Ultra-Dynamic Voltage Scaling (UDVS) Using Sub-Threshold Operation and Local Voltage Dithering

Benton H. Calhoun, *Student Member, IEEE,* and Anantha P. Chandrakasan, *Fellow, IEEE*

***Abstract*—Local voltage dithering provides near optimum savings when workload varies for fine-grained blocks. Combining this approach with sub-threshold operation permits ultra-dynamic voltage scaling from 1.1 V to below 300 mV for a 90-nm test chip. Operating at 330 mV provides minimum energy per cycle at 9X less energy than ideal shutdown for reduced performance scenarios. Measurements from the test chip characterize the impact of temperature on the minimum energy point.**

***Index Terms*—Dynamic voltage scaling, low-power circuits, sub-threshold operation, voltage dithering.**

I. Introduction

PROBLEMS arising from technology scaling have recently made power reduction an important design issue for circuits and applications that traditionally are driven almost solely by performance constraints. The increased importance of power is even more notable for a new class of energy-constrained systems. This type of application must achieve long system lifetimes from a limited energy source, so the need to reduce energy consumption whenever possible is paramount. Dynamic voltage scaling (DVS) has become a standard approach for reducing power when performance requirements vary. DVS systems lower the frequency and voltage together to reduce power when lower performance is allowed. Voltage dithering was proposed as a low overhead implementation of DVS to provide near-optimum power savings using only a few discrete voltage and frequency pairs [1]. The savings are only achievable if the voltage and frequency can change on the same time scale as the altering workload. Previous implementations apply voltage dithering to entire chips and require many microseconds to change operating voltage [1], [2]. This paper describes a 90-nm test chip that demonstrates the proposed concept of *local voltage dithering (LVD)* and couples LVD with sub-threshold operation to achieve *ultra-dynamic voltage scaling (UDVS)* [3].

Our work extends the state-of-the-art in several ways. First, LVD improves on existing voltage dithering systems by taking advantage of faster changes in workload and by allowing each block to optimize based on its own workload. Additionally, we show that the time and energy overhead of LVD are small based on measurements from the test chip. Second, UDVS provides a practical method for extending DVS into the sub-threshold region. The minimum energy per operation for a block occurs in the sub-threshold region [4], and sub-threshold circuits have been demonstrated in silicon [5]. For many emerging energy-constrained applications, lowering energy consumption is the primary constraint under most conditions. Thus, operating at the minimum energy point conserves energy at the cost of lower performance (frequency). This type of application works at the minimum energy point primarily, and only jumps to higher performance voltages in rare cases. We have shown the effectiveness of UDVS coupled with LVD using a 90-nm, bulk CMOS test chip. We also provide measurements of the effect of temperature on minimum energy operation for the 90-nm test chip.

Manuscript received May 14, 2005; revised September 7, 2005. This work was funded by Texas Instruments Incorporated and by the Defense Advanced Research Projects Agency (DARPA) through a subcontract with MIT Lincoln Laboratory.

The authors are with the Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: bcalhoun@mtl.mit.edu; anantha@mtl.mit.edu).

Digital Object Identifier 10.1109/JSSC.2005.859886

II. Dynamic Voltage Scaling and Voltage Dithering

Many signal processing systems process blocks of data that arrive at some regular rate, and sometimes the amount of data to process is less than the maximum amount. This corresponds to a fixed-throughput system whose workload requirements change in a time varying fashion. Examples of this type of application include MPEG video processing and FIR filtering with a variable number of taps [1]. In the video processing example, the maximum workload corresponds to a scene change in the video sequence. In this case, the entire new frame of data requires processing. In the absence of scene changes, only a fraction of the new frame requires processing. This case represents a reduced workload for the system. In a system without buffering, the lowest allowable rate (normalized operating frequency) equals the workload. If buffering is possible, then there are different strategies involving operation at different rates that can correctly perform the required processing on a block with a given workload by ensuring that the average rate equals the workload. There are many applications where workload varies with time [6], and policies for setting the rate based on incoming data have been explored [7]–[9].

Fig. 1 shows four approaches to power supply management for reducing energy consumption when the workload varies [1]. It plots the required rate of the system versus the normalized energy required to process one generic block of data. The most straightforward method for saving energy when the workload decreases is to operate at the maximum rate until all of the required processing is complete and then to shutdown. This approach only requires a single power supply voltage (corresponding to full rate operation), and it results in linear energy savings. The fixed power supply curve in Fig. 1 assumes ideal shutdown (i.e., no shutdown power). A variable supply voltage with infinite allowable levels provides the optimum curve for reducing energy. This curve in Fig. 1 corresponds to theoretically

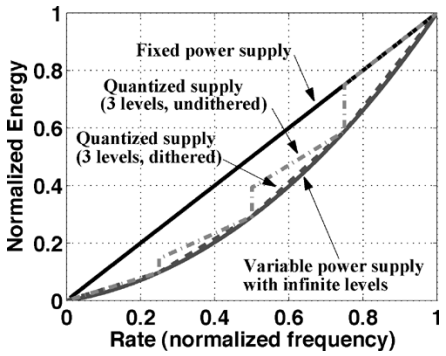

Fig. 1. Theoretical energy consumption versus rate for different power supply strategies [1].

ideal DVS according to the model in [1] where velocity saturation is omitted. When velocity saturation occurs, the energy savings for ideal DVS increase [2].

