mechanism for dynamically predicting circuit performance with respect to supply voltage and other operating condition variations, a control loop to adaptively regulate voltage, and an efficient power converter. The following section investigates how well the critical path of digital circuits tracks the delay of a simple inverter, which is used as an indicator of circuit performance. Then, we discuss efficient power-regulator design and explore control-loop issues in adaptively scaling supply voltage. Section IV describes a fully digital implementation of an adaptive supply-regulator controller, and Section V analyzes measured results from fabricated test chips. Measured data show that a digitally controlled adaptive supply regulator holds promise for low-power applications that require energy-efficient operation.

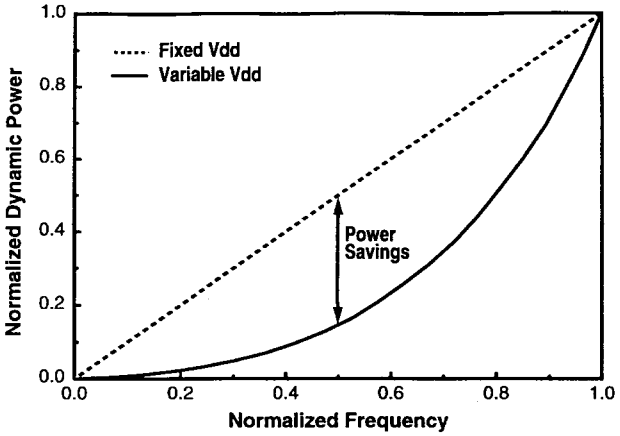

Fig. 1. Normalized power versus frequency.

II. Delay Matching

To dynamically measure digital circuit performance, a number of researchers have proposed using an inverter-based delay chain or ring oscillator to model the critical path of a circuit [5]–[8]. In digital CMOS circuits, the basic unit circuit element is an inverter. A fanout-of-4 (FO4) inverter, which uses an optimal fanout ratio to minimize delay for driving large loads [9], can be the basis for measuring the delay of various circuits. Simulated delays of various logic circuit elements and of a 64-bit adder's self-bypass path, presented in FO4 inverter delays, are plotted against voltage and temperature variations in Figs. 2 and 3, respectively. These plots reveal that basic static and dynamic digital circuits track well with the delay of an FO4 inverter. Therefore, the oscillation frequency of an FO4 inverter-based ring oscillator should accurately track the delay of the load circuit. Notice, however, that for this application, the oscillator frequency need not exactly match the desired frequency of operation. The critical path delay

Manuscript received February 16, 1998; revised August 20, 1998. This work was supported by the Advanced Research Project Agency under Contract J-FBI-92-194.

The authors are with the Computer Systems Laboratory, Stanford University, Stanford, CA 94305 USA.

Publisher Item Identifier S 0018-9200(99)02430-0.

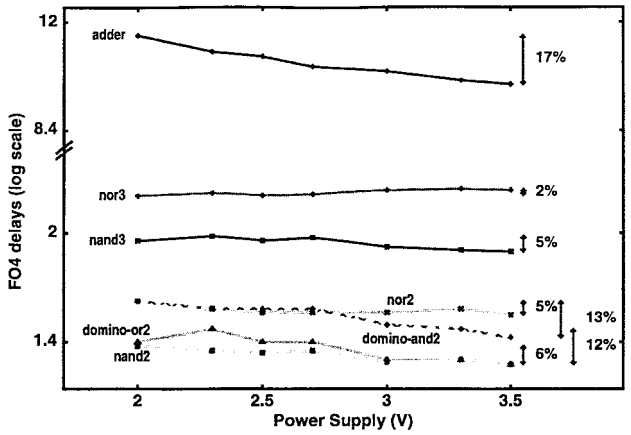

Fig. 2. Normalized gate delay versus $V_{\rm dd}$.

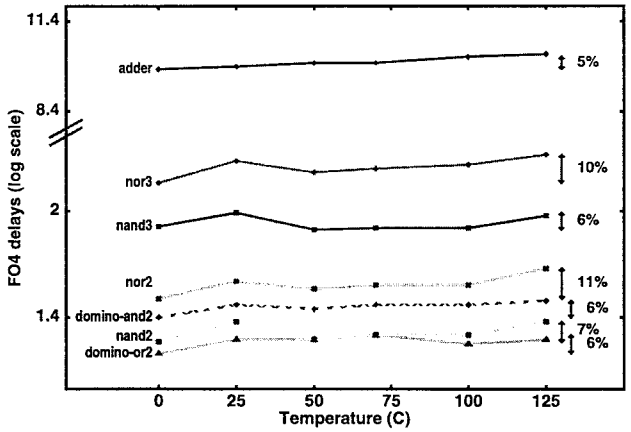

Fig. 3. Normalized gate delay versus temperature.

must simply be proportional to the delay through the ring. A gain factor can tune the output frequency to the frequency desired, and this factor can include enough overhead to account for any mismatches over temperature and voltage that may exist.

By fabricating the ring oscillator on the same die as the circuit it is intended to model, oscillator delay closely tracks the chip's critical path across process, temperature, and voltage variations. In the case of some memory circuits, dummy cell and delay paths, which are used for self-timing, may also be used to predict performance [10]. Examples of using tracking cells that replicate the gates, wires, and loading to construct local-clock generators have also been demonstrated in [7] and [8]. Although a replica of the functional unit's critical path most accurately models its delay, an exact match of the delay can be cumbersome and is not necessary. Rather, a mechanism for tracking delay variation with operating conditions is all that is required.

