A Fully Integrated On-Chip DC–DC Conversion and Power Management System

George Patounakis, *Student Member, IEEE*, Yee William Li, *Student Member, IEEE*, and Kenneth L. Shepard, *Senior Member, IEEE*

***Abstract*—It is widely recognized that adaptive control of the power supply is one of the most effective variables to achieve energy-efficient computation. Most on-chip dc–dc conversion systems have relied on buck converters with off-chip *LC* filters. In this paper, we describe the development of a software controllable, fully integrated on-chip dc–dc downconversion system that combines switched-capacitor voltage dividers and linear regulators to efficiently regulate from 2.5 V down to about 0.65 V. The use of switched-capacitor supplies offers better efficiencies than what is achievable with linear regulators alone.**

***Index Terms*—DC-to-DC converter, low power, power management, regulator.**

I. Introduction

POWER consumption has become one of the most important issues in processor design, not only in portable, battery-powered applications, but in high-performance desktop and server applications because of packaging and cooling requirements. Dynamic (or adaptive) voltage scaling (DVS) has been widely studied [1]–[3] and is being implemented commercially as one of the most effective means of achieving energy-efficient design. A given computation proceeds in the most energy-efficient manner when the supply voltage is scaled to the point of "just-in-time" operation.

Most DVS systems are based on the idea that multiple power grids are available to be "tapped into" to support multiple voltage operation, which comes at the cost of additional complexity and area [4]. Entire design methodologies have been developed around such a concept of voltage islands [5]. An alternative to a set of externally generated fixed voltage supplies that are switched into on-chip voltage domains is to provide for dynamic dc–dc conversion for each island, which would allow for continuous scaling and negate the need for multiple global power grids. The most efficient techniques for dc–dc downconversion are based on *buck* converters, which essentially filter a pulsewidth modulated (PMW) signal through an *LC* network to achieve a downconverted dc voltage [6]–[8]. Efficiencies of 80% to 90% can be easily achieved in these systems, although off-chip inductors and (usually) off-chip capacitors are needed. This need increases the pin requirements, reduces efficiency, and makes "fine-grain" voltage domains impractical and expensive.

The simplest on-chip dc–dc conversion scheme would be one using a single linear regulator to regulate over the entire desired voltage range. In addition to having poor "ideal" efficiencies at low voltages, the quiescent currents required to achieve the requisite high-voltage bandwidths result in further efficiency degradation at low voltages. In this paper, we address these efficiency issues by employing a "hybrid" voltage regulator approach, regulating down from an on-chip $V_{DD}/2$ switched-capacitor voltage supply for low-voltage conversion and using a separate linear regulator for dc–dc conversion at high supply voltages. The system is software-controlled with a dynamic control system to automatically adjust the supply voltage of a domain of digital logic to just meet specified performance requirements. An on-chip flash analog-to-digital converter (ADC) allows noninvasive transient monitoring of the regulated power supply to test the power management system functionality.

In Section II, we describe the system-level design of the power management unit, its integration into a prototype digital signal processor, and the associated design specifications for the dc–dc conversion and power-supply regulation. Section III provides an overview of the hybrid dc–dc conversion system. Design issues associated with the important constituents of the system—the high-voltage regulator, the watchdog, and the low-voltage regulator—are described in detail. Measurement results are presented in Section IV. Section V concludes.

II. Power Management System

The on-chip power management unit described here was developed to dynamically scale the supply voltage of an asynchronous micropipelined datapath as part of a prototype stream-based signal processing engine which is described in detail elsewhere [9]. The instruction word in this processor specifies a required throughput performance for the datapath and, in general, specifies a single-instruction multiple-data (SIMD) operation on a large data block. The power management system, as shown in Fig. 1, scales the supply for the datapath to just meet this performance requirement and is designed to robustly operate during the voltage-scaling transients. The "replica slice" in Fig. 1 tracks the performance of the asynchronous logic and consists of a ring structure configured to oscillate with a period approximately four times the datapath circuit cycle time. An asynchronous counter in the power management controller is used to count the "ticks" of this replica clock.

Manuscript received May 6, 2003; revised October 1, 2003. This work was supported by the National Science Foundation under Grant CCR-00-86007 and by the MARCO/DARPA C2S2 Center under subcontract from Carnegie-Mellon University.

The authors are with the Columbia Integrated Systems Laboratory, Department of Electrical Engineering, Columbia University, New York, NY 10027 USA (e-mail: shepard@ee.columbia.edu).

Digital Object Identifier 10.1109/JSSC.2003.822773

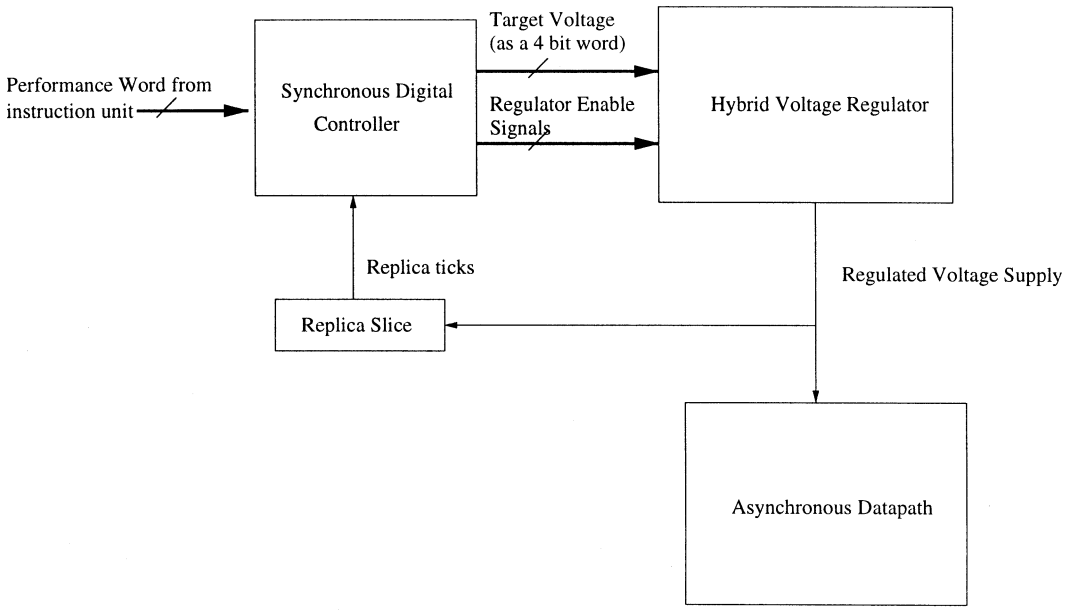

Fig. 1. Top-level diagram of the power management system.