One method that avoids the problem of creating an infinite number of supply voltages is to use quantized supply voltages. In Fig. 1, three levels of supply voltage quantization are used with two different policies. The undithered policy simply selects the lowest supply voltage for which the rate exceeds the desired rate, operates at that rate and voltage until all of the data in the block is processed, and then shuts down. This results in the stair-step energy characteristic. A better method is called voltage dithering [1]. The basic idea behind voltage dithering is to divide the computation of one block of data between operation at the quantized supply voltage and rate pairs that occur above and below the desired average rate. The energy profile for dithering between quantized voltage supplies linearly connects the quantized rate and energy pairs on the plot.

Implementations of systems that use voltage dithering apply it monolithically to an entire chip. The system in [1] uses an on-chip variable DC–DC converter to dither the voltage supplied to the entire chip. A chip containing header switches was used in [2] to select the voltage supplied to a different chip (off-the-shelf processor). A similar off-chip voltage hopping approach is used in [10] for a zero-V_{T} processor in FD-SOI. These implementations have shown the effectiveness of voltage dithering to save energy for high-performance applications with variable workload.

III. Local Voltage Dithering

We have proposed local voltage dithering (LVD) to improve upon chip-wide voltage dithering. This section discusses the advantages of voltage dithering at the local level and describes a test chip that demonstrates these improvements.

Previous chip-wide implementations using voltage dithering report that the transition between two different supply voltages takes hundreds of microseconds [1], [2]. This prevents the system from achieving any energy savings for faster changes in the workload. Dividing up the power supply grid into local regions reduces the capacitance that must be switched when the voltage supplied to a local block needs to change. This allows for faster changes in supply voltage and permits energy savings for changes in workload on the same timescale.

Chip-wide voltage dithering also restricts the extent to which varying workload may be leveraged because it must account for the highest workload from all of the blocks across the entire chip. The energy savings that are lost by using chip-wide voltage dithering only increase with more blocks and wider differences between the maximum and minimum workloads. In contrast, LVD lets each block operate according to its own workload.

Our implementation of LVD uses embedded power switches (PMOS header devices) to toggle among a small number of voltage levels at the local block level. One advantage of this implementation approach is that the local dithering switches can be turned off to provide fine-grained power gating essentially for free.

A. Test Chip

We have implemented a test chip in 90-nm bulk CMOS to demonstrate LVD and UDVS. Fig. 2 shows the primary block used for testing LVD on the test chip. The circuit of interest is a 32-bit Kogge–Stone adder that can be configured as an accumulator for testing. In this figure, two PMOS header switches select between two supply voltages for the adder block. Other adders on the chip have different numbers of header devices. A critical path replica ring oscillator shares the same dithered voltage supply as the adder and sets the frequency of the clock based on the selected voltage. The die photo in Fig. 14 shows the accumulators with different numbers of header switches used for testing, and the approximate area of a single header switch is highlighted for reference.

Placing a PMOS header switch in series with the power supply increases the delay of the circuit because of the voltage drop across the on resistance of the header. This effect is well-known and thoroughly analyzed in the context of power gating approaches such as multi-threshold CMOS (MTCMOS). Numerous methods for sizing such header devices are available, and most of them are designed to ensure that the circuit never exceeds some delay penalty relative to the circuit without any headers. The header switches on the test chip are sized to keep the delay penalty less than 10%.

Fig. 3 shows the architecture and circuit schematics for the adder block on the test chip. The inverse of the propagate and generate signals are calculated in the first stage, and these results are applied to the adder tree. Each reconverging point in the tree has a "dot operator" circuit that calculates the propagate and generate values for that stage. Each stage in our implementation is inverting, so the two flavors of dot operator are shown in the critical path schematic.

B. Measured LVD Results

Fig. 4 illustrates the savings that LVD provides for the adder block on the test chip when the rate varies. The dotted line shows operation at the highest rate followed by ideal shutdown. The solid line shows the measured energy versus rate for DVS assuming continuous voltage and frequency scaling. Selecting two rates from the curve, 1 and 0.5 in the figure, and operating for the correct fraction of time at each rate results in the dashed line that connects the quantized points, as described previously. A local block with three headers can achieve closer to optimum

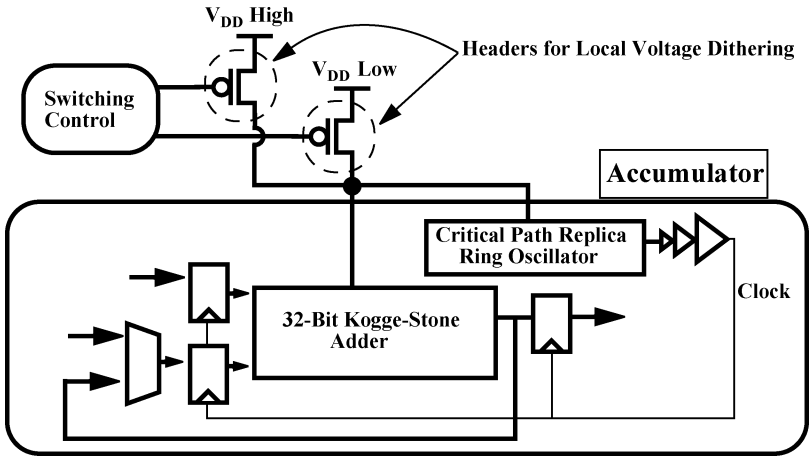

Fig. 2. Block diagram of voltage dithered adder and critical path replica using two local header switches for local voltage dithering (LVD).

