

Technologies for Ultradynamic Voltage Scaling

Circuits such as logic cells, static random access memories, analog-digital converters and dc–dc converters can be used as building blocks for applications that can function efficiently over a wide range of supply voltages.

By ANANTHA P. CHANDRAKASAN, *Fellow IEEE*, DENIS C. DALY, *Member IEEE*, DANIEL FREDERIC FINCHELSTEIN, *Member IEEE*, JOYCE KWONG, *Student Member IEEE*, YOGESH KUMAR RAMADASS, *Member IEEE*, MAHMUT ERSIN SINANGIL, *Student Member IEEE*, VIVIENNE SZE, *Student Member IEEE*, AND NAVEEN VERMA, *Member IEEE*

ABSTRACT | Energy efficiency of electronic circuits is a critical concern in a wide range of applications from mobile multimedia to biomedical monitoring. An added challenge is that many of these applications have dynamic workloads. To reduce the energy consumption under these variable computation requirements, the underlying circuits must function efficiently over a wide range of supply voltages. This paper presents voltage-scalable circuits such as logic cells, SRAMs, ADCs, and dc–dc converters. Using these circuits as building blocks, two different applications are highlighted. First, we describe an H.264/AVC video decoder that efficiently scales between QCIF and 1080p resolutions, using a supply voltage varying from 0.5 V to 0.85 V. Second, we describe a 0.3 V 16-bit microcontroller with on-chip SRAM, where the supply voltage is generated efficiently by an integrated dc–dc converter.

KEYWORDS | Integrated circuits; low power; ultradynamic voltage scaling

Manuscript received February 18, 2009; revised August 6, 2009 and September 7, 2009. Current version published January 20, 2010. This work was funded by the Defense Research Advanced Projects Agency (DARPA), the Focus Center for Circuit and System Solutions (C2S2), one of five research centers funded under the Focus Center Research Program, a Semiconductor Research Corporation program, Nokia and Texas Instruments (TI). The work was also supported by the Intel Ph.D. Fellowship Program, the Texas Instruments Graduate Women's Fellowship for Leadership in Microelectronics, and Natural Sciences and Engineering Research Council of Canada (NSERC) Fellowships. Chip fabrication was provided by National Semiconductor and Texas Instruments.

A. P. Chandrakasan, D. C. Daly, J. Kwong, Y. K. Ramadass, M. E. Sinangil, V. Sze, and **N. Verma** are with the Massachusetts Institute of Technology, Cambridge, MA 02142-1479 USA (e-mail: anantha@mit.edu; ddaly@alum.mit.edu; jyskwong@mit.edu; ryogesh@mit.edu; sinangil@mit.edu; sze@alum.mit.edu; nverma@princeton.edu).

D. F. Finchelstein is with Nvidia, Santa Clara, CA 95050 USA (e-mail: dfinchel@alum.mit.edu).

Digital Object Identifier: 10.1109/JPROC.2009.2033621

I. INTRODUCTION AND MOTIVATION

As we continue to target severely energy-limited applications, technology scaling, circuit topologies, and architecture trends have all aligned to specifically target low-power trade-offs through the use of fine-grained parallelism [1], [2] and low-voltage operation [3]. Power-management has evolved from static custom-hardware optimization to highly dynamic run-time monitoring, assessing, and adapting of hardware performance and energy with precise awareness of the instantaneous application demands.

In applications such as mobile and embedded computing, dynamic voltage scaling (DVS) [4], [5] serves as an effective means to reduce power through voltage scaling in response to varying performance requirements. However, a far more diverse and energy-constrained set of applications are emerging, including implanted biomedical, remote wireless sensing, and rich mobile multimedia, that have complex use-cases and stringent power limitations. These require ultradynamic voltage scaling (UDVS), where the range in ratio of supply-voltage to threshold-voltage must be extended aggressively to yield orders of magnitude in energy savings or performance increase.

In these systems, ultralow-voltage operation is required in addition to high-speed, making scalability and reliability the primary circuit concerns. Despite issues of variation and leakage-currents, advanced technologies are increasingly important in UDVS systems due to the elevating instantaneous performance demands and the increasing prevalence of energy-constraints in high-volume applications. Conventional DVS design requires moderate scalability in power-delivery and circuit V_{DD} , largely permitting the use of standard architectures and topologies. However, UDVS design requires new logic design

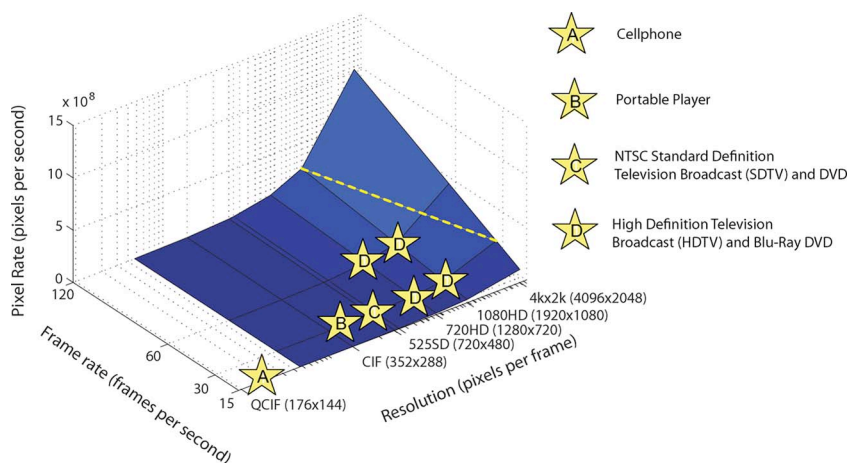


Fig. 1. Performance requirements for various applications based on frame rate and resolution [6]. Yellow dashed line shows limit of H.264/AVC standard. Next-generation standard is expected to reach above this line.

methodologies, new circuit topologies, and more aggressive use of parallelism in order to ensure reliability over a wide operating range. The necessary approach combines balanced circuit optimization considering device trade-offs over an extreme biasing range; application of specialized processing algorithms and architectures to broadly exploit parallelism; and reconfigurable hardware to selectively improve the power-performance trade-offs of the circuits in response to the target operating point.

We introduce the challenges of UDVS systems more concretely in the video decoding and biomedical monitoring applications described below. In the remainder of the paper, specific circuit components and techniques necessary to enable UDVS systems are considered starting with digital logic design and SRAMs. Subsequently, the practical and essential issue of efficient power delivery over an extreme range is considered. Then, the integration of highly scalable analog circuits is presented, and finally, the system applications are considered in detail.

A. Video Decoding

In video decoding, the frame rate and resolution of the playback video dictates the performance requirement of the video decoder hardware. Over the past years, video has become increasingly ubiquitous due to the reduction in storage and transmission costs. The number of different types of video content has been growing rapidly ranging from professional cinema to news reports to, most recently, user-generated content. In addition, the numerous modes of transmission of the video have also expanded from broadcast and playback of local storage material (e.g., DVD), to streaming across the internet and cellular network. Both of these factors cause the frame rate and resolution of today's video content to vary widely. Fig. 1 shows the resolutions and frame rates for different applications and their corresponding performance require-

ment (in terms of pixels/s). For instance, high definition (e.g., 720 or 1080 HD) is used for playback of movies and broadcast television on a high resolution monitor. A higher frame rate [e.g., 60 or 120 frames per second (fps)] is used for high-action sports. Video conferencing and streaming media can be done at lower resolutions (e.g., CIF or VGA) and frame rates (e.g., 15 or 30 fps) for display on a phone.