The principles of an adaptive voltage regulation technique were first demonstrated with a phase-locked loop (PLL) structure as described in [6]. A ring oscillator modeled the critical path delay. Since a linear regulator was used, a simple loop filter was sufficient for a stable configuration. Furthermore, since linear voltage regulation principally operates in the analog voltage domain, the loop utilized predominantly analog

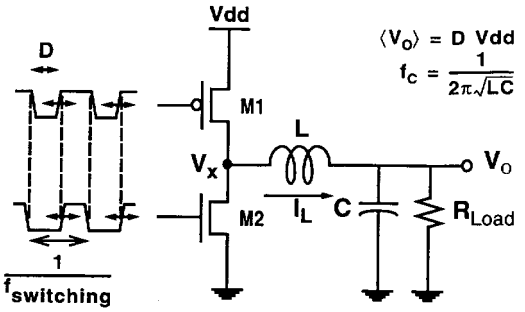

Fig. 4. Synchronous buck converter.

circuit components. While this chip successfully demonstrated a voltage-reduction technique for energy-efficient operation, overall efficiency for this type of implementation was limited by the poor converter efficiency of linear regulators and the static power invariably consumed by the analog circuits. Therefore, other power-conversion techniques must be investigated in order to achieve higher overall conversion efficiency.

III. Power-Regulator Design

Design of efficient power converters has been investigated by many researchers, and various methods for energy-efficient operation have been introduced [11], [12]. Conventional power-supply design concentrates on minimizing losses in order to efficiently deliver a fixed voltage over a wide range of load conditions. Furthermore, feedback loops are mainly employed to minimize output transients due to sudden load changes [13]. Different from fixed voltage regulation, however, adaptive voltage regulation requires a mechanism for detecting circuit delay variations due to variations in operating conditions and also for using that information to dynamically regulate the power-supply voltage. Since an inverter-chain ring oscillator demonstrates the ability to dynamically predict circuit performance, by placing it in the negative feedback path of a control loop that controls a power converter, adaptive supply-voltage scaling is achievable. The following subsection briefly reviews the switching power-supply converter used and subsequently explores loop-control design issues.

A. Synchronous Buck Converter

An integral component of adaptive power-supply design is the mechanism used to regulate the desired supply voltage. While linear regulators have been demonstrated in adaptive supply scaling applications, poor voltage regulation efficiency was a major limitation. Switching power-supply designs, however, exhibit exceptionally good efficiencies (greater than 90%) [13], [14]. Therefore, we employ a synchronous buck converter circuit shown in Fig. 4 for efficient power conversion. The regulated output of this type of converter is simply the average voltage of a pulse-width modulated (PWM) square wave at node V_x. Since ideally lossless reactive circuit components (L and C) acting as a low-pass filter do the averaging, power delivered to the load is close to the power pulled from the external supply (Vdd). The cutoff frequency (f_C) of the LC filter is chosen sufficiently below the switching

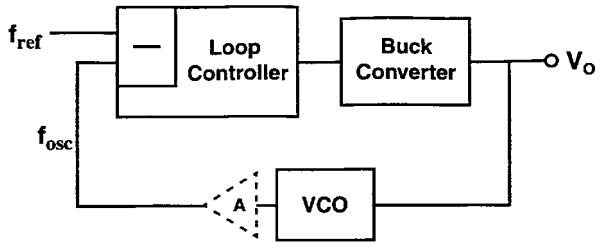

Fig. 5. Control-loop block diagram.

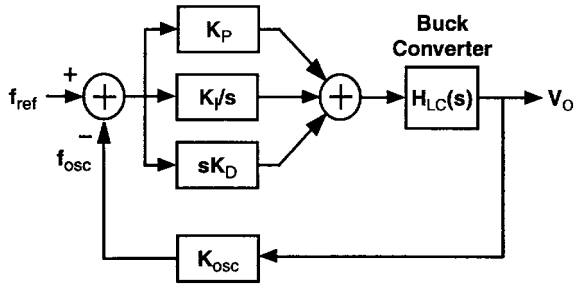

Fig. 6. Control-loop frequency-domain model.

frequency of the input pulses in order to keep the ripple on the regulated output below an acceptable level (less than 5%). Since this type of converter's performance has been thoroughly documented in [11]–[16], the reader is referred to these references for detailed description and analysis of this and other power-converter designs.

B. Control Loop

A voltage-controlled oscillator (VCO) in the negative feedback path of a closed loop monitors variations in circuit performance and adaptively scales the regulated voltage of the buck converter via a loop controller. A block diagram of the control loop is presented in Fig. 5. The feedback loop matches the oscillator frequency to the input reference frequency by adjusting the regulated supply voltage, the level required for the load circuit to operate at the desired frequency $(f_{\rm ref})$. The loop controller compares the frequency of the VCO with the reference. If the frequencies differ, the controller uses the detected error to appropriately adjust the input to the buck converter. The loop compensates for variations in temperature and voltage in order to guarantee proper functionality of the load circuit. Loop stability can be verified by analyzing the loop's frequency response.

A difficulty associated with designing this type of controller arises from the LC resonance of the buck converter. An open-loop frequency-response analysis exhibits a resonant peak at the cutoff frequency of the LC filter. A sharp peak, quantified by the quality factor (Q), is desirable for efficient power conversion. For simple integral control, this resonant peak must be kept below unity gain in the open-loop frequency response to ensure stability. Such a controller configuration has a low loop bandwidth and leads to slow transient response characteristics. A more complicated proportional, integral, and derivative (PID) control shown in Fig. 6, however, provides loop stability without sacrificing bandwidth and improves the loop's transient response. The proportional and derivative control blocks introduce compensation zeros that push unity-gain crossing beyond the resonant peak and eliminate the bandwidth limitation otherwise imposed by the resonant nature of the buck converter. The open-loop frequency-domain transfer function of the overall control is as follows:

$$H_{\text{open-loop}} = K_{\rm osc} H_{\rm LC}\left(K_p + \frac{K_l}{s} + sK_D\right)e^{-j\tau} \quad (1)$$

where the buck converter's transfer function (HLC) is given by the following equation:

$$H_{\rm LC} = \frac{A_{\rm Vdd}/LC}{s^2 + \left(\frac{R_s}{L} + \frac{1}{R_pC}\right)s + \frac{1}{LC}\left(1 + \frac{R_s}{R_p}\right)}. \quad (2)$$