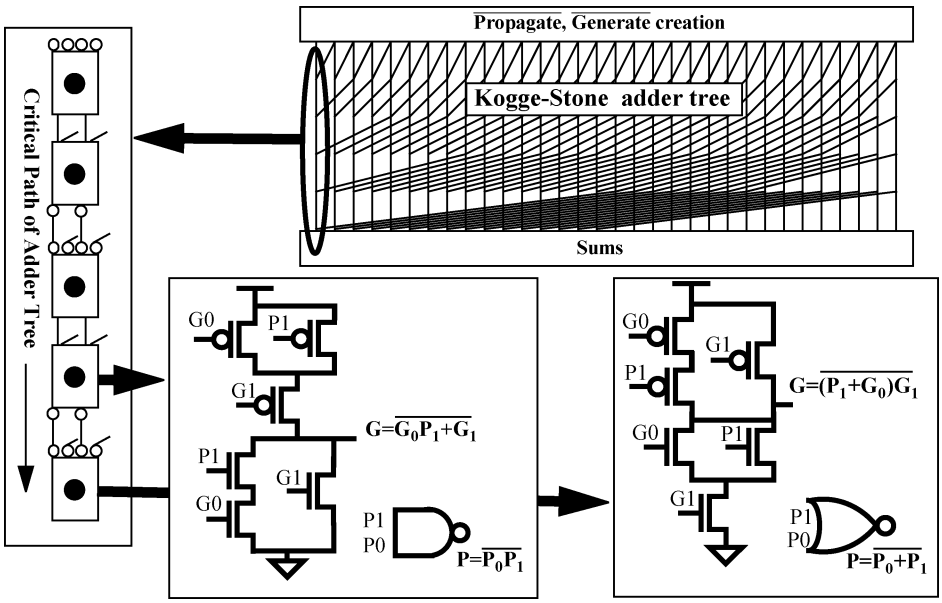

Fig. 3. Schematic of adder circuits. Kogge–Stone-based tree with inverting stages of dot operators (at each reconvergence of the tree). Circuit schematics show that the circuits do not require large stacks of transistors, which degrade sub-threshold operation.

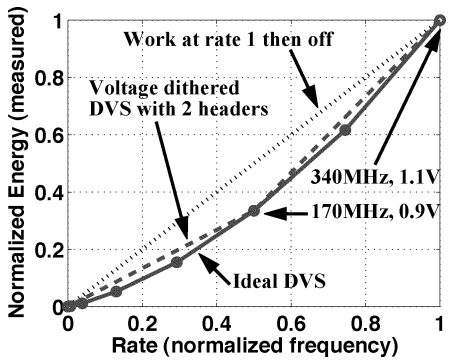

Fig. 4. Characterized local voltage dithering using measured results for 32-bit Kogge–Stone adder.

savings by selecting three rates and then dithering to connect those points on the plot.

While previously reported chip-wide approaches to voltage dithering have largely ignored the overhead energy of their schemes, we have investigated and measured the time and energy overhead of the LVD switching approach. The delay overhead is the amount of time that operations in a voltage dithered logic block must be suspended to allow for the voltage supply to settle to its new value during a V_{DD} change. Fig. 5 shows the test circuit used to measure the delay overhead of LVD for the test chip. This circuit stops the clock to the adder block for a specified number of cycles immediately following a change in V_{DD}. By decreasing the duration of clock gating at each V_{DD} transition until the adder fails, we can determine the number of cycles required for the supply voltage to settle sufficiently such that the adder operates correctly. Measurements showed that the adder continued to function correctly even with only 1/2 cycle (minimum possible using the test circuit) of clock gating for V_{DDL} above 0.6 V, which corresponds to a rate of 0.04. Thus, even conservative settling times for this LVD implementation are on the order of a few cycles. This measurement confirms that LVD can adjust to fast changes in the workload of the local blocks. The oscilloscope plot in Fig. 5

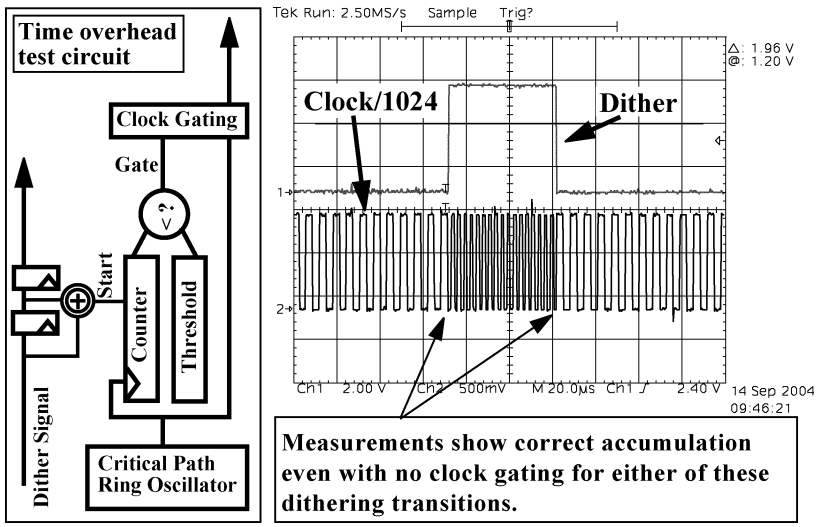

Fig. 5. Measuring timing overhead of LVD between rate 0.5 (170 MHz) and rate 1 (340 MHz).