Accordingly, the current state of the art video coding standard H.264/AVC [7] was designed to support a wide range of resolutions and frame rates as seen in Fig. 1. H.264/AVC supports videos from QCIF (176×144) at 15 fps [380 kpixels/s] up to $4 \text{ k} \times 2 \text{ k}$ (4096×2048) at 30 fps [252 Mpixels/s]; the performance requirement for $4 \text{ k} \times 2 \text{ k}$ at 30 fps is $662\times$ greater than QCIF at 15 fps. It is likely that in the next-generation standards [8], [9], both the lower and upper bound of this range will be increased supporting QVGA (320×240) at 24 fps [1.92 Mpixels/s] up to $4 \text{ k} \times 2 \text{ k}$ (4096×2048) at 120 fps [1 Gpixels/s], which covers a performance range of more than $500\times$. In addition to resolution and frame rate, workload variation can occur frame-to-frame depending on the form of compression used. A highly scalable video decoder is needed to support the wide variety of H.264/AVC sequences.

The use of video playback on handheld battery-operated devices is increasingly common. It is expected that a video decoder on a cellphone can playback different types of video under various use cases. For instance, it should be able to playback low to medium resolution/frame rate videos locally on the phone that perhaps were transmitted over a low bandwidth network; with the growing popularity of video capture on a cellphone, it may also be convenient to be able to connect the phone to a monitor for playback of high resolution and fast frame rate sequences. Having a single video decoder ASIC, which is scalable and supports all these applications, is convenient and cost effective.

Consequently, it is important to minimize and scale the power across this wide range. UDVS is an effective method to support the more than $100\times$ workload variation due to video content in an energy efficient manner. Section VI-A will describe how increased concurrency can improve power-performance trade-off of video decoding. This approach can also be applied to the video encoding hardware; unlike video decoding, where the video dictates the performance requirement, in video encoding, the user has the ability to select the performance-power point depending on the desired application.

B. Medical Monitoring

Likewise, UDVS enables sensor processors to support large workload variation. Several sensor processors have recently demonstrated voltage scalability down to the subthreshold region [10]–[12]. Consider incorporating a similar UDVS-capable processor into a platform for mobile medical monitoring, where intelligent sensors acquire biological signals from a subject through an analog front-end, perform local processing, then transmit results via a radio [13]. Since algorithms are continually evolving, having programmability in the circuits, especially for digital processing, is highly desirable. To extend the operating lifetime of this entire system from a limited energy source, a subset of vital signs can be observed in normal situations, while more detailed monitoring and analysis can be enabled should irregularities occur.

As an example, Fig. 2 shows three different scenarios for monitoring cardiac activity with varying real-time processing demands. Algorithms for each scenario were profiled on a processor based on the MSP430 instruction set [12], [14] with additional application-specific hardware. Pulseoximetry measures the subject's blood oxygen saturation and heart rate; the algorithm described in [15] requires a clock frequency of 331 kHz. This is compatible with minimum energy operation on low-voltage processors such as [12], which achieved optimum energy/cycle at




Fig. 2. Scenarios for monitoring cardiac activity with varying real-time processing demands. For each application, locations of electrodes/probes on the body are shown, as well as the required clock frequency of the sensor processor. (Photos courtesy of GANFYD.)

$V_{DD} = 500$ mV with the corresponding speed of 434 kHz, as detailed in Section VI-B. At the next level of monitoring, the subject's electrocardiogram (ECG) based on a subset of leads is analyzed. In particular, the algorithm from [16] extracts key features from the ECG and performs basic arrhythmia detection. Lastly, the full set of signals in a standard 12-lead ECG provides the highest level of detail. Several algorithms have been developed for automated analysis of the 12-lead ECG, as examined in [17]. In this instance, the work in [18] is applied to extract the QT interval and the QT dispersion, which, if abnormally prolonged, is associated with arrhythmia, fainting, and other adverse effects [19], [20]. Because the computations for one lead require data from other leads, it is preferable to consolidate the processing on a single node running at a higher clock rate. Overall, the three scenarios outlined here span a $78\times$ range in performance, which can be supported in an energy-efficient manner through UDVS.

II. ULTRALOW-VOLTAGE LOGIC DESIGN

A. Challenges and Approaches

In order to support a UDVS system, logic circuits must be capable of operating across a wide voltage range, from nominal V_{DD} down to the minimum energy point which optimizes the energy per operation. This optimum point typically lies in the subthreshold region [21], below the transistor threshold voltage (V_t). Although voltage scaling within the above-threshold region is a well-known technique [4], [22], extending this down to subthreshold poses particular challenges due to reduced I_{ON}/I_{OFF} and process variation. In subthreshold, drive current of the *on* devices (I_{ON}) is several orders of magnitude lower than in strong inversion. Correspondingly, the ratio of active to idle leakage currents (I_{ON}/I_{OFF}) is much reduced. In digital logic, this implies that the idle leakage in the *off* devices counteract the *on* devices, such that the *on* devices may not pull the output of a logic gate fully to V_{DD} or ground. Moreover, local process variation can further skew the relative strengths of transistors on the same chip, increasing delay variability and adversely impacting functionality of logic gates. A number of effects contribute to local variation, including random dopant fluctuation (RDF), line-edge roughness, and local oxide thickness variations [23]. Effects of RDF, in which placement and number of dopant atoms in the device channel cause random V_t shifts, are especially pronounced in subthreshold [24] since these V_t shifts lead directly to exponential changes in device currents. The impact on logic gate functionality is illustrated in Fig. 3, by the random perturbations in the voltage transfer curve (VTC) of a two-input NAND gate at 0.3 V.

To address these challenges, logic circuits in sub- V_t should be designed to ensure sufficient I_{ON}/I_{OFF} in the presence of global and local variation. For example, the

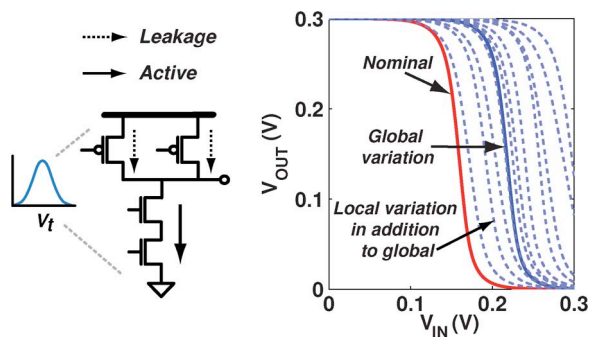


Fig. 3. Effects of variation and reduced I_{ON}/I_{OFF} on sub- V_t NAND voltage transfer curve.

authors of [25] provided a logic gate design methodology accounting for global process corners, and identified logic gates with severely asymmetric pullup/pulldown networks which should be avoided in sub- V_t . In [26], analytical models were derived for the output voltage and minimum functional V_{DD} of circuits, such as in register files [27], where many parallel leaking devices oppose the active device. One approach to mitigate local variation is to increase the sizes of transistors [28] at a cost of higher leakage and switched capacitance. Accordingly, a transistor sizing methodology is described in [29] to manage the trade-off between reducing variability and minimizing energy overhead.