All the resistive loss components (e.g., power transistor's "on" resistance) in series with the inductor are lumped together as R_s, and R_p is the effective parallel resistance of the load circuit and the nonideal resistance of the filter capacitor. A gain factor $A_{\rm Vdd}$ models the effect of the buck converter's external supply voltage variations on the loop's overall gain. The loop parameters are carefully designed with enough phase margin to account for the negative phase shift caused by delay (τ) through the loop, as indicated by the exponential term in (1), and a wide range of R_p values to account for different load environments.

Since the external supply is not always a fixed value, but can vary over time in many applications, it is important to briefly comment on how $A_{\rm Vdd}$ affects loop dynamics. The external input voltage to the buck converter has a $1:1$ correspondence to overall loop gain and therefore can constrain loop operation. As external supply diminishes, the loop gain can droop to a point where bandwidth falls below the converter's resonant peak and significantly slows down transient response. Frequency-response analysis in MATLAB shows the lower bound of $A_{\rm Vdd}$ at 0.65, 65% being the nominal external supply setting. While this lower limit constrains operating conditions, the loop can tolerate up to and beyond $A_{\rm Vdd} = 1.5$. To guarantee functionality, we can either require the external supply to always fall within tolerable limits or employ a detection circuit to monitor this gain factor and compensate with another gain stage incorporated into the loop.

In conventional voltage regulators, feedback loops are implemented with analog circuit blocks, which may take advantage of micropower techniques to minimize power consumption and achieve high conversion efficiency [13], [14]. The loop described above for adaptive supply regulation could also be fully implemented in analog, but would require a frequency-to-voltage converter. Since the inputs to the controller arrive as digital signals, however, a digital implementation is possible. We decided to explore this digital approach to better understand the tradeoffs in building a fully digital controller. These design tradeoffs are described next.

IV. Digital Control

Since the reference and feedback signals come into the controller as variable frequency clocks, both feed into counters,

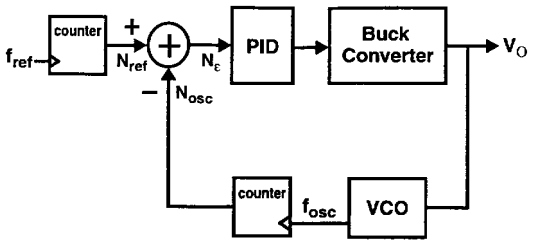

Fig. 7. Digital-loop architecture.

and the number of transitions are counted for a fixed period of time. Since the ring oscillator used to sense variations in circuit delay generates digital pulses with varying frequency, counting the pulses out of the oscillator over a known period results in a sampled average of the oscillation frequency. The difference between the outputs of the two counters corresponds to the error difference between the reference and oscillator frequencies in digital form. Since this adaptive regulator is intended for a digital system, the controller may be integrated on the same die as the load circuit and ideally also powered off the adaptively regulated supply in order to minimize its power consumption.

This digital controller approach takes an inherently analog loop and implements a sampled-data system with fully digital components that closely mimic the functionality of the underlying analog loop. In doing so, a general approach to adaptive supply scaling is achievable. Other digital approaches used specifically for digital signal-processing (DSP) applications have also been demonstrated [17], [18]. The digital controller described in [18] uses an open-loop approach to coarsely set the regulated voltage to discrete voltage levels optimized for the DSP core. It employs one feedback loop to adjust for variations in the operating conditions and another to dampen out the output voltage transients that result from sudden changes in performance and frequency requirements. This approach for adaptive supply-voltage scaling, however, does not provide a continuous range of regulated voltage levels and requires analog blocks to implement the damping circuit. To target general-purpose processor applications, the only assumption made on the functionality of the load circuit is that it requires a granular variable-frequency clock and operates off the adaptively regulated supply.

A block diagram of a control-loop architecture for such a general-purpose system that can regulate voltage with fine resolution and be implemented completely in digital is illustrated in Fig. 7. A binary equivalent of the error between the reference and oscillator frequencies, N_ϵ, feeds into the PID control block, which in turn uses the error value to make appropriate corrections to the output. The output of the buck converter is therefore adaptively regulated to the required voltage for the load circuit to operate at the desired reference frequency. The ring oscillator in the feedback path operates off the regulated supply voltage and generates a frequency that predicts circuit performance variations due to changes in operating conditions to complete the loop. The digital implementation of the PID blocks mimics equivalent analog blocks with the following proportional, integral, and derivative

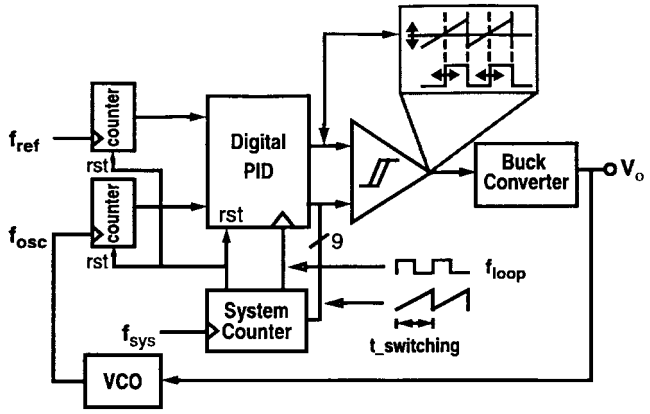

Fig. 8. Digital fixed-frequency-loop block diagram.

control equations:

$$y_P(n) = K_P x(n) \tag{3}$$

$$y_i(n) = y_I(n-1) + \frac{1}{K_I} x(n) \tag{4}$$

$$y_D(n) = K_D[x(n) - x(n-1)]. \tag{5}$$

$x(n)$ represents the error N_ϵ between the reference and oscillator frequencies. To reduce hardware complexity and power, all coefficients K_P, K_I, and K_D are binary factors and are implemented with shifters. Thus, a series of shifters, adders, and latches implement each of the above equations. The sum of (3)–(5) corresponds to the binary digit equivalent of the duty cycle for the PWM square wave input to the buck converter at time step n.