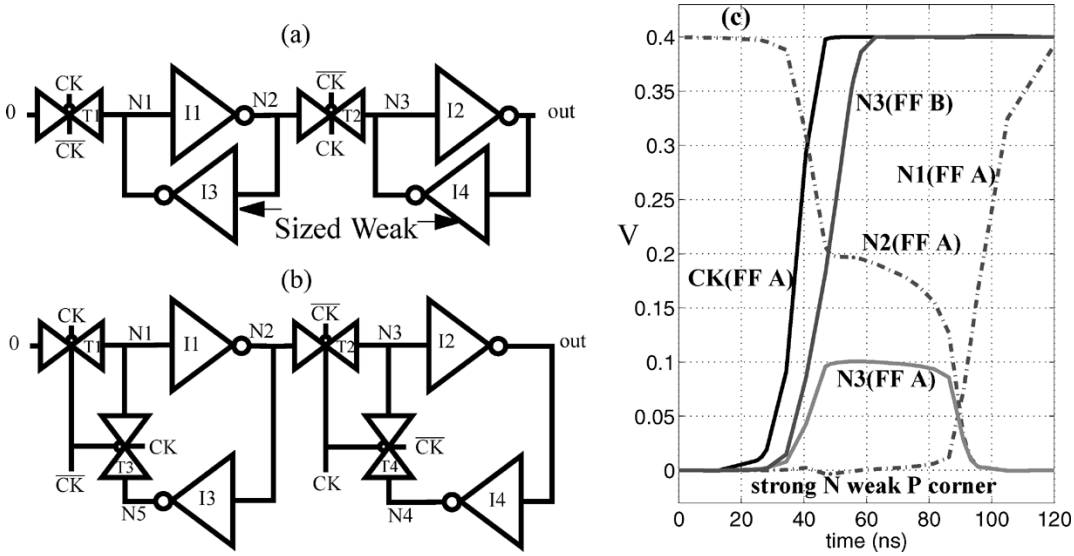

Fig. 6. Flip-flop selection for sub-threshold operation. Ratioed circuits such as in (a) cannot function across process corners in sub-$V_{\rm T}$ at the minimum energy voltage because the exponential dependence of current on $V_{\rm T}$ becomes more important than sizing. Cutting the feedback loop for writing a latch (b) is robust across all corners for operation at the minimum energy voltage. (c) shows simulation of flip-flop A failing.

shows the divided ring oscillator output and the signal that selects the supply voltage (dither) for a dithering cycle between full and 0.5 rate.

In addition to timing overhead, there is energy overhead associated with the LVD approach. The buffer network and control circuits that drive the header switches consume energy every time they toggle the header switches to select a new supply voltage for the adder circuit. We can relate this overhead switching energy to the active switching energy of the adder block to determine its impact on overall energy savings from the LVD approach. To this end, we normalize the effective overhead switched capacitance of the control and buffer circuits and header switches, $C_{\rm OVERHEAD}$, to the effective switched capacitance of the adder. The expression in (1) shows the relation that must hold true in order for LVD to provide energy savings for a given transition.

$$NV_{\rm DDH}^2 \geq NV_{\rm DDL}^2 + C_{\rm OVERHEAD}V_{\rm DDH}^2 \qquad (1)$$

Solving (1) for N gives the number of cycles that must occur at $V_{\rm DDL}$ in order to make switching to $V_{\rm DDL}$ worthwhile for saving energy. Measurements of the test chip show that $C_{\rm OVERHEAD} = 3.7$ for the adder, so N is only 12 for the adder block with $V_{\rm DDH} = 1.1$ V and $V_{\rm DDL} = 0.9$ (rate $= 0.5$). Since the control circuits on the test chip are relatively simple, the overhead energy for more complicated control schemes, such as those that calculate the effective workload, has the effect of increasing N.

IV. Sub-Threshold Operation

This section examines some of the issues related to sub-threshold operation. For energy-constrained applications, lowering the energy per operation is more important than operating at high performance during most of the system's active lifetime. Thus, circuits for such energy-constrained

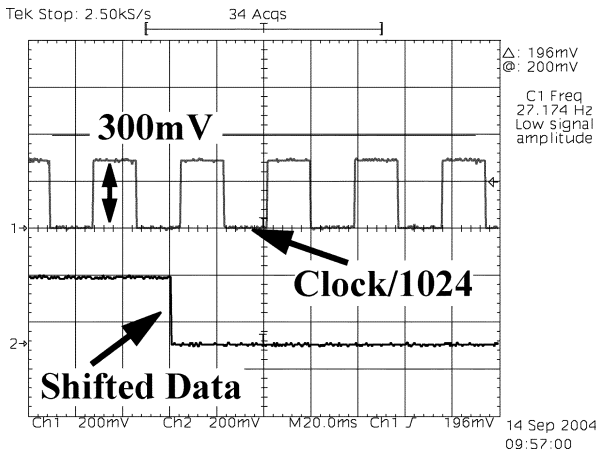

Fig. 7. Oscilloscope plot showing the clock and data from the 90-nm test chip operating at 300 mV, just below the minimum energy point. The adders functioned to 200 mV.