In addition to affecting logic functionality, process variation increases circuit delay uncertainty by up to an order of magnitude in sub- V_t . To illustrate, Fig. 4(a) shows delay histograms of a 65 nm logic timing path at 0.3 V (sub-threshold) and 1.2 V (nominal). Similarly, Fig. 4(b) plots distributions of 30 000 timing paths in a 0.3 V microcontroller under local variation [12]. Each horizontal cross section represents the delay distribution of one

timing path. Note that adjacent paths with similar means can exhibit substantially different variances, as reflected by the lengths of the distribution tails. Statistical methodologies are thus needed to fully capture the wide delay variations seen at very low voltages. A significant body of work, and, recently, commercial computer-aided design tools, have been developed for statistical static timing analysis in the above-threshold region. On the other hand, whereas the relationship between delay and V_t is approximately linear in above-threshold, it becomes exponential in sub- V_t , and timing analysis techniques for low-voltage designs must adapt accordingly. To this end, the work in [30] presented nominal delay and delay variability models valid in both above- and subthreshold regions, while analytical expressions for sub- V_t logic gate and logic path delays were derived in [24].

B. Trends in the Minimum Energy Point

Since UDVS systems spend periods of low activity at the minimum energy point (MEP), trends in the MEP with process technology scaling is examined here. The minimum energy point arises from opposing trends in the dynamic and leakage energy per clock cycle as V_{DD} scales down. The dynamic (CV_{DD}^2) energy decreases quadratically, but in the subthreshold region, the leakage energy per cycle increases as a result of the leakage power being integrated over exponentially longer clock periods. With process scaling, the shrinking of feature sizes implies smaller switching capacitances and thus lower dynamic energy. At the same time, leakage current in recent technology generations have increased substantially, in part due to V_t being decreased to maintain performance while the nominal supply voltage is scaled down. Fig. 5 examines the net effect of these trends on the energy of a 32-bit adder simulated with predictive technology models [31] and interconnect parasitics. The W/L of devices in the adder are kept constant as the lengths are scaled to the 65 nm, 32 nm, and

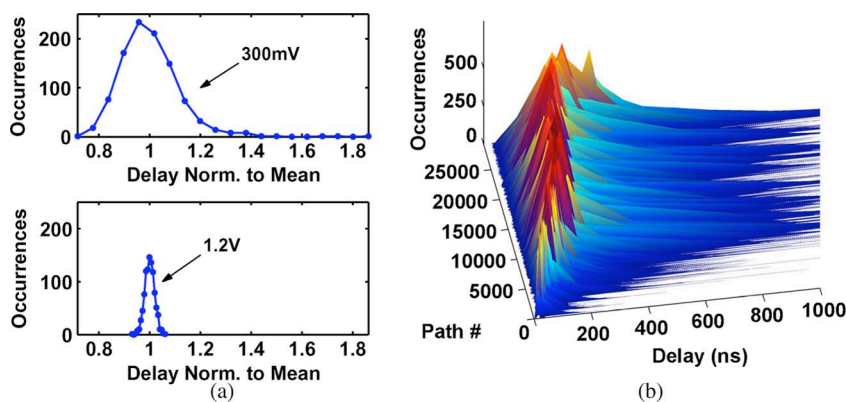


Fig. 4. (a) Delay histograms of a microcontroller logic path, each normalized to sample mean to highlight the difference in variability. Both histograms contain 1000 samples. (b) Delay distributions of 30 k microcontroller timing paths at 300 mV, fast corner. Each horizontal cross section represents distribution of one path.




Fig. 5. Trend in minimum energy point of a 32 b adder with process scaling using predictive models [31].

22 nm nodes. At nominal V_{DD} (0.8 V to 1 V), the reduction in active energy with process scaling is apparent. Importantly, the MEP occurs at a higher voltage at the deeply scaled nodes, due to the larger relative contribution of leakage energy. Nevertheless, in this instance, the MEP continues to occur in the sub- V_t region for several technology nodes, suggesting that low-voltage operation will remain a viable strategy for energy reduction.

C. Survey of Ultralow-Voltage Logic Designs

The idea of exploiting subthreshold operation for low-power circuits was presented by Vittoz in the 1960s [32]. Further, Swanson and Meindl [33] derived the minimum V_{DD} at which CMOS digital circuits can function, and demonstrated an inverter VTC down to 0.2 V. In recent years, researchers have demonstrated a variety of systems functioning at very low voltages. For example, a sub- V_t DLMS filter was designed for a hearing-aid application, and an 8×8 carry save array multiplier test-chip was fabricated in $0.35 \mu\text{m}$ [34]. This test-chip published in 2003 explored adaptive body biasing and operated down to 0.3 V. In 2004, a 180 mV FFT processor in $0.18 \mu\text{m}$ was demonstrated in [25]. The processor featured an energy-aware scalable architecture supporting variable bit precision and FFT length. Device sizing strategies for logic gates accounted for global process variation, and the register file design employed a multiplexer-based hierarchical-read-bitline scheme to address weak I_{ON}/I_{OFF} .

A 200 mV, $0.13 \mu\text{m}$ sub- V_t sensor processor was demonstrated in [10] and featured an 8 b ALU, 32 b accumulator, and 2 kb SRAM. The standard cell library used to implement the processor was carefully selected to exclude cells with a fan-in of more than 2 as well as pass-transistor logic. Another sub-200 mV processor in $0.13 \mu\text{m}$ examined the effectiveness of body biasing in mitigating variation and several logic gate sizing strategies for performance tuning [11].

When combined with the appropriate architectures, ultralow-voltage operation can be applied not only to low-performance, energy-constrained systems, but also to those

with dynamic workloads or high throughput constraints. To accommodate dynamic workloads, a UDVS test-chip was demonstrated in [35] featuring a 32 b Kogge-Stone adder in 90 nm. By operating at the minimum energy point of 330 mV instead of the nominal V_{DD} , the test-chip achieved $9\times$ energy savings at reduced performance scenarios. When higher speeds are needed, V_{DD} can be dynamically elevated up to 1.1 V for $6800\times$ performance increase. Likewise, an FIR filter in [36] achieved functionality from 1.2 V down to 85 mV with 6 orders of magnitude in performance scaling. Since the relative strengths of PMOS and NMOS can change as V_{DD} scales, the FIR filter employed body biasing to adjust device strengths in order to improve noise margin at very low voltages. Finally, systems with constant, high-throughput requirements can leverage extreme parallelism to compensate for speed decrease at low V_{DD} . A 400 mV baseband radio processor [37], for instance, supported 500 M samples/s throughput by distributing computations to 620 parallel hardware blocks. Further, the optimal degree of parallelism and number of pipeline stages for lowest power at constant throughput, as well as the choice of logic style to improve speed in subthreshold, were explored in [38].

Looking forward, technology scaling enables reduced CV_{DD}^2 energy and increased density, but presents heightened process variation. Recently in 2008, a 320 mV 411 GOPS/watt motion estimation accelerator in 65 nm [39] employed optimized datapath circuits to address variation and weak I_{ON}/I_{OFF} , as well as to improve performance. A 65 nm system-on-a-chip was presented in [12] which featured a 16-b microcontroller, a 128-kb 8T SRAM, and an on-chip switched capacitor dc-dc converter. The system demonstrated approaches for library design and timing analysis, as well as circuit techniques to enable sub- V_t operation down to 300 mV.