The system clock frequency, called f_{loop}, clocks the digital controller and sets the time base for the overall system. Over a period of f_{loop}, counters count the reference and oscillator clock pulses to generate the equivalent numeric representations of the frequency, and it is over this same period that the variable duty cycle square wave pulse into the buck converter is switched. Acting as a time base for the controller, the loop parameters of the digital PID blocks are proportional to the loop frequency. The cutoff frequency for the LC filter of the buck converter, however, is fixed independent of loop frequency. Therefore, in order to maintain a stable control loop over the dynamic range of regulated voltage levels, a fixed loop frequency is required. The most obvious approach is to use a fixed frequency controller. As we will see in the following subsection, this results in significant power overhead. Although a simple approach to reducing this overhead is to reduce the resolution and operating frequency, only proportional savings is achievable. An interesting approach that uses a variable-frequency controller is then described that requires a frequency-shifting technique in order for overhead power to track frequency and significantly improves overall power conversion efficiency.

A. Fixed-Frequency Control

An external frequency generates the loop frequency through the system counter as shown by the block diagram in Fig. 8. The output and reference voltages are converted to equivalent

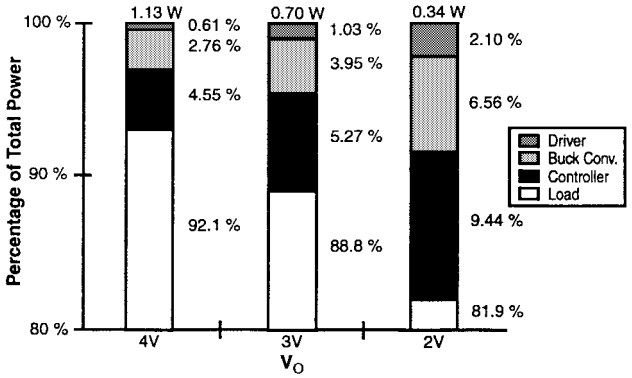

Fig. 9. Fixed-frequency controller power breakdown.

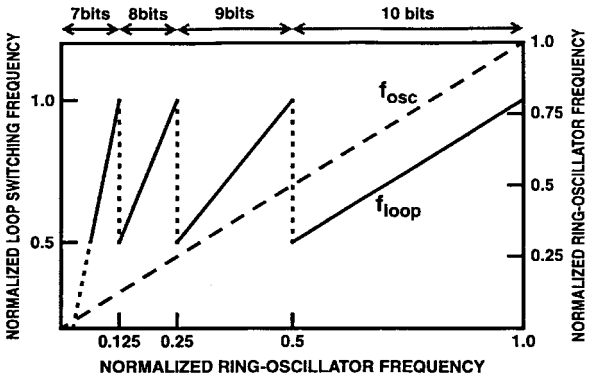

Fig. 10. Variable-bit counter frequency.

9-bit binary values by counting the pulses out of the ring oscillators over a fixed sampling period set by f_{loop}. A fixed external clock f_{sys} feeds into the synchronous counter where the highest order bit sets the loop sampling frequency and the switching frequency of the square wave input to the converter. An effective digital-to-analog (D/A) conversion is implemented by comparing the numerical value of the duty cycle, calculated by the PID block, to the output of the free-running synchronous counter. The comparator generates a fixed-frequency square wave, and the constantly updated PID block output proportionally modulates the duty cycle of the square wave input to the buck converter.

A detailed analysis of the fixed-frequency controller's performance can be found in [15], but results are best summarized by Fig. 9. A decomposition of different components of power consumed as a percentage of the total power is shown. The graph illustrates that the percentage of overhead power consumption increases and conversion efficiency degrades as regulated voltage (V_O) decreases. The largest component of overhead power is due to the controller. While parts of the controller operating at the slower loop frequency could be powered off the regulated voltage, components operating at a high fixed frequency (sys_clk) require a fixed voltage and consume a fixed amount of CV^2f power. Reducing counter resolution and utilizing a slower system clock proportionally reduces the controller's power, but some fixed power overhead is unavoidable. With variable frequency and voltage operation, however, this power can be significantly reduced to be only a small fixed percentage of the load's power.