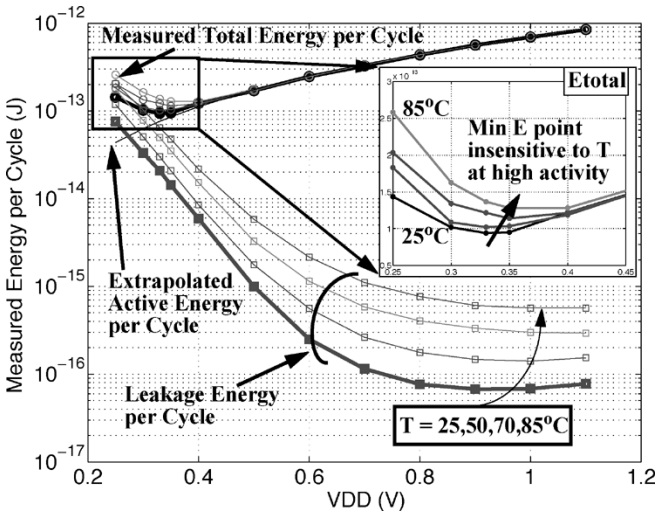

Fig. 8. Measured energy per cycle in the sub-threshold region for input activity of one. Minimum energy point occurs at 330 mV (50 kHz) and 0.1 pJ per operation at 25 °C. The optimum supply voltage is insensitive to temperature variation.

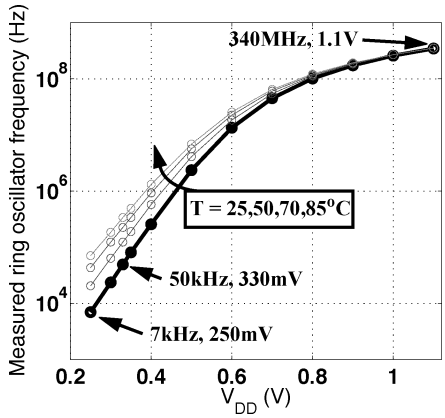

Fig. 9. Measured ring oscillator frequency versus $V_{\rm DD}$ and temperature.

systems need to operate in sub-threshold to minimize energy per operation.

Generally, static CMOS circuits are robust for sub-threshold operation. The large sensitivity of sub-threshold circuits to process variation makes some circuit designs more suitable than others. For example, Fig. 6 demonstrates an important consideration about circuit selection for sub-threshold. Fig. 6 shows a ratioed flip-flop [flip-flop A in (a)] and a flip-flop with transmission gates for cutting off feedback in the latches [flip-flop B in (b)]. The ratioed flip-flop uses devices sized to be strong (wide) on the critical path and inverters sized to be weak for the feedback in the latches. This approach depends on the sizing ratios in the circuit to guarantee functionality, and it works well above-threshold. In sub-threshold, however, current depends exponentially on $V_{\rm T}$. This makes process variation and mismatch quite significant relative to device sizes. Fig. 6(c) shows a plot of flip-flop A failing to write node N3 to a one at the strong NMOS, weak PMOS corner. The pull-up path through I1 and T2 is weak because of the higher PMOS threshold voltage and cannot overcome the pull-down path through I4 and T2 that is made strong by the lower NMOS threshold voltage. Instead, the pull-down path overcomes I1 and actually flips the state of the master stage by pulling N2 to zero. Thus, flip-flop A fails below 450 mV at this corner, preventing the circuit from reaching its minimum energy voltage. Massively upsizing the PMOS devices can correct this problem, but a better choice is to eliminate the ratioed fight by adopting flip-flop B. This flip-flop cuts off the feedback path before writing a latch, allowing the $V_{\rm GS}$ applied to on-transistors to increase current beyond any device off-currents even at process corners. By examining the circuits used in the chip to make sure that sub-$V_{\rm T}$ currents are never too imbalanced, we designed a block that operates to below 200 mV. Fig. 7 shows an oscilloscope plot of the adder on the 90-nm test chip operating in sub-threshold at 300 mV, just below the minimum energy voltage.

The minimum energy per operation point measured for the adder appears in Fig. 8 at $V_{\rm DD} = 330\,{\rm mV}$ ($f = 50\,{\rm kHz}$) and 0.1 pJ per addition for 25 °C. Fig. 8 also shows the measured effect of temperature on the total energy per cycle and leakage energy per cycle. An increase in temperature lowers the mobility of MOSFETs and decreases the threshold voltage according to: $\mu(T) = \mu(T_o)(T/T_o)^{-M}$ and $V_T(T) = V_T(T_o) - KT$. For above-threshold operation, the decreased mobility dominates, and circuits slow down as they heat up. The leakage energy increases quickly with temperature for $V_{\rm DD} > V_{\rm T}$ because of the exponential dependence on the lower threshold voltage. In the sub-threshold region, however, the increased current also decreases the cycle delay, which causes the higher leakage currents to integrate over a shorter cycle time. As a result, the leakage energy does not change enough with temperature to greatly impact the optimum supply voltage. Fig. 8 shows that the measured effect of temperature on the minimum energy point is small, validating the model in [11]. Fig. 9 shows the measured frequency of one of the critical path ring oscillators on the test chip versus $V_{\rm DD}$ and temperature, confirming the increase of performance at higher temperatures in the sub-threshold region. Different process corners also can have a strong impact on leakage energy, changing it by up to a few orders of magnitude in the above-threshold region. In sub-threshold, process corners that increase leakage current also tend to decrease delay because of larger on currents. Thus, although the impact on the minimum energy point is larger than for temperature, it is still in the range of 100 mV even for orders of magnitude changes in leakage current.