Using this counter gauge as input, the digital controller (see Fig. 1) is responsible for adjusting the voltage to meet the desired performance target as specified by the instruction word, making smooth transitions across regulator boundaries and selecting the appropriate regulator for the target voltage. The controller uses a monotonic linear search on eleven voltage steps of approximately 150 mV between 2.5 V and 650 mV starting from the current voltage to find the voltage yielding the target performance. If the desired performance is greater than (less than) the current performance, the controller increases (decreases) the target voltage. An upward (downward) search terminates when the performance is greater than (less than) the desired performance. Once a downward search terminates, the voltage level is increased by one code word to guarantee the performance target is met. The controller remains idle and the target voltage constant until the instruction unit changes the desired performance. The monotonicity of the controller eliminates the possibility of instability and saves power because the controller only adjusts the voltage when necessitated by a change in the performance as specified by the instruction unit. This "search once" control algorithm does not allow the voltage to adjust to self-heating effects, but these can be easily accommodated by margins in the software-specified performance targets. Running at 10 MHz, the controller consumes an average power of 470 μW with a 200-MHz replica oscillator output (typical of 2.5-V operation of the datapath) and 65 μW with a 20-MHz replica oscillator output (typical of 650-mV operation of the datapath).

In addition to the target voltage code, the controller also supplies a set of control signals (labeled "regulator enable signals" in Fig. 1) to the hybrid voltage regulator. These include enable signals to indicate which of the regulators should be engaged to supply current to the load for the target voltage value. The low and high constituent regulators also receive two other control signals—**turbo** and **warmup**—which are used to improve regulator response during the performance search without requiring larger quiescent currents, thereby improving average efficiency. These circuit "knobs" exploit information known at the system level.

The **turbo** signal is asserted during the target performance search to improve the settling time associated with downward target voltage changes. Reaching the target performance under these conditions can be slow if the load circuits are drawing little current, because the regulator alone must discharge the capacitance on the regulated node. As discussed in more detail in Section III, the turbo mode enables an additional current source to help discharge this capacitance. A 50-ns target settle time is chosen for this turbo mode based on a 10-MHz controller clock.

Crossing regulator boundaries while searching for the target voltage can lead to severe loss of voltage regulation. To mitigate this problem, the regulators have a **warmup** signal to prime them for operation before they carry the full load. For example, during a downward search, immediately before the controller crosses the boundary between the high-voltage and low-voltage regulators, the **warmup** signal of the low-voltage regulator is asserted. During the transition, both voltage regulators are enabled and are jointly supplying the load. Once the regulator boundary is crossed, the high-voltage regulator is disabled; the low-voltage regulator handles this event as if were a load current transient. This approach assures smooth transitions but does consume some additional power since both regulators are on simultaneously for the period of a controller clock cycle. The circuit-level use of the **warmup** control line is described in Section III.

This regulation system takes full advantage of the fairly deterministic current waveforms presented by the asynchronous load circuits. dI/dt transients for synchronous designs are generally associated with the charging or discharging of capacitance with the active edge of the clock. In contrast, circuit activity in the deeply pipelined asynchronous datapath is "spread out" in time. As a result, the largest dI/dt transients are associated with the complete filling or emptying of the pipeline at the beginning and

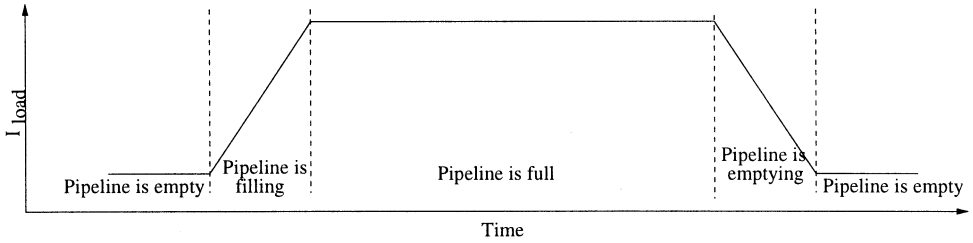

Fig. 2. Typical asynchronous current waveform characteristics.

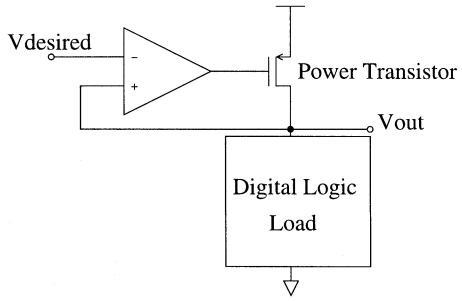

Fig. 3. Linear regulator high-level schematic.

end, respectively, of SIMD instruction execution. This results in the trapezoidal load current profile shown in Fig. 2. The slopes associated with filling and emptying are inversely proportional to the depth of the pipeline and are, consequently, less steep for deep pipelines. The robust nature of the asynchronous digital circuits easily allows them to accommodate $\pm 10\%$ variations around the target voltage, which we establish as the specification for our regulation.

III. Hybrid Regulator System

Linear regulators are the most easily integrable dc–dc converters. Because of this, they have found broad application in low-power digital design [10]–[13] despite "ideal" efficiencies limited to $\eta = V_{\text{out}}/V_{DD}$, where V_{out} is the regulated output voltage and V_{DD} is the supply voltage. Most linear regulators use a pFET output device as shown in Fig. 3 controlled by an amplifier with negative feedback that monitors the output voltage and compares it to the target voltage. The pFET device is usually saturated, except when regulating V_{out} close to V_{DD}. The output impedance increases with increasing frequency because of the rolloff of the control loop gain, necessitating adequate decoupling capacitance on the regulated node. Alternately, a source-follower output stage could be employed (nFET output device) but this necessitates that V_{out} be at least one body-affected gate–source voltage below V_{DD}, resulting in a high dropout voltage, unless a bootstrapped gate bias is employed [10]. Source-follower topologies also require very large power transistors to limit output ripple. In general, the bias currents of the linear regulator must be increased if a fast response time is required resulting in low efficiencies during periods of light loading. The design of linear regulators is also complicated by the wide range of loading characteristics that a digital circuit produces during operation.

The hybrid voltage regulator, proposed here and shown in Fig. 4, receives commands from the power management controller and consists of two separate regulation units: a high-voltage regulator and a low-voltage regulator. In addition, a high-voltage switch allows direct connection of the output to the maximum core supply of 2.5 V through a large pFET device. The digital representation of the target supply voltage from the power management system is converted to an analog voltage with a low-power amplifierless charge redistribution digital-to-analog converter (DAC).