B. Variable-Frequency Control

The key to improving efficiency while maintaining loop stability is to allow the internal frequency and voltage of the controller to change while keeping the controller's loop parameters relatively constant. Since these parameters are proportional to the update rate (or loop frequency) of the controller, it must be kept relatively constant (within a factor of two). To accomplish this, a frequency detection circuit is required. In the previous design, the switching frequency was set by the binary ramp wave, generated by a 9-bit synchronous system counter, which was clocked with an external fixed frequency. In a variable-frequency implementation, the system counter is clocked by the variable frequency of the VCO or reference frequency, and the number of bits up to which the counter counts is also variable. A system counter that can count up to a maximum of seven to ten bits, determined by a frequency-detect circuit and a variable-frequency clock, is used. Fig. 10 plots the variable ring-oscillator frequency and resulting loop frequency generated by the counter in order to graphically illustrate frequency shifting. The controller monitors the system counter's input clock frequency and adaptively determines the number of bits up to which the counter needs to count in order to maintain a switching frequency that maximally varies by a factor of two. For example, when the oscillator frequency is half the maximum rate, the counter will only use nine bits, restoring the loop switching frequency to the maximum rate. To further limit the loop parameters' excursions with frequency, the PID blocks' gain coefficients, implemented with shifters and adders, dynamically adjust within this factor-of-two variation in loop frequency. By incorporating these two methods into the controller, a stable configuration for the loop can be achieved. The simulated open-loop frequency response of Fig. 11 verifies stability. The two curves show that even at the extreme ends of the frequency excursion, the loop exhibits sufficient margin for stability. As mentioned in Section III, reducing A_{Vdd} shifts down the magnitude plots, and the lower bound at 0.65 corresponds to when open-loop magnitude crosses unity gain at a frequency lower than the resonant peak. Fig. 12 presents the resulting behavioral-model simulation of the loop's transient response to step changes in the reference.

Given a variable frequency controller powered off the regulated supply, it is possible to keep the controller's power-consumption overhead at a fixed percentage of the total power consumed. Components of the controller that operate at the switching frequency can also be powered off the regulated supply as long as the timing requirements of this much slower frequency can be met at the lowest regulated voltage level. Since this controller's power is nominally proportional to V^2f, the power consumed should be on the order of 1 mW at 1.5 V and should no longer be the dominant limitation for efficiency due to fixed overhead power observed in the fixed-frequency implementation.

Using a ring-oscillator and counter combination has been the prevalent method for measuring variations in circuit delay performance in the presence of variations in operating

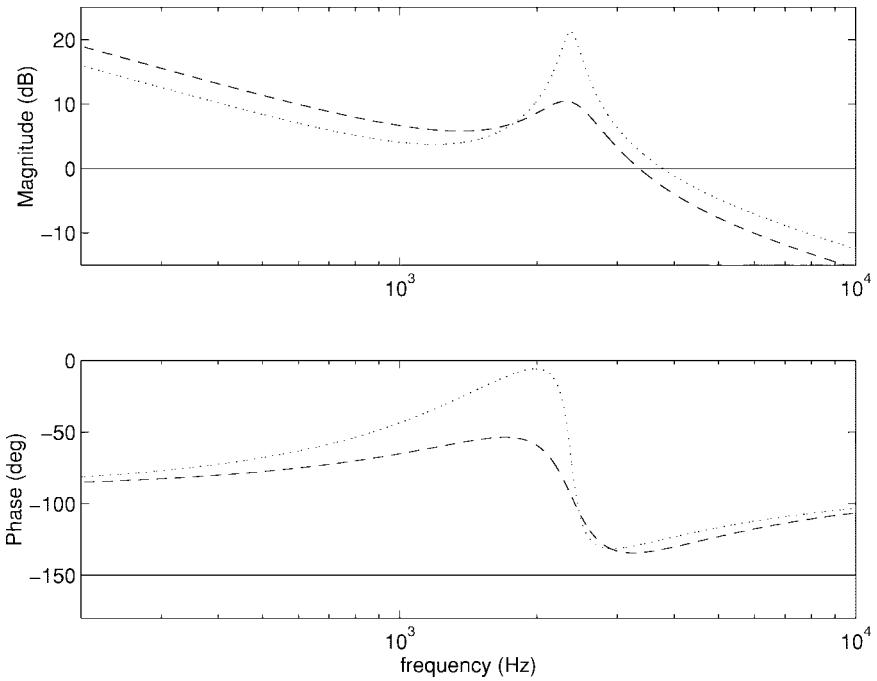

Fig. 11. Variable-frequency control open-loop frequency response.

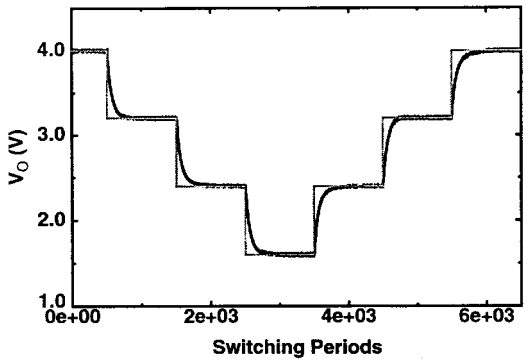

Fig. 12. Variable-frequency control closed-loop transient response.

condition of digital adaptive voltage-scaling controllers. The ring oscillator and counter, however, exhibit a high activity factor such that power consumption can be considerable. A ring-oscillator and counter pair also requires additional logic to detect the error between the desired and regulated voltage levels. Therefore, other methods of predicting and comparing performance should be considered. Techniques that compare the input and output pulses through a delay chain as described in [19] and [20] may potentially reduce power. Reference [19] has the advantage of utilizing fewer circuit elements and less dynamic power, but requires analog circuit blocks and is not suitable for a fully digital implementation. Reference [20] implements a simple bang-bang control that uses a critical path replica to determine the actual performance of the RISC core. This requires the replica path to accurately represent the worst case delay across all possible data, processing, and operating conditions. To efficiently generate the PWM control signal for the buck converter, [21] describes a clever technique that utilizes a tapped delay line to implement the effective D/A conversion with low power overhead. Total delay through the delay line matches the switching period, and the digital number that corresponds to the pulse width sets the fractional delay through the tapped delay line. A set–reset flip-flop constructs a square wave with the desired duty cycle by using the appropriate tap digitally selected along the delay line. This technique significantly reduces power consumption since it obviates the need for the highly active system counter and comparator operating at high frequencies. The tapped delay line still requires a digital value for the duty cycle, and so this method can be utilized with the rest of the digital controller described thus far. Since reducing power consumption is the key design goal, other digital techniques that further improve converter efficiency are investigated next.