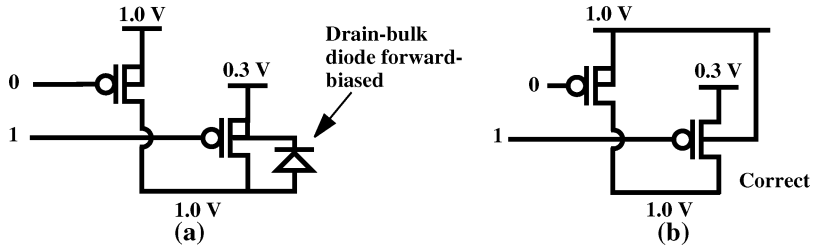

Fig. 10. For UDVS, tying the bulk connections of the PMOS header switches to their own sources can forward bias the drain-bulk diode (a), so the bulk terminals of the headers have to connect to the highest supply voltage as in (b).

V. Ultra-Dynamic Voltage Scaling

The discussion to this point has assumed that the varying rate remains above roughly a few percent of one. As previously mentioned, sub-threshold operation has proven to minimize energy for low performance applications. While scaling to sub-threshold is rarely advantageous for full processors [12], local blocks or special applications that require brief periods of high performance spend significant amounts of time operating at effective rates that are orders of magnitude below one. Examples of these applications include microsensor nodes, medical devices, wake-up circuitry for processors, and local blocks on active processors. When performance is noncritical, energy is minimized by operating at the minimum energy point that occurs because of increased leakage energy at low frequency and then shutting down if there is more timing slack.

Since LVD works well for high-speed operation and operating at the minimum energy point is optimal for low performance situations, we propose ultra-dynamic voltage scaling (UDVS) using local power switches [3]. This approach uses local headers to perform LVD when high performance is necessary and selects a low voltage for sub-threshold operation at the minimum energy point whenever performance is not critical. As with LVD, all of the headers for a given block can turn off when the block is idle to conserve standby power using power gating. Since the power switches in the UDVS approach connect to different voltages that can differ substantially, they must be configured carefully to prevent forward biasing the junction diodes. Fig. 10(a) shows that connecting the bulk terminal of the header transistors to the source can forward bias the drain-bulk diode when the V_{DDL} switch is off. One solution to this problem is to tie the bulk of all of the header switches to the highest supply voltage, as in Fig. 10(b).

Fig. 11 provides one example of measured UDVS characteristics for the adder. In this example, dithered voltages are chosen at 1.1, 0.8, and 0.33 V, which is the optimum voltage for minimum energy. When the adder block is performing operations with no timing deadline, it functions at the minimum energy point at 50 kHz and saves 9× the energy versus the ideal shutdown scenario. When performance becomes important, the adder dithers between 1.1 and 0.8 V within 30% of the optimal energy consumption while adjusting for variations in the rate above 0.1. It was shown in [13] that significant extra savings are available if you can match the selected dithered rates to the prominent average rates in the data. This brings the dithered curve closer to the optimum DVS curve for the common cases. Fig. 12 shows two additional examples in which the supply

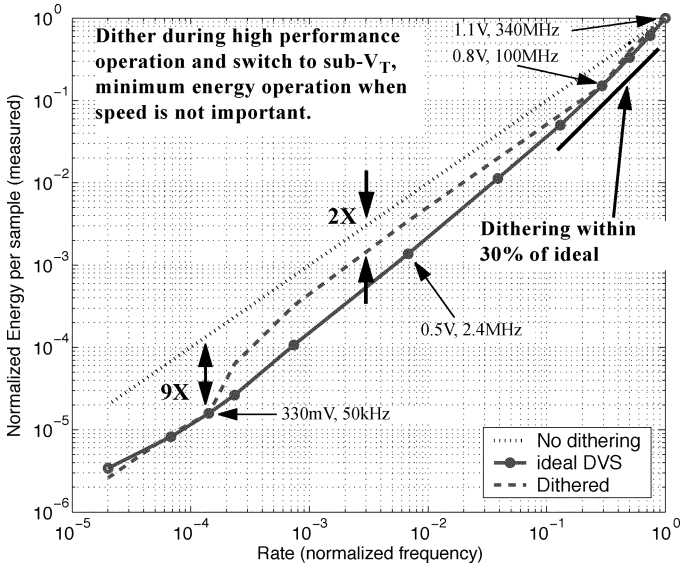

Fig. 11. Ultra-dynamic voltage scaling (UDVS) using two headers with one variable DC–DC converter or using three headers (see Fig. 13).

voltages are chosen for different scenarios. If performance constraints prevent a system from ever operating at the minimum energy point, the supply voltage can be adjusted to higher voltages to achieve near optimal energy operation over the range of higher rates as in Fig. 12(a). For a system whose rate requirements vary evenly over the full range, the voltage choices in Fig. 12(b) provide a better match to the ideal energy profile, but the minimum energy per operation is not achievable.

Fig. 13 shows two options for implementing the power supplies and headers in a UDVS system. The straightforward option is to distribute three supply voltages around the chip and to use three header switches at each block as in Fig. 13(a). Although only one block is shown, multiple blocks can share the same power rails and still operate at different effective rates by dithering based on their own workloads. The voltages V_{DDH}, V_{DDM}, and V_{DDL} can be selected based on the system workload statistics as we described above. The only advantage to using more than three power supplies is to pin the UDVS energy profile to the ideal variable supply profile in more places. Fig. 13(b) offers a second option for implementing UDVS. When transitions between high and low performance mode are infrequent and the amount of time in between transitions is long, two header switches may be paired with one adjustable DC–DC converter for the same functionality. For example, during high-speed operation, the headers dither between 1.1 and 0.8 V. When the rare transition to low speed occurs, the