III. VOLTAGE SCALABLE SRAM

The ability to operate SRAMs at low supply-voltages is critical for minimization of both active and leakage energy. However, as V_{DD} is scaled, SRAMs posed the first point of failure, and, as a result, specialized circuits and architectures are required for UDVS SRAMs that specifically address the severe effects of variation at low-voltages. Since SRAM design trade-offs vary greatly over the UDVS operating range, and, further, since their density remains a critical concern, care must be taken to ensure that ultralow-voltage circuit assists do not impose excessive performance, power, or density overheads in any of the operating regimes. As a result, UDVS SRAM design requires a combination of balanced optimization to support operation under diverse biasing conditions, alternate bit-cell and architecture topologies to specifically relieve over-constrained trade-offs at low-voltages, and selective use of reconfigurability to address the most urgent needs in each operating regime.

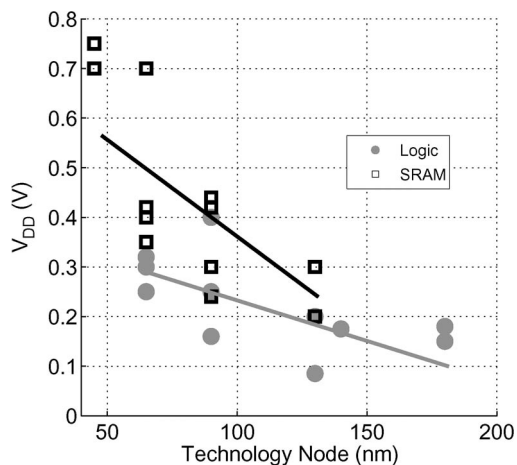


Fig. 6. Minimum reported V_{DD} for recent designs, highlighting UDVS limitation posed by SRAMs compared with logic.

Fig. 6 shows the minimum V_{DD} for recently reported SRAM and logic designs, highlighting the rapidly widening gap as technology scales. As shown, SRAMs pose a critical limitation in UDVS systems, and this is attributed to two key challenges when compared with logic: 1) their functionality margins are lower, leading to an increased dependence on variation and 2) their performance degrades more quickly with V_{DD} scaling.

Standard CMOS logic gates are designed to yield noise-margins that are nearly half of V_{DD} . Standard 6T SRAM bit-cells, however, have a read static-noise-margin (SNM) [40] and write-margin that is set by ratiometric interactions between devices, and these are thus greatly reduced. As V_{DD} is scaled, the read and write margins can easily be violated, particularly since bit-cells exhibit extreme variation [41], [42] owing to their use of minimally sized devices and the tendency to maximize array size for improved area efficiency.

Extreme variation also leads to a more severe effect on SRAM performance with respect to V_{DD} . Typical logic paths are composed of several gate-stages, so variation in any single device has reduced impact [24]. However, in SRAM arrays, the critical delay corresponds to bit-line discharge by the worst case bit-cell, precluding the benefit of averaging over multiple stages. Further, in order to manage their high leakage-currents, particularly in low-power data-retention modes [43]–[45], SRAM devices used in low-power applications are often specially designed with high-threshold voltages [46]. This leads to further degradation of the noise-margins and increased sensitivity to variation at low-voltages.

A. Low-Voltage Array Design

The emphasis in typical UDVS systems on ultralow-energy modes motivates bit-cell design optimized preferentially for low-voltage operation. Though enhancements

via peripheral assists are desirable to minimize the array area overhead, unfortunately, the 6T bit-cell suffers from an unavoidable trade-off between its drive-strength (i.e., read-current, I_{RD}), which requires strong bit-line access-devices, and its read SNM, which requires weak access-devices [47]. Hence, bit-cell design to overcome the trade-off between I_{RD} and read SNM is important.

1) *Non-Buffered-Read Bit-Cells*: Minimally, the read SNM trade-off can be alleviated by ensuring that the bit-line does not remain statically at V_{DD} during a read-access, but rather that it discharges rapidly to avoid the disruptive effect on the internal bit-cell storage nodes. As shown in Fig. 7(a), this can be accomplished by strengthening the cell I_{RD} , via device up-sizing, and, simultaneously, reducing the read bit-line (RdBLT) capacitance. Since the read bit-line voltage-swing is intentionally enhanced, data-sensing can be performed through a large-signal detection circuit. As a result, single-ended sensing can be achieved, precluding the need to up-size the second pull-down path. Hence, aggressive asymmetric up-sizing of $M6/M2$ by $3.7\times$ increases the mean bit-cell I_{RD} and greatly reduces the effect of variation at low-voltages, leading to 22% area savings (compared with symmetric sizing) and 1 GHz operation down to 0.7 V in a 45 nm design [48].

Further voltage scalability can be afforded by eliminating the read SNM entirely. The 7T cell in [47] [Fig. 7(b)] also employs single-ended read but uses an additional device ($M7$) to break the problematic cross-coupled feedback path that can flip its stored data state. For reads, WL is asserted (WLB lowered), disabling the pull-down path through $M7/M1$. As a result, even during disruptive transients on NT, NC remains high, albeit *dynamically*. To minimize susceptibility to leakage currents during this time, however, rapid RdBL discharge must be ensured. For the design of [47], 8 cells are used per RdBLT. Efficient layout of all of the sense-amplifiers required is achieved by exploiting the *L*-shaped layout of the bit-cell. The bit-cell layout is shown in Fig. 7(b), and its *L*-shape is obtained in order to accommodate the additional device $M7$. As a result, the 8×2 layout block shown is used, and the NMOS and PMOS devices for the two sense-amplifiers are distributed among the spaces in between the bit-cells, which provide NWELL and non-NWELL regions. The 90 nm design in [47] achieves operation down to 0.5 V with an access-time of 20 ns.

Another approach to increasing the read-margin of non-buffered-read cells is to selectively adjust the switching threshold of the cross-coupled inverter structure in order to make the stored data-state more difficult to disrupt. The design in [50] uses four additional devices (leading to a 10T cell) to implement hysteresis in the transfer functions of the cross-coupled inverters. As a result, the read-margin is widened, and operation down to 0.175 V is achieved for a 4 kb array in a 130 nm process.

2) *Buffered-Read Bit-Cells*: The functionality of non-buffered-read bit-cells depends on dynamic capacitive discharge rates or ratiometric device strengths, which are