C. Minimizing Switching Losses

Although variable frequency control yields an energy-efficient digital controller, overhead associated with the buck converter losses and buffer drivers limits efficiency. Since the optimal sizing for the power transistors of the buck converter depends on the power's being delivered to the load, on-chip segmented power transistors have been implemented. Assuming a cubed dependence of load power on voltage, by progressively turning off sections of the power transistor for lower levels of regulated voltage, the transistor width is adaptively set closer to the optimum, and the switched capacitances of the buck converter's power transistors and buffers also adaptively vary. Thus, power consumed by the buck converter is no longer fixed, but decreases for lower levels of regulated voltage and further improves efficiency. As a result, high conversion efficiency is achievable over a dynamic range of voltages.

Maintaining high conversion efficiency is especially challenging for low load current conditions, but additional op-

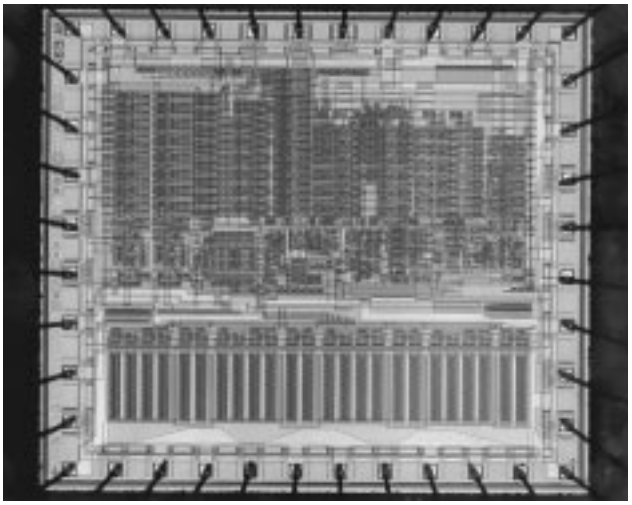

Fig. 13. Variable-frequency test chip die photo.

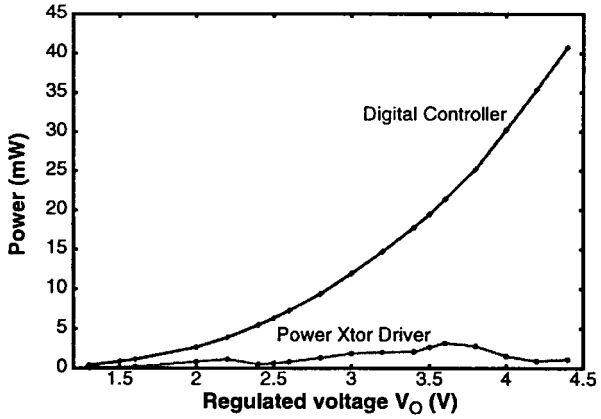

Fig. 14. Measured power components.

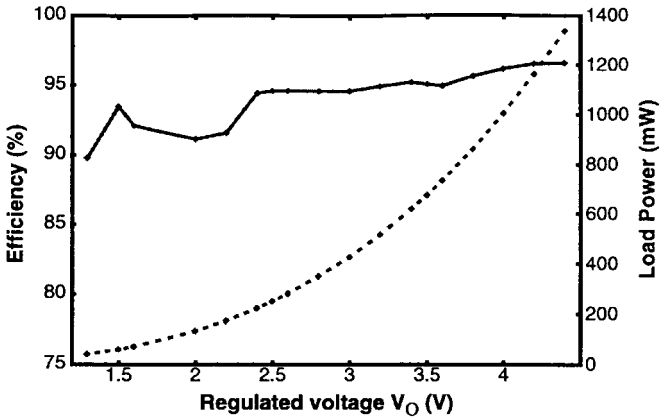

Fig. 15. Power-conversion efficiency versus regulated voltage.

timizations further improve efficiency. The buck converter consumes a minimum energy each cycle set by the CV^2 energy required to switch the power transistors' gates and parasitic drain capacitances and the resistive losses dissipated to support the average load and ripple filter currents. Under low load conditions, ripple current through the inductor cycles from the filter capacitor back to the supply and dissipates power as it flows through the power transistor switches. This cycling current, however, does not contribute to the overall conversion effort to supply power to the load. Therefore, minimizing this component of energy overhead improves efficiency. Taking advantage of the relative ease with which nonlinear methods can be implemented in digital controllers, a pulse-squashing technique eliminates this cycling current. Additional control logic detects current circulating back to the supplies from the filter capacitor and squashes subsequent pulses to the power transistors until the regulated voltage falls below a specified tolerance. Current polarity through the inductor is detected at the end of each switching period by momentarily turning off both NMOS and PMOS power transistor switches and sensing the drain voltage. Under such low load conditions, the loop enters a nonlinear mode of operation and the buck converter regulates discontinuously. In addition to eliminating losses attributed to the circulating current, the buck converter's switching frequency effectively decreases and reduces the overhead associated with switching the power transistors.