A. High-Voltage Regulator

The high-voltage regulator shown in Fig. 5 regulates target voltages between 2.15 V and 1.16 V, using a pFET common source stage (transistor M5 with current source M6).[1] Transistors M1, M2, M3, M4, and M7 form an nFET differential pair with current-mirror load. With 25 pF of explicit thin-oxide decoupling capacitance added to the output node, the regulator is compensated with Miller capacitor C_c of 2 pF. The explicit decoupling capacitance augments the large intrinsic decoupling added by the drain capacitance of the power transistors of the switch and the two regulators: high-voltage switch (5-mm gate width), high-voltage regulator (3-mm gate width), and low-voltage regulator (2.75-mm gate width). Transistors M12 and M13 preset the gate voltage of the power transistor for smooth regulator transitions when the **warmup** signal is asserted. The linear regulator's power consumption can be reduced to leakage power when it is not enabled by means of transistors M9, M10, and M11. The high-voltage linear regulator's gain-bandwidth product under typical loading conditions varies from 40 MHz at 2.15 V with a 100-mA load to 125 MHz at 1.16 V with a 100-μA load. At 2.15 V, transistor M5 becomes trioded under heavy loading, reducing the power-supply rejection ratio (PSRR).

The bias current value through current source M6 is determined by competing requirements of power dissipation and settle time during target voltage transitions. The use of the **turbo** control signal eases this tradeoff. Under light loading conditions, a minimum current through M6 is necessary for proper settling at each voltage step during target voltage changes. Unfortunately, a fast transition time requires more power to be dissipated in the current source M6 even during periods of constant regulated voltage. This tradeoff between power consumption and settle time is achieved by adding a second current source (M16) in parallel with current source M6. This auxiliary current source is enabled by the power management logic during voltage transitions with the **turbo**

[1]By contrast, a source-follower output stage would have only allowed regulation up to 1.7 V.

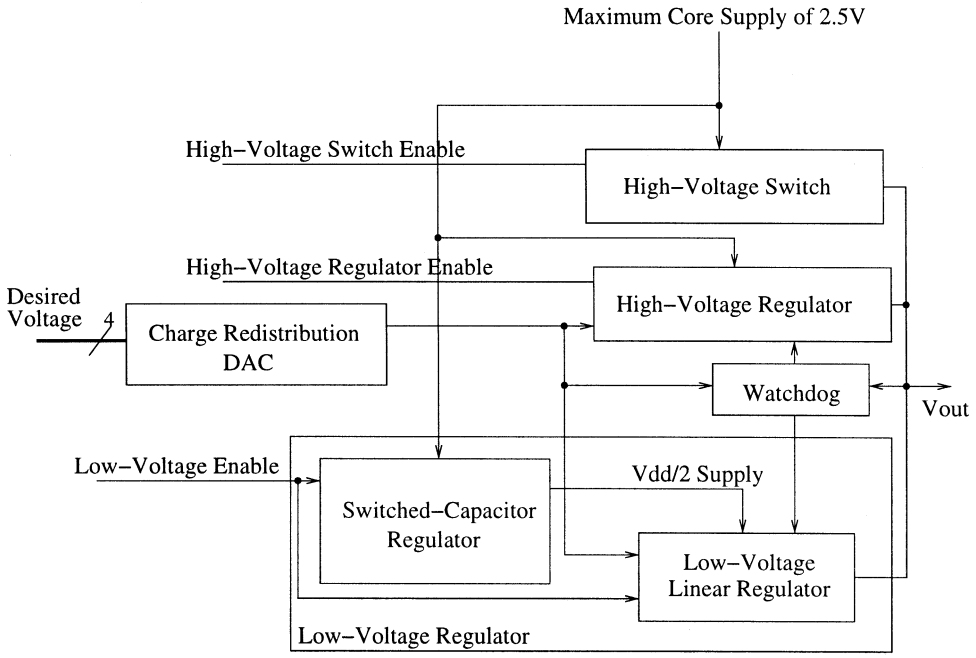

Fig. 4. High-level schematic of hybrid voltage regulator.

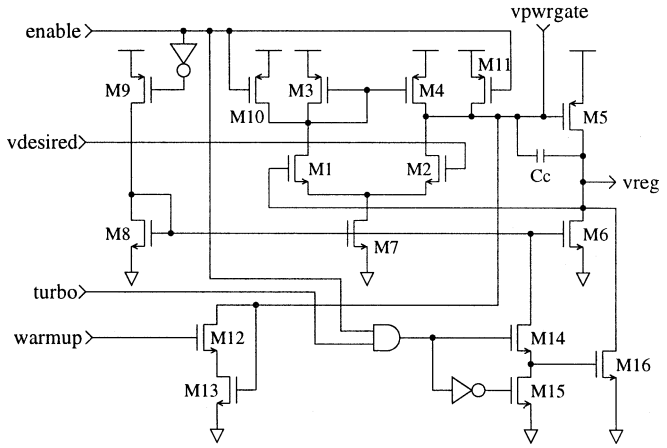

Fig. 5. High-voltage linear regulator schematic.

signal and disabled during all other times. Transistors M14 and M15 turn on current source M16 when **turbo** is asserted.

Another consideration in setting the output stage bias current is to prevent slewing from degrading the response time of the regulator in the presence of large load current transients. High-frequency, large-amplitude transients can cause slewing in the first stage of the amplifier. The effects of slewing are reduced by choosing an appropriate bias current for this stage (500 μA) in conjunction with the discrete time "watchdog" circuitry described below, which effectively augments available bias current during large-signal transients.

B. Watchdog

A discrete-time "watchdog" provides an energy-efficient means to improve the large-signal response times of the regulators. The watchdog output node injects a pulse of charge on to and off of the pFET power transistor gates in the linear regulators. The watchdog consists of two clocked comparators, two dynamic AND gates, a pulse generator, and one set of pull-up/pull-down transistors, as shown in Fig. 6. The watchdog controls the gate voltage of the power transistor in the linear regulator only when the regulator output voltage lies outside a "dead zone" around the target voltage, indicating a transient lack of regulation. This dead zone is implemented by intentionally increasing the offset of two comparators in opposite directions and combining the outputs of the comparators using dynamic AND gates. The comparator offset magnitude and direction is controlled by increasing the capacitance on the integrating nodes of gate isolated sense amplifiers [14]. To save power at low regulated supply voltages, the monitoring frequency of the watchdog is controlled by a digitally tunable current-starved voltage-controlled oscillator (VCO) (rather than by a fixed frequency clock) with four different delay settings between 1.42 and 4.56 ns. When the watchdog is not being used, the VCO can be completely disabled to save power.