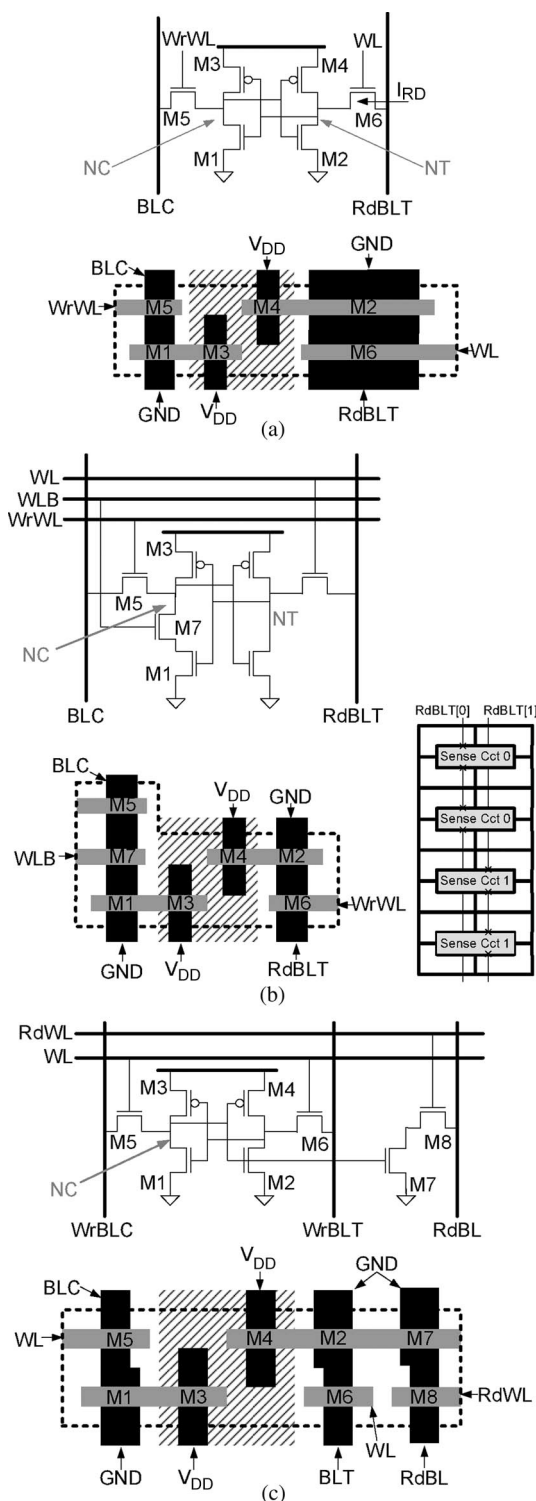


Fig. 7. Ultra-low-voltage bit-cells utilizing: (a) 6T topology [48], (b) 7T topology [47], and (c) 8T topology [49].

difficult to ensure in low-speed, high-variability ultralow-voltage modes. Further, strict limits on the maximum number of cells tolerable per read bit-line improves performance, but degrades array area efficiency. Accordingly, the read-buffer formed by M7–8 of the 8T cell in Fig. 7(c) [49] provides a means to isolate NT/NC from disruptive transients during reads without introducing additional dynamic nodes. Nonetheless, as with 6T cells, the RdBL must be able to *dynamically* retain a logic “1” level when it is not actively pulled low by the accessed cell. Hence, to mitigate bit-line leakage currents from unaccessed cells sharing RdBL, 10T cells, operating below 0.4 V, are proposed with enhanced read-buffers that gate their own leakage, supporting the integration of up to 1024 cells per RdBL [51], [52]. Other modifications include the use of differential read-buffers [53], [54] and internal switches to enable column-wise cell selection in order to support a column interleaved layout in which only desired cells in the row are accessed during both reads and writes [54].

To evaluate the merits of the 8T topology for low-voltage, its area overhead must be compared against the approach of up-sizing a standard 6T cell for reduced variation. Further, for comparable I_{RD} , its read-buffer (M7–8) must be sized equivalently to the 6T pull-down paths, leaving very little area for its storage-cell (M1–6). However, the storage-cell is required to meet only the hold SNM, rather than the far more stringent read SNM, and it can thus tolerate much higher variability. Fig. 8(a) shows the result of Monte Carlo simulations with a 65 nm technology comparing the read SNM of 6T cells and the hold SNM of 8T cells with equivalent area and driver sizes. The improved operating-margin of the 8T topology during low-voltage operation points to its benefit in UDVS systems.

In addition to better read stability, the 8T cell provides the possibility for improved write-margin and performance. Write-ability relies on the storage-cell access-devices (M5–6) being stronger than the load-devices (M3–4) so that the required storage-node (NT/NC) can be pulled low by WrBLT/WrBLC in order to over-write old data. In the 6T cell, the access-devices must simultaneously be kept much weaker than the driver-devices (M1–2) to avoid disruptive read transients on NC/NT. This leads to a strict two-sided strength constraint for the access-devices. In the 8T cell, however, M5–6 can be made arbitrarily strong (especially through device V_t engineering). Similarly, the need to limit the strength of M5–6 for the 6T cell’s read SNM also degrades I_{RD} . Since this constraint does not apply to the read-buffer access device (M8), the 8T cell also benefits from greatly improved performance.

In addition to alleviating the basic operating constraints, the 8T cell can also take advantage of write-functionality assists that have been used to extend the operating margins of 6T cells [56], thereby supporting robust operation at aggressively low-voltages. The bit-cell

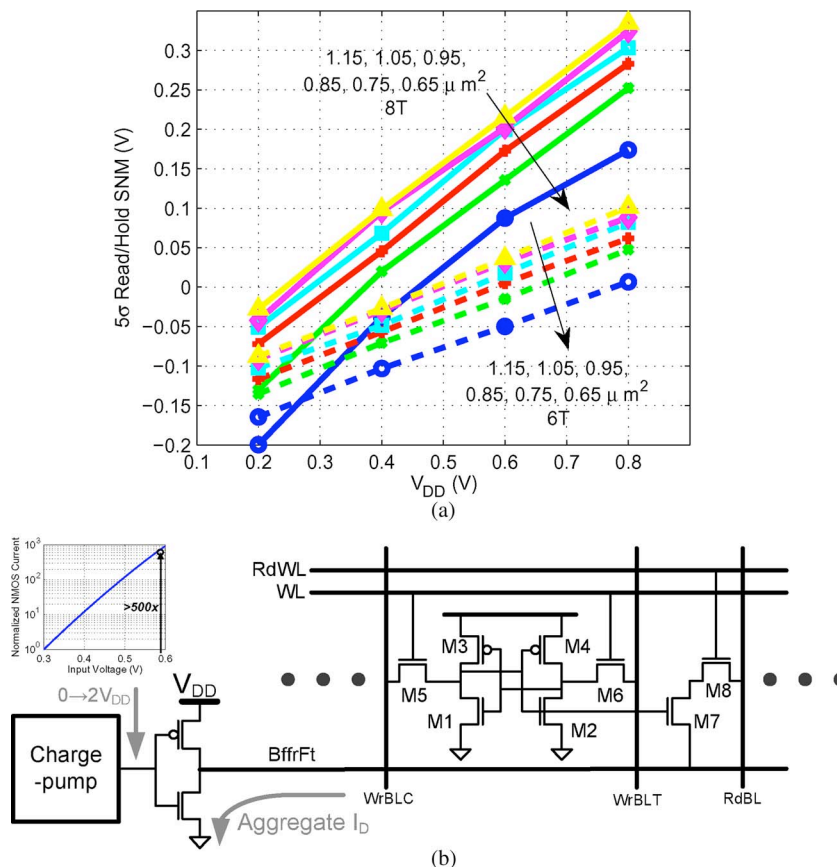


Fig. 8. 8T bit-cell (a) low-voltage analysis compared to (b) implementation for 0.35 V 65 nm SRAM [55].

in Fig. 8(b) operates below 0.35 V in a 65 nm technology [55] by applying write-margin enhancing and bit-line leakage reducing circuit-assists to an 8T cell. Importantly, however, since these perform selective biasing on the bit-cell from the periphery, their area overhead is minimal. For instance, during writes, the V_{DD} node is actively lowered for all cells in the row to specifically weaken the load-devices (M3–4) and improve write-ability. In this case, all cells in the row share V_{DD} and are over-written together. To manage bit-line leakage during reads, the BffrFt node, which is also shared by all cells in the row, is held high until the row is accessed. As a result, the unaccessed read-buffers impose no subthreshold leakage currents on RdBL. A limitation that must be overcome, however, is that the peripheral BffrFt driver [shown in Fig. 8(c)], must sink the read-current of all cells in the accessed row. Simple up-sizing of the driver to provide the required current would prohibitively increase area and leakage-power. Instead, the peripheral driver uses the charge-pump shown, which, in weak-inversion, yields a 500× drive current increase.