V. Measurement and Analysis

A digitally controlled power converter with variable-frequency operation was fabricated in a MOSIS 0.8-μm technology. A die photo of the prototype chip is presented in Fig. 13. The controller was tested with a 15-stage ring oscillator, which oscillates at a peak frequency of 116.5 MHz at 5 V, and the loop frequency maximally varied from 56.9 to 113.8 kHz. The converter's external LC-filter component sizes were 70 μH and 70 μF and powered off a 5-V external supply. A cubic dependence of the controller's power consumption on voltage is apparent from the plot of measured power versus regulated voltage in Fig. 14. Since the digital controller is also powered off the regulated voltage, the controller's power tracks the CV^2f power of the load circuit and is no longer the limiting component of the efficiency, even at low regulated voltage levels—the major limitation in the fixed-frequency implementation. Although this controller's absolute power consumed can be considered high [18], the intended target of having it track the load's cubic dependence on voltage has been met. Controller power is further reducible by implementing a more power-efficient PWM generator [21] and by running the converter at a slower frequency, but at the cost of slower transient response. The digital controller consumed 12 mW at 3 V with the configuration described above. For comparison, the controller's dynamic power extrapolates down to 4.4 mW at 1.3 V for a converter implemented in a 0.25-μm technology and powered off a 2.5-V external supply. Fig. 15 plots the conversion efficiency for a configuration with a CV^2f load targeted to consume 1 W at 4 V. Variable frequency control demonstrates high conversion efficiency over a wide range of regulated voltages. Improved process technology and other low-power circuit techniques can further reduce the power overhead so that digitally controlled power converters with very high efficiencies are achievable.

High conversion efficiencies demonstrated over a dynamic range of regulated voltages rely, however, on consistent and appreciable load power consumption. Furthermore, due to some fixed overhead power, high conversion efficiency is

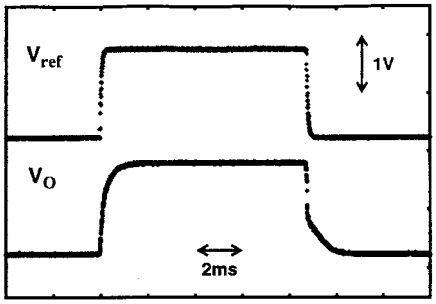

Fig. 16. Measured voltage-transient response.

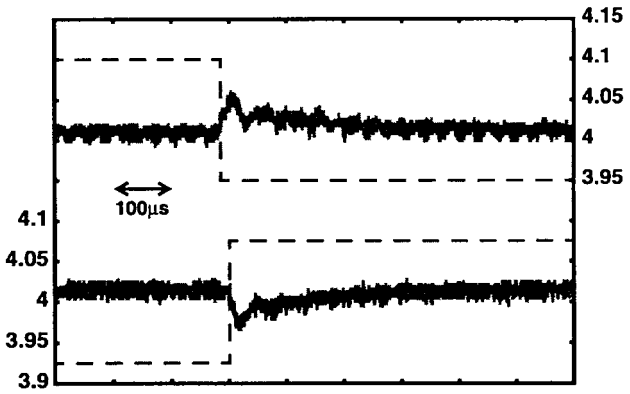

Fig. 17. Measured load-transient response.

difficult to guarantee for extremely low-load power conditions. One example is if the load circuit enters sleep mode; it consumes a negligible amount of power. Although discontinuous operation with pulse squashing reduces switching regulator power, the controller dissipates overhead power. Active control periodically supplies packets of charge to compensate for leakage losses and monitors when the circuit comes out of sleep mode. Again, the digital controller's power limits efficiency. To abate this condition, the load's sleep-mode signal could trigger the controller to also enter into its own form of sleep mode. Since very low power consumption is anticipated during sleep mode, the filter capacitor discharges slowly and the controller can operate at a slower rate. By waking periodically (one-tenth the normal operating frequency), the controller can sustain the required voltage level but consume significantly less power. Employing these low-power techniques leads to additional power savings and improves light load converter efficiency.

The adaptive power supply's response to voltage and load transients is shown in Figs. 16 and 17, respectively, which plot data downloaded from a sampling digital oscilloscope (HP54601A). To facilitate voltage-transient measurements, an additional ring oscillator was fabricated on the controller chip, which operates off a separate reference voltage ($V_{\rm ref}$ to generate the reference frequency used by the digital controller. A regulated voltage that closely tracks a step change in the desired reference is observed for the tested configuration in Fig. 16. Since a frequency-shifting boundary exists halfway through the voltage transient, the loop frequency changes across that boundary and results in different loop bandwidths for the falling and rising transitions. Fig. 17 presents the loop's response to load current transients where dotted lines delineate the step changes in the load current between 0 and 100 mA at 4 V. Minimizing the magnitude of the controller's response to load transients is a key design constraint since performance overhead margins incorporated into the model of the load circuit's critical path must account for these worst case voltage deviations.

VI. Conclusion

To achieve energy-efficient digital circuits operating at the highest speed-power point, a fully digital controller design for adaptive supply-voltage regulation has been described and measured results from test silicon presented. A digital controller that can be embedded within the digital system to which power is being supplied has the advantage of tracking its own power consumption with that of the rest of the system. The first test silicon demonstrated the feasibility of a fully digital implementation, but the efficiency of fixed-frequency operation was limited by fixed overhead power consumption. To achieve high conversion efficiency over a wide range of voltage and power, a variable-frequency controller was designed and fabricated. By dynamically adjusting the switching frequency and gain coefficients to maintain relatively constant loop parameters, a stable loop was achieved. Since the design is fully digital, the variable-frequency controller can also take advantage of the power-savings potential available through adaptive supply regulation. Thus, a digital controller whose power consumption is no longer fixed but nominally tracks that of the load circuit has been demonstrated, and high conversion efficiency over a dynamic range of regulated voltage levels has been achieved. The measured controller's power is nonnegligible, however, and efficiency may degrade for very low-load power usage. To improve efficiency under these conditions, implementing an additional power-saving sleep mode is a viable solution. Other lower power techniques for generating the PWM signal can also further improve efficiency. Therefore, this fully digital technique holds promise as a controller for adaptive supply-voltage regulation in fully digital system applications that present a hostile environment for noise-sensitive analog circuits.

Acknowledgment

The authors would like to thank B. Amrutur, K. Yang, S. Sidiropoulos, R. Ho, D. Harris, and D. Weinlader for invaluable discussions and insight.