The watchdog allows both the monitoring frequency and magnitude of charge injected to be independently controlled by 2-bit digital words allowing the designer to trade off power dissipation and large-signal response time. There is further digital tunability of the dead zone by varying extra capacitance present on the integrating nodes of the comparators. A different set of these parameters is associated with each of the eleven different target voltage values. Under relatively constant loading current conditions, the watchdog is disabled, leaving control of the power transistors to the continuous-time amplifier. As with the **turbo** and **warmup** signal, circuit "knobs" exposed to the system level allow improved system efficiency.

C. Low-Voltage Regulator

Regulation in the voltage range of 0.99–0.65 V is achieved by linearly regulating from a $V_{DD}/2$ supply generated by a switched-capacitor regulator (SCR).

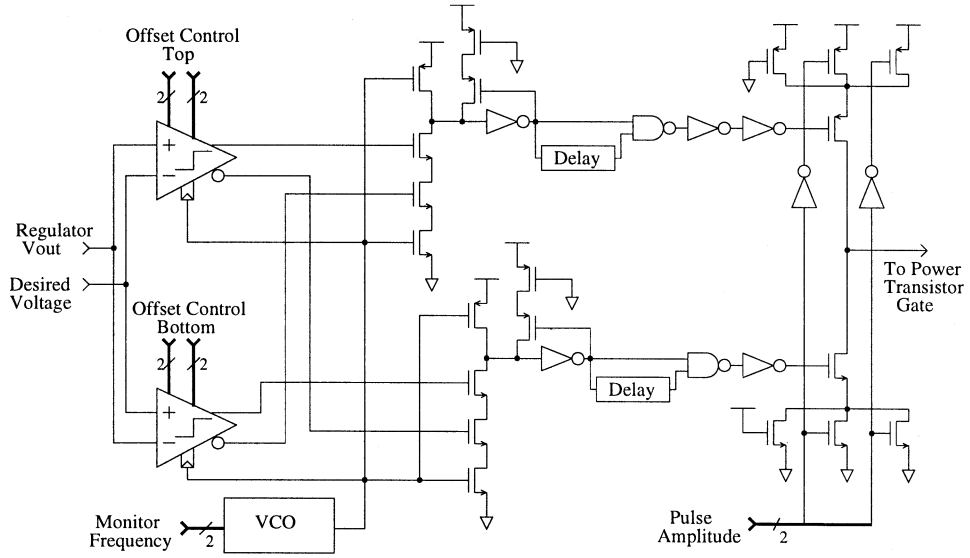

Fig. 6. Watchdog schematic.

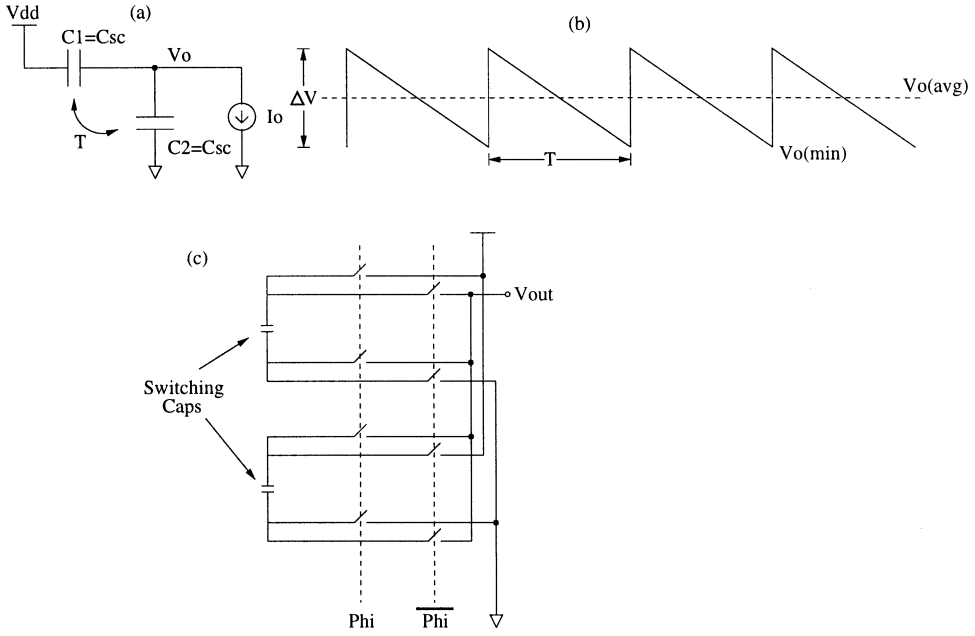

Fig. 7. (a) Simple SCVD. (b) Output waveform. (c) Actual implementation.

Switched-capacitor divider: Switched-capacitor voltage dividers (SCVDs) have been applied to low-power medical implants [15] and can be viewed conceptually as shown in Fig. 7(a). Capacitors $C_{\rm sc}$ are "switched" periodically, resulting in the output voltage waveform shown in Fig. 7(b). The SCVD can be implemented with clock-controlled switches as shown in Fig. 7(c). $V_{o(\rm avg)} = V_{DD}/2$ and for a load current of I_o and switching period of T, $\Delta V = I_o T/2C_{\rm sc}$. For this ideal switched-capacitor voltage divider in which none of the components has loss, conservation of energy shows that the efficiency is given by $\eta_{\rm ideal} = (1+\Delta V/V_{DD})^{-1}$. Note that the efficiency of this *ideal* SCVD is inversely related to the output voltage ripple. Efficiencies are high when the switching frequency is high, when the load is light, and when the switching capacitors are large.