3) *Low-Voltage Sensing:* In UDVS SRAMs, the reduced voltage-levels, increased impact of bit-line leakage, and

lower read-currents lead to drastically degraded bit-line sensing margins. Further, the most promising bit-cell topologies for low-voltage operation employ single-ended read-ports, eliminating the noise-margin and sensitivity benefits of differential sensing. 9T buffered-read bit-cells with differential read-ports have been proposed [57], though the odd number of devices these require makes it challenging to achieve area-efficient layouts.

To reduce sensitivity to bit-line leakage, the leakage replica sensing scheme [52], shown in Fig. 9(a) can be used. In the bit-cell of this design, leakage compromises the logic “0” level. Hence, a replica column is used to approximately track this level and, accordingly, adjust the trip-point of the sensing inverters. Although this reduces the sensing margin required by calibrating the sensing levels, it does not alleviate unpredictable shifts in the sense amplifier trip-point caused by variation. The technique of sense-amplifier redundancy [55] [shown in Fig. 9(b)] uses several smaller sense-amplifiers and selectively enables one that has sufficiently small offset. Although the smaller size of the individual sense-amplifiers leads to increased variation and, therefore, increased failure probability ($P_{ERR,N}$), the overall sensing network failure probability ($P_{ERR,TOT}$) depends on the joint failure

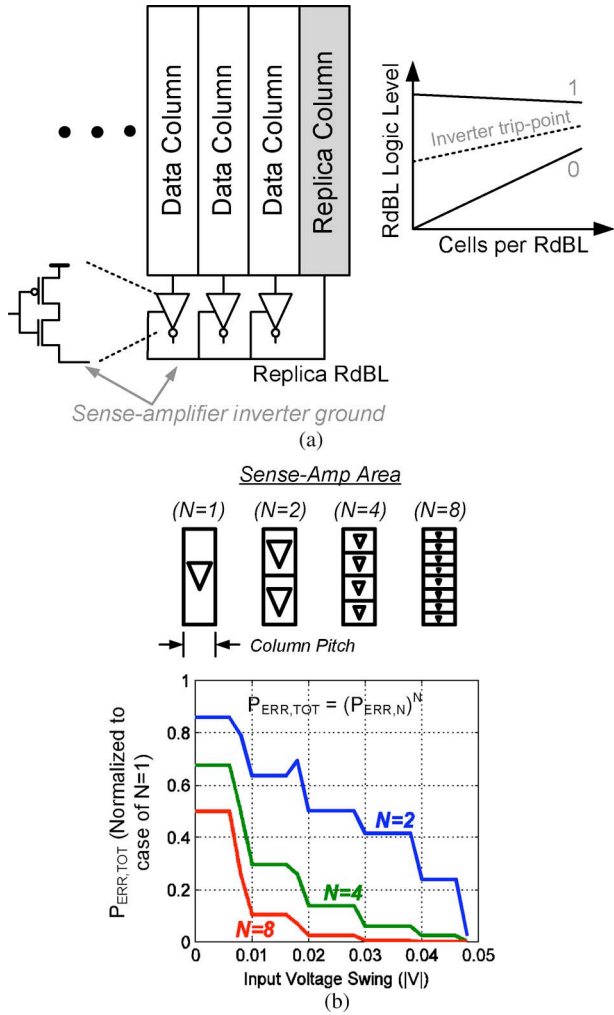


Fig. 9. Ultralow-voltage sensing techniques using: (a) leakage replica sensing [52] (b) sense-amplifier redundancy [55].

probabilities of the individual sense-amplifiers, and [as shown in Fig. 9(b)], it is thus significantly reduced for the sizes considered.

B. SRAM Design for UDVS

Over the voltage range targeted by UDVS systems, the constraints relevant to SRAM operation change drastically. For instance, from 0.3 V to 1.2 V, the bit-cell I_{RD} 's of a UDVS SRAM in 65 nm CMOS [58] change by over four orders of magnitude. At the extreme low-voltage end, reliable data-retention and sensing requires highly selective bit-cell biasing on a per-cycle basis to specifically address either the read, write, or sensing margin. On the other hand, at the high-voltage end accommodating high I_{RD} 's and minimizing the performance overhead of the low-voltage assists is critical.

SRAM optimization over the extreme range must also consider second-order effects, whose impact can be mild in

one regime but greatly elevated in an alternate regime. For instance, increasing the lengths of the bit-cell read-buffers mitigates variation, having an amplified appeal at low-voltages. However, although length up-sizing reduces variation and even increases the mean I_{RD} at low-voltages (through the reverse short-channel-effect (RSCE), which effectively reduces device V_i 's), the opposite impact on cell-current is observed in high-voltage high-performance modes. Consequently, bit-cell sizing optimizations must consider the total effect over the entire range, as shown in Fig. 10(a) [58].

Additional options to optimize operation over the wide range exist at the array architecture level. For instance, the bit-line leakage and performance degradation concerns at low-voltages are addressed in [59] by using very short (8 cell) read-bit-lines. However, to maximize density the write-bit-lines are amortized over 512 cells, and scalable functionality is achieved from 1.2 V down to 0.41 V in 65 nm CMOS.





Fig. 10. (a) Effect of read-buffer device channel length on worst case I_{RD} and (b) reconfigurable write-assist scheme [58].

For extreme voltage scalability, however, reconfigurable circuit-assists are employed in [58] to enable operation from 1.2 V down to 0.25 V. For instance, to optimize write-ability and power consumption, a dynamically selectable write-assist is used to actively lower, passively float, or leave unchanged the bit-cell supply voltage (MCHd) depending on the operating voltage [as shown in Fig. 10(b)]. Further reconfigurable sensing networks dynamically switch between sense-amplifier topologies to optimize bit-line sensing.

IV. POWER DELIVERY FOR ADAPTIVE SYSTEMS

While dynamic voltage scaling is a popular method to minimize power consumption in digital circuits given a performance constraint, the same circuits are not always constrained to their performance-intensive mode during regular operation. There are long spans of time when the performance requirement is highly relaxed. There are also certain emerging energy-constrained applications where minimizing the energy required to complete operations is the main concern. For both these scenarios, operating at the minimum energy operating voltage of digital circuits has been proposed as a solution to minimize energy [60]. The minimum energy point (MEP) is defined as the operating voltage at which the total energy consumed per desired operation of a digital circuit is minimized. The MEP is not a fixed voltage for a given circuit, and can vary widely depending on its workload and environmental conditions (e.g., temperature). Any relative increase in the active energy component of the circuit due to an increase in the workload or activity of the circuit decreases the minimum energy operating voltage. On the other hand, a relative increase of the leakage energy component due to an increase in temperature or the duration of leakage over an operation pushes the minimum energy operating voltage to go up. This makes the circuit go faster thereby not allowing the circuit to leak for a longer time. By tracking the MEP as it varies, energy savings of 50%–100% has been demonstrated [61] and even greater savings can be achieved in circuits dominated by leakage. This motivates the design of a minimum energy tracking loop that can dynamically adjust the operating voltage of arbitrary digital circuits to their MEP.