References

[1] M. Horowitz, T. Indermaur, and R. Gonzalez, "Low-power digital design," in *IEEE Symp. Low Power Electronics Dig. Tech. Papers*, Oct. 1994, pp. 8–11.
[2] T. D. Burd and R. W. Broderson, "Processor design for portable systems," *J. VLSI Signal Process.*, vol. 13, no. 2–3, pp. 203–221, Aug.–Sept. 1996.
[3] A. P. Chandrakasan, S. Sheng, and R. W. Broderson, "Low-power CMOS digital design," *IEEE J. Solid-State Circuits*, vol. 27, pp. 473–484, Apr. 1992.
[4] R. Gonzalez, "Low-power processor design," Ph.D. dissertation, Stanford University, Stanford, CA, June 1997.
[5] M. Horowitz, "Low power processor design using self-clocking," presented at the Workshop on Low-Power Electronics, 1993.

[6] P. Macken, M. Degrauwe, M. Van Paemel, and H. Oguey, "A voltage reduction technique for digital systems," in *IEEE ISSCC 1990 Dig. Tech. Papers*, Feb. 1990, pp. 238–239.
[7] M. Santoro and M. Horowitz, "SPIM: a pipelined 64 $*$ 64-bit iterative multiplier," *IEEE J. Solid-State Circuits*, vol. 24, pp. 487–93, Apr. 1989.
[8] M. Dean, "STRIP: A self-timed RISC Processor," Ph.D. dissertation, Stanford University, Stanford, CA, July 1992.
[9] L. Gal, "Reply to 'Comments on the optimum CMOS tapered buffer problem,'" *IEEE J. Solid-State Circuits*, vol. 29, pp. 158–9, Feb. 1994.
[10] B. Amrutur and M. Horowitz, "Techniques to reduce power in fast wide memories," in *IEEE Symp. Low Power Electronics Dig. Tech. Papers*, Oct. 1994, pp. 92–93.
[11] A. J. Stratakos, "High-efficiency low-voltage DC-DC conversion for portable applications," Ph.D. dissertation, University of California at Berkeley, Dec. 1998.
[12] J. G. Kassakian and M. F. Schlecht, "High-frequency high-density converters for distributed power supply Systems," *Proc. IEEE*, vol. 76, no. 4, Apr. 1988.
[13] D. Monticelli, "Switching power supply design," in *Proc. IEEE Symp. Low Power Electronics*, Oct. 1995, p. 64.
[14] Maxim Integrated Products, Inc., *1995 New Releases Data Book*, CA, vol. IV, 1995.
[15] G. Wei and M. Horowitz, "A low power switching power supply for self-clocked systems," in *IEEE Int. Symp. Low Power Electronics and Design Dig. Tech. Papers*, Aug. 1996, pp. 313–317.
[16] A. P. Charandrakasan and R. W. Broderson, *Low Power Digital CMOS Design*. Norwell, MA: Kluwer Academic, 1995.
[17] L. Nielsen, C. Niessen, J. Sparso, and K. van Berkel, "Low-power operation using self-timed circuits and adaptive scaling of supply voltage," *IEEE Trans. VLSI Syst.*, vol. 2, pp. 391–397, Dec. 1994.
[18] V. Gutnik and A. P. Chandrakasan, "An efficient controller for variable supply-voltage low power processing," in *IEEE Symp. VLSI Circuits Dig. Tech. Papers*, June 1996, pp. 158–159.
[19] S. Sakayima, H. Nakahira, M. Fukuda, A. Yamamoto, M. Kinoshita, A. Matsuzawa, H. Yamamoto, Y. Kato, Y. Matsuya, S. Mutoh, H. Fukuda, Y. Nishino, and T. Sakurai, "A lean power management technique: The lowest power consumption for the given operating speed of LSI's," in *IEEE Symp. VLSI Circuits Dig. Tech. Papers*, June 1997, pp. 99–100.
[20] K. Suzuki, S. Mita, T. Fujita, F. Yamane, F. Sano, A. Chiba, Y. Watanabe, K. Matsuda, T. Maeda, and T. Kuroda, "A 300 MIPS/W RISC core processor with variable supply-voltage scheme in variable threshold-voltage CMOS," in *Proc. IEEE Custom Integrated Circuits Conf.*, May 1997, pp. 587–590.
[21] A. Dancy and A. P. Chandrakasan, "Techniques for aggressive supply voltage scaling and efficient regulation," in *Proc. IEEE Custom Integrated Circuits Conf.*, May 1997, pp. 579–586.

Gu-Yeon Wei was born in Seoul, Korea, in 1972. He received the B.S. and M.S. degrees in electrical engineering from Stanford University, Stanford, CA, in 1994 and 1997, respectively, where he is currently pursuing the Ph.D. degree in the Computer Systems Laboratory.

His research interests are in low-power and high-performance circuits and systems design, with specific interest in energy-efficient adaptive supply regulation and low-power clocking techniques.

Mark Horowitz received the B.S. and M.S. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, in 1978 and the Ph.D. degree from Stanford University, Stanford, CA, in 1984.

He is the Yahoo Founders Professor of Electrical Engineering and Computer Science at Stanford University. His research area is in digital system design, and he has led a number of processor designs, including MIPS-X, one of the first processors to include an on-chip instruction cache; TORCH, a statically scheduled, superscalar processor; and FLASH, a flexible DSM machine. He has also worked in a number of other chip design areas, including high-speed memory design, high-bandwidth interfaces, and fast floating point. In 1990, he took leave from Stanford to help start Rambus, Inc., a company designing high-bandwidth memory interface technology. His current research includes multiprocessor design, low-power circuits, memory design, and high-speed links.

Dr. Horowitz received a 1985 Presidential Young Investigator Award and an IBM Faculty Development Award, as well as the 1993 Best Paper Award at the International Solid State Circuits Conference.