Real SCVDs have finite switch resistances, parasitic capacitances to ground, and output decoupling that modify the expressions for $V_{o(\rm avg)}$ and ΔV. Fig. 8 shows the ideal SCVD of Fig. 7(a) with the addition of parasitic elements. $C_{\rm sc}$ is the switching capacitor with a parasitic to ground $C_{\rm scp}$, C_L is the decoupling capacitance on the load, $R_{\rm sw}$ is the average switch resistance, and $C_{\rm dsw}$ is the source/drain capacitance of the switch. The average output voltage of a real SCVD (neglecting the switch resistance) can be shown to be $V_{o(\rm avg)} = V_{o(\rm min)} + \Delta V/2$ where

$$V_{o(\rm min)} = \frac{V_{DD}(2C_{\rm dsw} + 2C_{\rm sc}) - I_o T}{4C_{\rm sc} + C_{\rm scp} + 4C_{\rm dsw}} \tag{1}$$

and

$$\Delta V = \frac{I_o T}{2C_{\rm sc} + C_{\rm scp} + 8C_{\rm dsw} + C_L}. \tag{2}$$

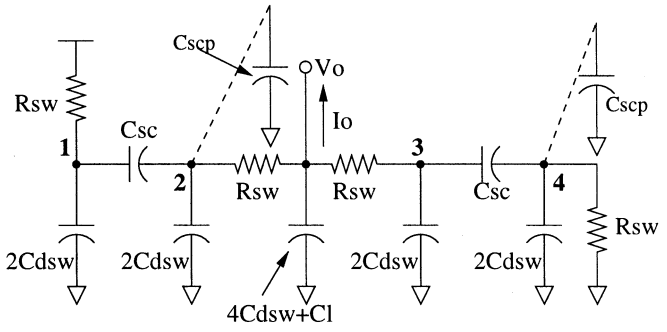

Fig. 8. SCVD equivalent circuit.

$C_{\rm scp}$ will be small if the switching capacitors are implemented with metal–insulator–metal (MIM) capacitors. If the capacitors are instead implemented with MOS structure, these parasitic capacitances to ground can be significant.

These parasitics along with the energy loss due to the charging and discharging of the gate capacitances of the switches degrades the efficiency of a real SCVD from $\eta_{\rm ideal}$. The energy dissipated during a single switching event due to the resistive drops in the switches is given by

$$E_{R_{\rm sw}} = R_{\rm sw}\left(\frac{(\Delta V)^2}{T}\right)\left[(C_{\rm sc}+2C_{\rm dsw})^2+2C_{\rm sc}^2 +(C_{\rm sc}+2C_{\rm dsw}+C_{\rm scp})^2\right].$$

The energy dissipated per switching event due to charging and discharging of the gate capacitances of the eight switches is given by

$$E_{C_{\rm sw}} = 8\times\left(\frac{1}{2}\right)C_{\rm sw}V_{DD}^2 = 4\times C_{\rm sw}V_{DD}^2 \tag{3}$$

where $C_{\rm sw}$ is the gate capacitance of the switch.

Energy is also dissipated in the switch resistances when the parasitic capacitances to ground are transiently charged or discharged with switching:

$$\begin{aligned}E_{C_p} =& \frac{1}{2}(2C_{\rm dsw})\Big[(V_{DD}-V_{o(\min)})^2 \\ &+(V_{DD}-(V_{o(\min)}+\Delta V))^2\Big] \\ &+\frac{1}{2}(2C_{\rm dsw}+C_{\rm scp}) \\ &\times\Big[(V_{o(\min)})^2+(V_{o(\min)}+\Delta V)^2\Big].\end{aligned} \tag{4}$$

Summing all the losses, the efficiency of the SCVD is:

$$\eta_{\rm sc} = \left(1+\frac{(E_{R_{\rm sw}}+E_{C_{\rm sw}}+E_{C_p})}{I_oTV_{o(\rm avg)}}\right)^{-1}. \tag{5}$$

Low-voltage linear regulator: The low-voltage linear regulator has virtually the same topology as in Fig. 5. To accommodate lower common mode voltages, transistors M1, M2, and M7 become pFET devices while transistors M3, M4, and M10 become nFET devices. The current source bias and bias enable logic are also modified to accommodate this change in the differential pair. The supply voltage of the power transistor (M5) is connected to the output of the SCVD.

The bandwidth of the regulator is designed to be lower than the high-voltage regulator because the load circuits are significantly slower in the low-voltage range. In addition, the supply to the power transistor is the triangular waveform shown in Fig. 7. Therefore, the bandwidth depends not only on the target voltage and load current but also on this changing power transistor supply voltage. The gain-bandwidth product when the supply is at 1.25 V is fairly constant over the regulated voltage range but varies from 50 MHz at an 11-mA load up to 70 MHz for loads of less than 1 mA. The regulator gain-bandwidth product varies considerably when the supply voltage is 1.05 V, from 12 MHz at a target voltage of 0.99 V with an 11-mA load to 67 MHz at a target voltage of 0.65 V with a 1-mA load.

The highest regulated voltage (0.99 V) must be at least 50 mV below $V_{o(\min)}$, which is consequently constrained in the SCVD design to be no lower than 1.05 V. While the power transistor is trioded in this case, leading to a poor PSRR, the ripple of the SCVD is designed to be low enough that the 10% voltage noise requirement at the regulated supply output can be satisfied. Included at the output of the SCVD is 100 pF of explicit thin-oxide decoupling capacitance. This, along with the significant intrinsic decoupling provided by the source capacitance of the power transistor of the low-voltage linear regulator, reduces ΔV to approximately 200 mV for $V_{o(\min)} = 1.05$ V and a load current of 11 mA.

The overall efficiency of the SCVD and linear regulator can be computed by allowing for a varying supply in the ideal linear regulator efficiency expression. The average efficiency of a linear regulator with a changing supply voltage is given by

$$\eta_{lin} = \frac{V_{\rm desired}}{\Delta V}\ln\left(1+\frac{\Delta V}{V_{o(\min)}}\right). \tag{6}$$

Therefore, the theoretical efficiency of the low-voltage regulator is $\eta_{\rm lvreg} = \eta_{\rm sc}\eta_{\rm lin}$.

Complete switched-capacitor regulator system: The complete SCR system consists of a divide-by-two SCVD, digital integrator, VCO, and nonoverlapping clock generator as shown in Fig. 9. The VCO is controlled by a thermometer code produced by the digital integrator. The height of this code is determined by comparing $V_{o(\min)}$ of the SCR to a reference ($V_{\rm ref}$). As such, the code is a direct measure of the current demands of the load, scaling the VCO's oscillation frequency higher in periods of high load current. The two comparators shown in Fig. 9 introduce hysteresis in order to suppress dithering in the thermometer code for $V_{o(\min)}$ values immediately around $V_{\rm ref}$. The integrator is implemented as a simple shift register. The code is increased (left shift) for voltages less than the reference and decreased (right shift) for voltages greater than the reference. Further power savings are achieved by varying the digital integrator's frequency along with the SCVD. The **warmup** signal momentarily sets the SCVD's switch frequency to the highest value to ensure the SCVD can provide sufficient current even under worst case loading conditions. The thermometer code also scales the switches of the SCVD; larger switches are only enabled in periods of high VCO oscillation frequencies. To reduce the capacitance that would be necessary for fully complementary switches, pFET