A. Reference Voltage Selection for Highly Energy-Constrained Applications

To effectively use DVS to reduce power consumption, a system controller that determines the required operating speed of the processor at runtime is needed. The system controller makes use of algorithms, termed voltage schedulers, to determine the operating speed of the processor at run-time. For general purpose processors, these algorithms effectively determine the overall workload of the processor and suggest the required operating speed

to handle the user-requests. Some of the commonly used algorithms have been described in [62]. For DSP systems like video processors, the speed of the system is typically measured by looking at the buffer length occupied. Once this operating speed has been determined, the operating voltage of the circuit needs to be changed so that it can meet the required speed of operation. DVS systems usually employ a suitable control strategy to change the load voltage supplied by a dc–dc converter to maintain the operating speed. Reference [63] presents a closed loop architecture to change the output voltage of a voltage scalable dc–dc converter to make the load circuit operate at the desired rate. [5] uses a hybrid approach employing both look-up tables and a phase-locked loop (PLL) to enable fast transitions in load voltage with change in the desired rate. While the look-up table aids in the fast transition, the PLL helps in tracking process variations and operating conditions. For DVS systems, maintaining the operating speed is the primary constraint and hence the control strategies they employ are geared for it. For energy constrained applications however, since the operating speed constraint is relaxed, the reference voltage to the dc–dc converter can be chosen so as to consume the minimum energy possible.

Fig. 11 shows the architecture of a minimum energy tracking loop. The objective of this loop is to first detect the correct minimum energy reference voltage to the dc–dc converter and to then track the MEP of the load circuit as it changes. The load circuit is powered from an off-chip voltage source through a dc–dc converter and is clocked by a critical path replica ring oscillator which automatically scales the clock frequency of the FIR filter with change in load voltage. The energy sensor circuitry calculates on-chip, the energy consumed per operation of the load circuit at a particular operating voltage. It then passes the estimate of the energy/operation (E_{op}) to the energy minimizing algorithm, which uses the E_{op} to suitably adjust the reference voltage to the dc–dc converter. The dc–dc converter then tries to get V_{DD} close to the new reference voltage and the cycle repeats till the minimum energy point is achieved. The only off-chip components of this entire loop are the filter passives of the inductor-based switching dc–dc converter.

The key element in the minimum energy tracking loop is the energy sensor circuit which computes the E_{op} of the load circuit at a given reference voltage. Methods to measure E_{op} , by sensing the current flow through the dc–dc converter's inductor, dissipate a significant amount of overhead power. The approach is more complicated at subthreshold voltages because the current levels are very low. Further, an estimate of the energy consumed per operation is what is required and not just the current which only gives an idea of the load power. Reference [64] estimates E_{op} in a digital fashion that does not require any high-gain amplifiers or analog circuit blocks.

Once the estimate of the energy per operation is obtained, the minimum energy tracking algorithm uses

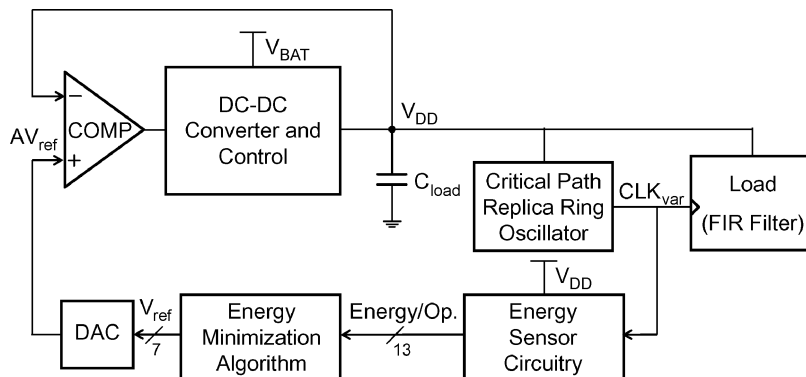


Fig. 11. Architecture of the minimum energy tracking loop.

this to suitably adjust the reference voltage to the dc–dc converter. The minimum energy tracking loop uses a slope-tracking algorithm which makes use of the single minimum, concave nature of the E_{op} vs. V_{DD} curve (see Fig. 5). The voltage step used by the tracking algorithm is usually set to 50 mV. A large voltage step leads to coarse tracking of the MEP, with the possibility of missing the MEP. On the other hand, keeping the voltage step too small might lead to the loop settling at the nonminimum voltage due to errors involved in computing E_{op} [64]. The E_{op} vs. V_{DD} curve is shallow near the MEP and hence a 50 mV step leads to a very close approximation of the actual minimum energy consumed per operation. The MEP tracking loop can be enabled by a system controller as needed depending on the application, or periodically by a timer to track temperature variations.

B. Inductor-Based DC–DC Converter for Minimum Energy Operation

A dc–dc converter that can deliver variable load voltages is one of the key requirements of a UDVS system. Since UDVS systems have widely varying current consumption as their V_{DD} changes, the dc–dc converter must also be able to maintain high efficiency as the load current varies by orders of magnitude. This section talks about the design of the dc–dc converter that enables minimum energy operation. Since, the minimum energy operating voltage usually falls in the subthreshold regime of operation, the dc–dc converter is designed to deliver load voltages from 250 mV to around 700 mV. The power consumed by digital circuits at these subthreshold voltages is exponentially smaller and hence the dc–dc converter needs to efficiently deliver load power levels of the order of micro-watts. This demands extremely simple control circuitry design with minimal overhead power to get good efficiency. The dc–dc converter shown in Fig. 12(a) is a synchronous rectifier buck converter with off-chip filter components and operates in the discontinuous conduction mode (DCM). It employs a pulse frequency modulation

(PFM) [65] mode of control in order to get good efficiency at the ultralow load power levels that the converter needs to deliver. The reference voltage to the converter is set digitally by the minimum energy tracking loop. The comparator compares V_{DD} with this reference voltage and when it is found to be smaller, generates a pulse of fixed width to turn the PMOS power transistor ON and ramp up the inductor current. A variable pulse width generator to achieve zero-current switching is used for the NMOS power transistor. The approximate zero-current switching method used helps to maintain high efficiency while not significantly increasing the control power. The efficiency achieved by the dc–dc converter with change in load power is shown in Fig. 12(b). Owing to the small overhead and low control circuit power, the dc–dc converter was able to achieve greater than 80% efficiency even at $1 \mu\text{W}$ load power levels.

C. Switched Capacitor DC–DC Converter

UDVS systems often require multiple on-chip voltage domains with each domain having specific power requirements. A switched capacitor (SC) dc–dc converter is a good choice for such battery operated systems because it can minimize the number of off-chip components and does not require any inductors. Compared to commonly used linear regulators for on-die voltage conversion, SC converters can provide 15%–30% higher efficiencies when the output voltage delivered is less than half the input voltage [66]. Previous implementations of SC converters (charge-pumps) have commonly used off-chip charge-transfer capacitors [67] to output high load power levels. SC dc–dc converters which integrate the charge-transfer capacitors are described in [68] and [69].