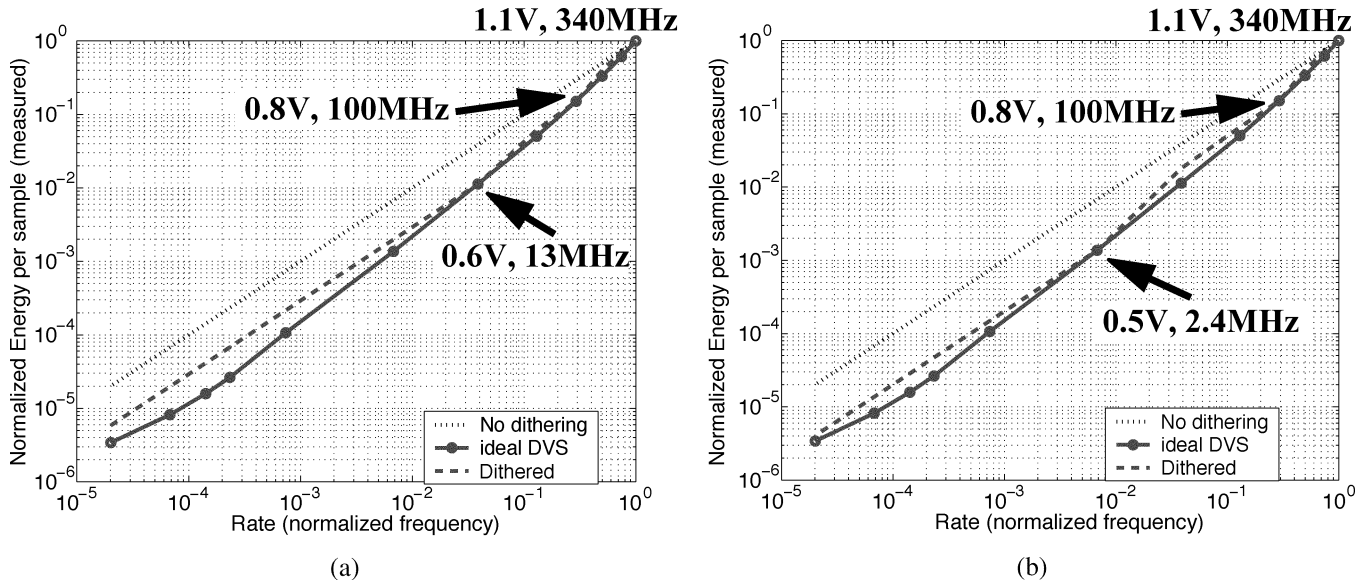

Fig. 12. Different choices of dithered voltages. (a) Closer to ideal over high-performance range. (b) Closer to ideal over full range.

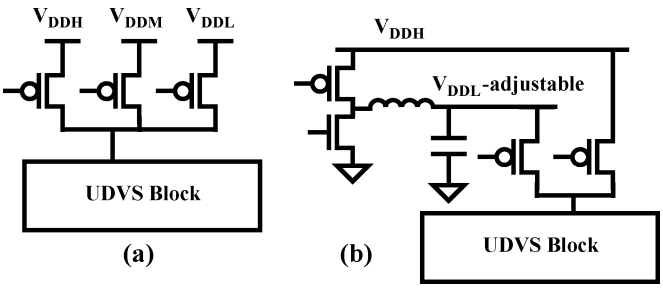

Fig. 13. Options for UDVS headers at the system level using three always-available voltages (a) or one unchanging voltage and one adjustable voltage (b).

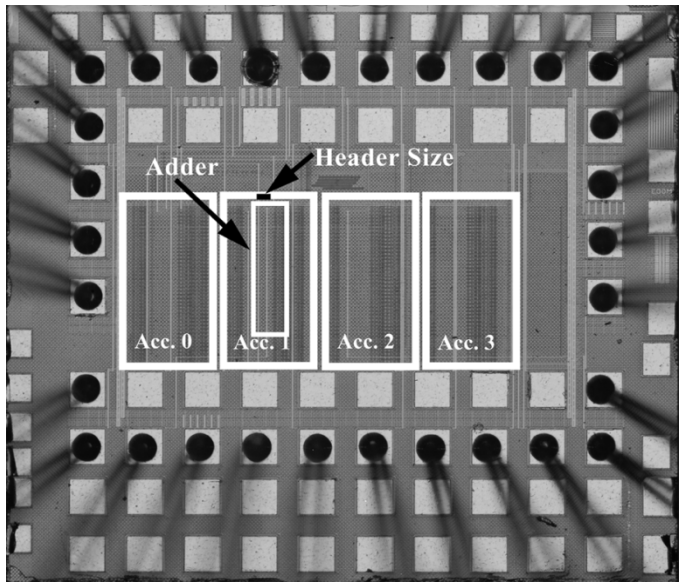

Fig. 14. Annotated die photograph showing accumulators with 0, 1, 2, and 3 headers. The size of one header is highlighted for reference.

DC–DC converter switches V_{DDL} to 0.35 V so that all of the blocks can operate near their minimum energy points.

For UDVS and LVD applications where blocks operating at different voltages have to communicate with each other, level converting circuits are required at the low voltage region to high-voltage region interface. The type of level converters to be used will depend on how the block interfaces to surrounding blocks. Ample previous work on level converter circuits offers many choices for implementing the required interfaces.