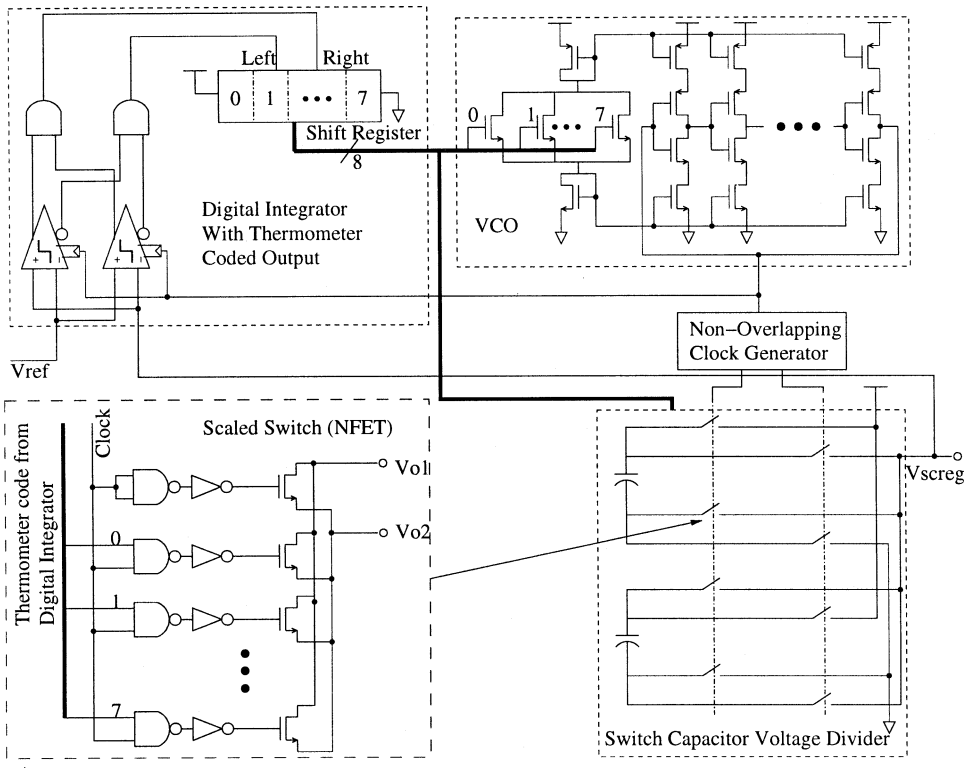

Fig. 9. High-level schematic of hybrid switched-capacitor regulator.

switches are used for connecting the capacitors to the 2.5-V supply, while nFET switches are used for all other connections.

The eight different VCO period (T) and switch-width $(W_{\rm sw})$ settings are chosen to maximize the efficiency for eight different current loading levels between 0.5 and 11 mA. The VCO is run at the minimum frequency to achieve the required $V_{o(\min)}$ target according to (1). The switch width is then chosen to maximize $\eta_{\rm sc}$. To accomplish this, the switch resistances, switch capacitances, and switch parasitics can be expressed in terms of a transistor switch width $W_{\rm sw}$.

Fig. 10 shows $\eta_{\rm sc}$ and $\eta_{\rm lvreg}$ as a function of load current calculated using these simple analytic models for parameter values closely matching the SCVD implemented here: $C_{\rm dsw} = 81$ fF, $C_{\rm sc} = 50$ pF, $C_{\rm scp} = 180$ fF, $C_L = 100$ pF, and $V_{DD} = 2.5$ V. $\eta_{\rm lvreg}$ is presented for three different target voltage values. The discontinuities are associated with the stepped changes in T and $W_{\rm sw}$. The $W_{\rm sw}$ and T optimizations are performed at the supply current levels denoted by the arrows in Fig. 10, constrained by $V_{o(\min)}$.

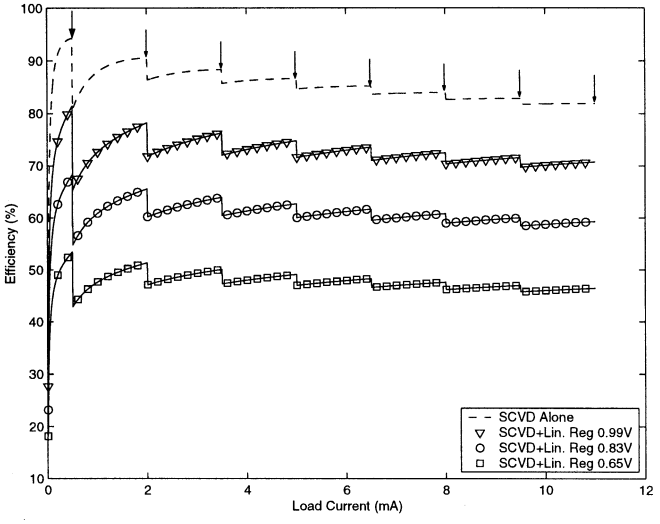

Fig. 10. Efficiency of the low-voltage regulator as a function of load current. The efficiency of the SCVD alone is shown as the dashed curve. The solid curves display the total efficiency of the low-voltage regulator for three different regulated voltage levels.

IV. Results

Fig. 11 shows the complete die photograph of the prototype chip fabricated in a TSMC 0.25-μm technology. The complete system, including linear regulators, switched-capacitor regulator, digital watchdog, and digital controller, occupies 0.4 mm^2. The chip is measured in a 108-pin PGA package. Table I summarizes the important characteristics and performance of the power management system.

Package parasitics preclude external measurement of the regulator output because it is sensitive to additional capacitive loading; therefore, the design incorporates an on-chip 40-MHz 6-bit flash ADC to allow for noninvasive measurement of the regulated supply voltage. Fig. 12 shows the regulated voltage as measured by this ADC for a series of four SIMD instructions with different performance targets. The load is configured to operate at maximum throughput for the target voltage and is briefly reset (stopped) by the instruction unit between instructions. The upper inset of Fig. 12 zooms into the transition from 2.5 to 0.9 V (0.99-V target). It shows the seamless transitions across regulator boundaries—the step from 2.5 to 2.15 V for the high-voltage switch to high-voltage regulator transition, and the

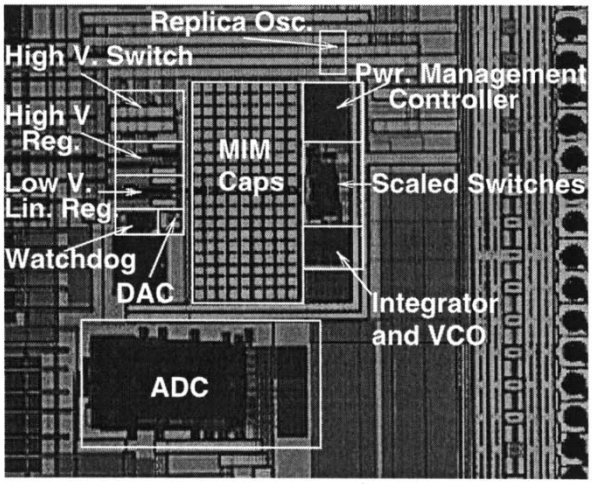

Fig. 11. Die photograph.

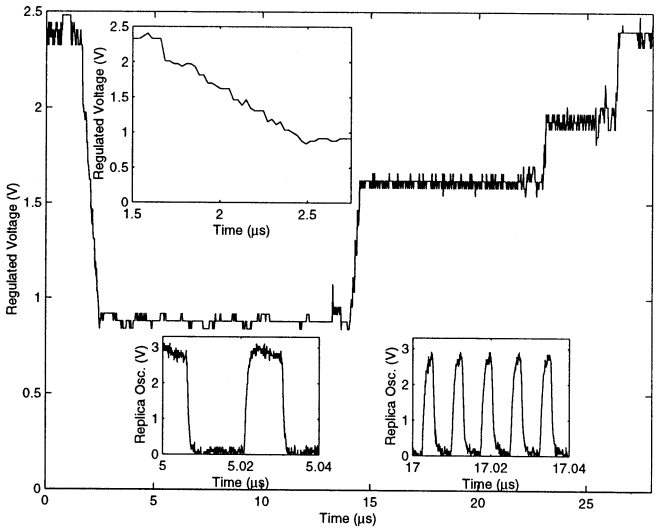

Fig. 12. Measured regulator voltage via ADC and buffered replica oscillator output.

TABLE I
SUMMARY OF CHARACTERISTICS AND MEASURED PERFORMANCE

Measured Efficiency (Heavy Loading)	High-voltage regulator @ 2.35V	93%
	High-voltage regulator@ 1.62V	61%
	Low-voltage regulator @ 0.872V	41%
Area	Total Area	$0.42mm^2$
	Pwr. Management Ctrl.	$0.026mm^2$
	SCR Switching Network	$0.012mm^2$
	Switching Capacitors (MIM)	$0.16mm^2$
	Lin. Regs. and HV pFET	$0.047mm^2$
	Watchdog	$0.0041mm^2$
Pwr. Tran. Widths	High-voltage switch	$5mm$
	High-voltage regulator	$3mm$
	Low-voltage regulator	$2.75mm$
HV Lin. Reg. Bias	Diff. Pair	$530\mu A$
	Pwr. Tran. w/o turbo	$825\mu A$
	Pwr. Tran. w/ turbo	$1650\mu A$
LV Lin. Reg. Bias	Diff. Pair	$280\mu A$
	Pwr. Tran. w/o turbo	$420\mu A$
	Pwr. Tran. w/ turbo	$840\mu A$
Max. load current	High-voltage regulator	100 mA
	Low-voltage regulator	11 mA

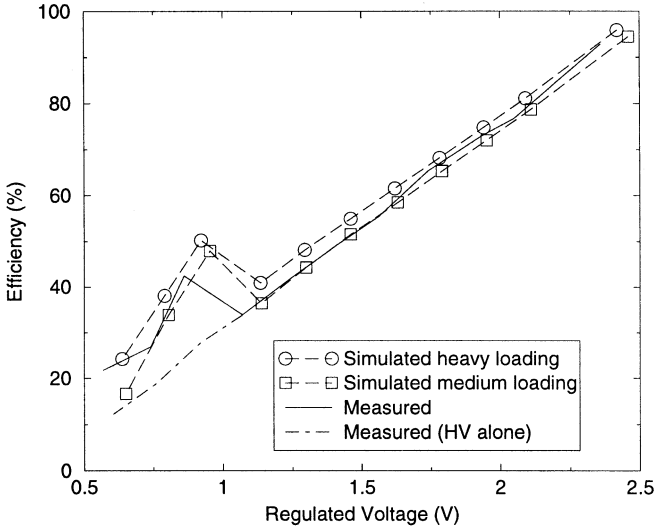

Fig. 13. Simulated and measured efficiency versus voltage.

high-voltage regulator to low-voltage regulator transition for 1.15 to 0.99 V. The 40-MHz sampling frequency of the ADC allows only two samples to be taken per voltage step. The two lower insets of Fig. 12 show the replica oscillator output at two different supply voltages, conforming to a software-specified performance target of two and five replica periods for a 50-ns interval.

The measured and simulated efficiencies of the power management system are shown in Fig. 13 as a function of the target output voltage. The simulated medium loading and heavy loading curves are generated using a diode-connected nFET transistor of 65 and 150 μm, respectively, chosen to approximately model the currents demands of the load. The measured efficiency is calculated as the ratio of the power delivered with the regulators disabled (and the voltage delivered off chip) to the power delivered with the regulators enabled. Override signals are incorporated into the power management system to allow this external control, overriding the power management logic and enabling the high-voltage switch to the external power supply.

Two measured curves are shown in Fig. 13. The curve labeled "measured" shows the efficiency of the complete hybrid regulator system, while the curved labeled "measured (HV alone)" shows the efficiency of using the high-voltage regulator alone for regulation all the way down to 650 mV. The curves diverge below 0.99 V because of the increased efficiency offered by the low-voltage regulator system. At a target output voltage of 872 mV, the hybrid system delivers an efficiency of 41%, while the high-voltage linear regulator alone offers

only 23% efficiency. These efficiencies are lower than those predicted from the simple analytic models of Section III (and indicated in Fig. 10) because of additional sources of loss not considered in these models, specifically the quiescent current of the linear regulators and dynamic power in the integrator, clock generator, and VCO of the SCR.

V. Conclusion

We have described the development of a software-controllable fully integrated on-chip dc–dc downconversion system that combines switched-capacitor voltage dividers and linear regulators to efficiently regulate from 2.5 V down to approximately 650 mV. The use of switched-capacitor supplies allows for better efficiency than what is achievable with linear regulators alone. Multiple linear regulators are used, each optimized to meet the loading requirements for different supply voltage ranges. Techniques for reducing the overhead of linear regulators, such as a mixed analog–digital control loops and auxiliary current sources, are also presented.