VI. Conclusion

Local voltage dithering provides a flexible approach for saving energy in fixed throughput systems. When saving energy becomes more important than meeting a performance constraint, operating at the minimum energy point in sub-threshold offers the best option. For many applications, sub-threshold operation will only make sense when the rest of the chip is off (i.e., wake-up circuits) or when the entire chip is in sub-threshold (i.e., microsensor node). This paper proposes ultra-dynamic voltage scaling, which combines dithering at high performance and minimum energy operation for low performance scenarios. A 90-nm bulk CMOS test chip, shown in Fig. 14, verifies the energy savings achievable by LVD and UDVS. This flexible approach allows scaling across the full range of supply voltage and frequency in an energy efficient fashion.

Acknowledgment

The authors acknowledge Texas Instruments Incorporated for chip fabrication.

REFERENCES

[1] V. Gutnik and A. Chandrakasan, "Embedded power supply for low-power DSP," *IEEE Trans. VLSI Syst.*, vol. 5, no. 4, pp. 425–435, Dec. 1997.

[2] H. Kawaguchi, G. Zhang, S. Lee, Y. Shin, and T. Sakurai, "A controller LSI for realizing VDD-hopping scheme with off-the-shelf processors and its application to MPEG4 system," *IEICE Trans. Electron.*, vol. E85-C, no. 2, pp. 263–271, Feb. 2002.

[3] B. H. Calhoun and A. Chandrakasan, "Ultra-dynamic voltage scaling using sub-threshold operation and local voltage dithering in 90nm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, 2005, pp. 300–302.

[4] A. Wang, A. Chandrakasan, and S. Kosonocky, "Optimal supply and threshold scaling for subthreshold CMOS circuits," in *Symp. VLSI Circuits*, Apr. 2002.

[5] A. Wang and A. Chandrakasan, "A 180mV FFT processor using subthreshold circuit techniques," in *IEEE ISSCC Dig. Tech. Papers*, 2004, p. 292.

[6] A. Chandrakasan, V. Gutnik, and T. Xanthopoulos, "Data driven signal processing: an approach for energy efficient computing," in *Proc. ISLPED*, 1996, pp. 347–352.

[7] L. Chandrasena, P. Chandrasena, and M. Liebelt, "An energy efficient rate selection algorithm for voltage quantized dynamic voltage scaling," in *Proc. ISSS*, 2001, pp. 124–129.

[8] S. Lee and T. Sakurai, "Run-time power control scheme using software feedback loop for low-power real-time applications," in *Proc. ASP-DAC*, 2000, pp. 381–386.

[9] S. Lee and T. Sakurai, "Run-time voltage hopping for low-power real-time systems," in *Proc. DAC*, 2000, pp. 806–809.

[10] H. Kawaguchi *et al.*, "A 0.5 V, 400 MHz, VDD-hopping processor with zero-VTH FD-SOI technology," in *IEEE ISSCC Dig. Tech. Papers*, 2003, pp. 106–107.

[11] B. Calhoun and A. Chandrakasan, "Characterizing and modeling minimum energy operation for subthreshold circuits," in *Proc. ISLPED*, 2004, pp. 90–95.

[12] B. Zhai, D. Blaauw, D. Sylvester, and K. Flautner, "Theoretical and practical limits of dynamic voltage scaling," in *Proc. DAC*, 2004, pp. 868–873.

[13] L. Chandrasena and M. Liebelt, "Energy minimization in dynamic supply voltage scaling systems using data dependent voltage level selection," in *Proc. ISCAS*, 2000, pp. 525–528.

Benton H. Calhoun (S'05) received the B.S. degree in electrical engineering with a concentration in computer science from the University of Virginia, Charlottesville, in 2000, and the M.S. degree from the Massachusetts Institute of Technology, Cambridge, in 2002. He is currently pursuing the Ph.D. degree at the Massachusetts Institute of Technology.

His research interests include leakage reduction, sensor networks, energy-efficient circuits, and sub-threshold operation.

Anantha P. Chandrakasan (M'95–SM'01–F'04) received the B.S, M.S., and Ph.D. degrees in electrical engineering and computer sciences from the University of California, Berkeley, in 1989, 1990, and 1994, respectively.

Since September 1994, he has been with the Massachusetts Institute of Technology, Cambridge, where he is currently the Joseph F. and Nancy P. Keithley Professor of Electrical Engineering and Computer Science. His research interests include low-power digital integrated circuit design, wireless microsensors, ultra-wideband radios, and emerging technologies. He is a co-author of *Low Power Digital CMOS Design* (Kluwer, 1995) and *Digital Integrated Circuits* (Pearson Prentice-Hall, 2003, 2nd edition). He is also a co-editor of *Low Power CMOS Design* (IEEE Press, 1998) and *Design of High-Performance Microprocessor Circuits* (IEEE Press, 2000).

Dr. Chandrakasan has received several awards, including the 1993 IEEE Communications Society's Best Tutorial Paper Award, the IEEE Electron Devices Society's 1997 Paul Rappaport Award for the Best Paper in an EDS publication during 1997, and the 1999 Design Automation Conference Design Contest Award. He has served as a technical program co-chair for the 1997 International Symposium on Low Power Electronics and Design (ISLPED), VLSI Design '98, and the 1998 IEEE Workshop on Signal Processing Systems. He was the Signal Processing Sub-committee Chair for ISSCC 1999–2001, the Program Vice-Chair for ISSCC 2002, the Program Chair for ISSCC 2003, and the Technology Directions Sub-committee Chair for ISSCC 2004–2005. He was an Associate Editor for the IEEE JOURNAL OF SOLID-STATE CIRCUITS from 1998 to 2001. He serves on the SSCS AdCom and is the meetings committee chair. He is the Technology Directions Chair for ISSCC 2006.