References

[1] G. Qu *et al.*, "Energy minimization of system pipelines using multiple voltages," in *Proc. IEEE Int. Symp. Circuits and Systems*, vol. 1, 1999, pp. 362–365.

[2] G. Semeraro, G. Magklis, R. Balasubramonian, D. H. Albonesi, S. Dwaradasas, and M. L. Scott, "Energy-efficient processor design using multiple clock domains with dynamic voltage and frequency scaling," in *Proc. Int. Symp. High Performance Computer Architecture*, 2002, pp. 29–42.

[3] T. D. Burd and R. W. Broderson, "Design issues for dynamic voltage scaling," in *Proc. Int. Symp. Low Power Electronics and Design*, 2000, pp. 9–14.

[4] K. Usami and M. Horowitz, "Clustered voltage scaling technique for low-power design," in *Proc. Workshop Low Power Design*, Aug. 1995, pp. 3–8.

[5] D. E. Lackey, P. S. Zuchowski, T. R. Bednar, D. W. Stout, S. W. Gould, and J. M. Cohn, "Managing power and performance for system-on-chip designs using voltage islands," in *Proc. IEEE/ACM Int. Conf. Computer-Aided Design*, 2002, pp. 195–202.

[6] A. P. Dancy, R. Amirtharajah, and A. P. Chandrakasan, "High-efficiency multiple-output dc–dc conversion for low-voltage systems," *IEEE Trans. VLSI*, vol. 8, pp. 252–263, June 2000.

[7] A. J. Stratakos, S. R. Sanders, and R. W. Broderson, "A low-voltage CMOS dc–dc converter for a portable battery-operated system," in *Proc. Power Electronics Specialists Conf.*, vol. 1, June 1994, pp. 619–626.

[8] G.-Y. Wei and M. Horowitz, "A fully digital, energy-efficient adaptive power-supply regulator," *IEEE J. Solid-State Circuits*, vol. 35, pp. 520–528, Apr. 2000.

[9] Y. W. Li, G. Patounakis, and K. L. Shepard, "High-throughput asynchronous datapath with software-controlled voltage scaling," in *Symp. VLSI Circuits Dig. Tech. Papers*, June 2003, pp. 49–52.

[10] G. W. den Besten and B. Nauta, "Embedded 5 V-to-3.3 V voltage regulator for supplying digital IC's in 3.3 V CMOS technology," *IEEE J. Solid-State Circuits*, vol. 33, pp. 956–962, July 1998.

[11] S. Jou and T. Chen, "On-chip voltage down converter for low-power digital system," *IEEE Trans. Circuits Syst. II*, vol. 45, pp. 617–625, May 1998.

[12] L. R. Carley and A. Aggarwal, "A completely on-chip voltage regulation technique for low power digital circuits," in *Proc. Int. Symp. Low-Power Electronics and Design (ISLPED)*, 1999, pp. 109–111.

[13] N. Dragone, A. Aggarwal, and L. R. Carley, "An adaptive on-chip voltage regulation technique for low-power applications," in *Proc. Int. Symp. Low-Power Electronics and Design (ISLPED)*, 2000, pp. 20–24.

[14] M.-J. E. Lee, W. J. Dally, and P. Chiang, "Low-power area-efficient high-speed I/O circuit techniques," *IEEE J. Solid-State Circuits*, vol. 35, pp. 1591–1599, Nov. 2000.

[15] J. G. Ryan, K. J. Carroll, and B. D. Pless, "A four chip implantable defibrillator/pacemaker chipset," in *Proc. Custom Integrated Circuits Conf.*, 1989, pp. 7.6/1–7.6/4.

George Patounakis (S'00) received the B.S. degree in electrical engineering from Rutgers University, New Brunswick, NJ, in 2000 and the M.S. degree in electrical engineering from Columbia University, New York, NY, in 2001. He is currently working toward the Ph.D. degree at Columbia University in the Columbia Integrated Systems Laboratory.

His research interests include interfacing biological molecules with silicon microelectronics, on-chip power management, and high-speed intrachip interconnect.

Mr. Patounakis is a recipient of the Columbia SEAS Presidential Fellowship and the Intel Ph.D. Fellowship.

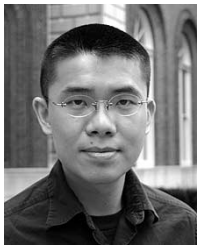

Yee William Li (S'00) received the B.Eng. degree (First Class Honors) in computer engineering from The University of Hong Kong, Pokfulam, Hong Kong, in 2000 and the M.S. degree in electrical engineering from Columbia University, New York, NY, in 2001. He is currently working toward the Ph.D. degree at Columbia University.

His research focuses on low-power circuit design techniques for signal processing applications.

Mr. Li was the receipient of a Croucher Foundation Scholarship. He received the Engineering Excellence Award while with Motorola Semiconductors. His chip was a winning entry in the 2003 ISLPED design contest.

Kenneth L. Shepard (S'85–M'92–SM'03) received the B.S.E. degree from Princeton University, Princeton, NJ, in 1987 and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1988 and 1992, respectively.

From 1992 to 1997, he was a Research Staff Member and Manager in the VLSI Design Department, IBM T. J. Watson Research Center, Yorktown Heights, NY, where he was responsible for the design methodology for IBM's G4 S/390 microprocessors. Since 1997, he has been with Columbia University, New York, NY, where he is now an Associate Professor. He also served as Chief Technology Officer of CadMOS Design Technology, San Jose, CA, until its acquisition by Cadence Design Systems in 2001. His current research interests include design tools for advanced CMOS technology, on-chip test and measurement circuitry, low-power design techniques for digital signal processing, low-power intrachip communications, and CMOS imaging applied to biological applications.

Dr. Shepard received the Fannie and John Hertz Foundation Doctoral Thesis Prize in 1992. At IBM, he received Research Division Awards in 1995 and 1997. He was also the recipient of an NSF CAREER Award in 1998 and IBM University Parnership Awards in 1998, 1999, 2000, and 2001. He also received the 1999 Distinguished Faculty Teaching Award from the Columbia Engineering School Alumni Association. He has been an Associate Editor of the IEEE Transactions on Very Large Scale Integration Systems and was the technical program chair and general chair for the International Conference on Computer Design in 2002 and 2003, respectively. He has served on the program committees for ICCAD, DAC, ISCAS, ISQED, GLS-VLSI, TAU, and ICCD.