Fig. 13(a) shows the major energy loss mechanisms in a switched capacitor dc–dc converter. SC dc–dc converters are hampered by a fundamental problem of charge transfer using only capacitors and switches. This introduces a linear efficiency loss that is similar to linear regulators. However, with SC converters, it is possible to

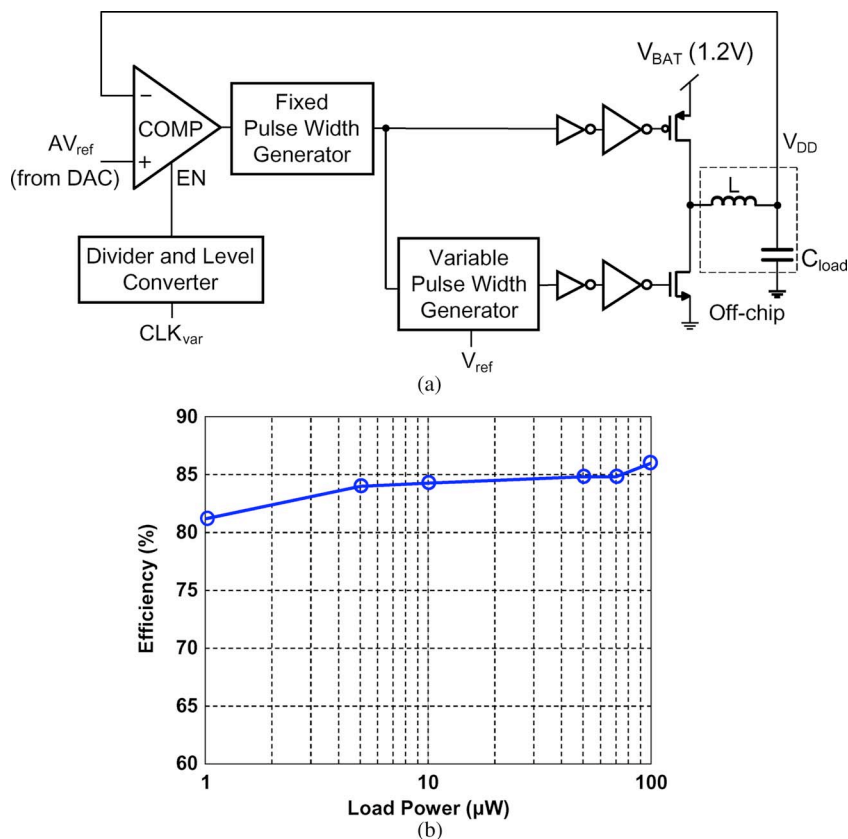


Fig. 12. (a) Architecture of the inductor-based dc-dc converter (b) efficiency plot with change in load power.

switch in different gain-settings whose no-load output voltage is closer to the load voltage desired. The loss due to bottom-plate parasitics occurs because of the energy stored in the parasitic capacitors being dumped to ground every cycle. This loss is more pronounced if on-chip capacitors are used. Common capacitors in CMOS processes may have up to 20% parasitic at the bottom-plate. On top of these losses are the switching and control losses. The efficiency achievable in a switched capacitor system is in general smaller than which can be achieved in an inductor based switching regulator with off-chip passives. Further, multiple gain settings and associated control circuitry is required in a SC dc-dc converter to maintain efficiency over a wide voltage range. However, for on-chip dc-dc converters, a SC solution might be a better choice, when the trade-offs relating to area and efficiency are considered. Further, the area occupied by the switched capacitor dc-dc converter is scalable with the load power demand, and hence the switched capacitor dc-dc converter is a good solution for low-power on-chip applications.

Switched capacitor dc-dc converters that employ closed-loop control to regulate the output voltage as the input voltage varies are described in [70]–[72]. The converter described in [66] regulates the output voltage over a wide range of load voltages with a fixed input

voltage. The converter employs five different gain settings with ratios 1 : 1, 3 : 4, 2 : 3, 1 : 2, and 1 : 3. The switchable gain settings help the converter to maintain a good efficiency as the load voltage delivered varies from 300 mV to 1.1 V. At the core of the system is a switch matrix which contains the charge-transfer capacitors, and the charge-transfer switches. A suitable gain-setting is chosen depending on the reference voltage which is set digitally. A pulse frequency modulation (PFM) mode control is used to regulate the output voltage to the desired value. Bottom-plate parasitics of the on-chip capacitors significantly affect the efficiency of the converter. A divide-by-3 switching scheme [66] was employed to mitigate the effect due to bottom-plate parasitics and improve efficiency. The switching losses are scaled with change in load power by the help of an automatic frequency scaler block. This block changes the switching frequency as the load power delivered changes, thereby reducing the switching losses at low load.

The efficiency of the SC converter with change in load voltage while delivering 100 μW to the load from a 1.2 V supply is shown in Fig. 13(b). The converter was able to achieve > 70% efficiency over a wide range of load voltages. An increase in efficiency of close of 5% can be achieved by using divide-by-3 switching.

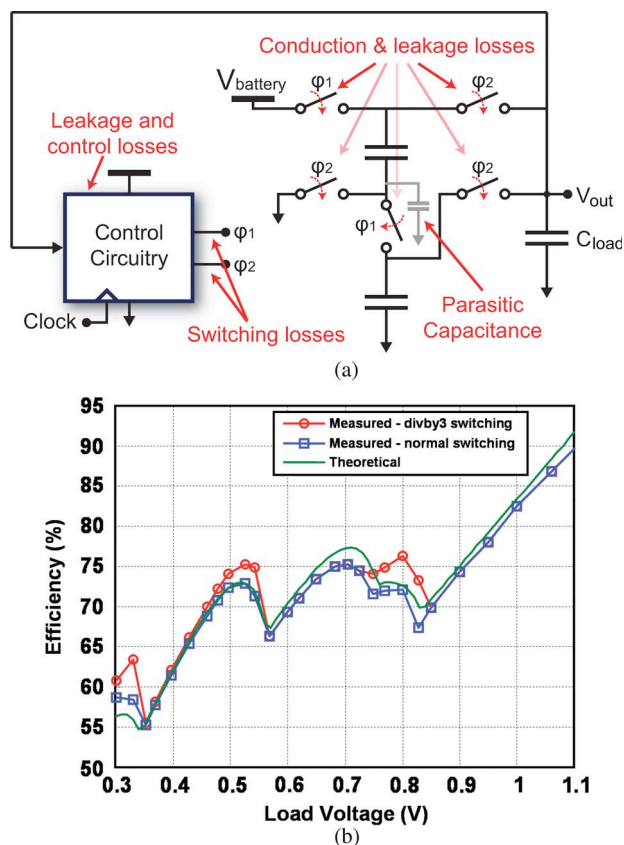


Fig. 13. (a) Energy loss mechanisms in a switched capacitor dc-dc converter (b) efficiency plot with change in load voltage.

V. VOLTAGE-SCALABLE, RECONFIGURABLE ADC

In nearly all systems that interface to the physical world, analog-to-digital converters (ADCs) are required. In the majority of applications involving an ADC, there is usually a fixed performance specification in terms of resolution and input bandwidth that must be met. However, there are some emerging applications in which a reconfigurable ADC is desired, including wireless sensor networks, software defined radios, and other highly reconfigurable chips. Additionally, in the existing commercial ADC market, there are a huge assortment of parts, each optimized for a different, specific performance specification. By implementing performance scaling in part through voltage scaling, a single chip could potentially be sold for a wide range of applications, thereby offering the potential for significant market impact. This section describes ADC techniques to achieve reconfigurability and voltage scalability. As a case study, a voltage scalable flash ADC is described that operates from 0.2 V to 0.9 V.

A. Reconfigurable ADC Techniques

Voltage scaling is a technique that can be leveraged to minimize energy consumption in reconfigurable ADCs;

however, historically, reconfigurability has been implemented through alternate methods, such as varying bias currents, duty cycling and reconfiguring architectures. These methods are important to identify, as they can be used in conjunction with voltage scalability to minimize energy consumption.

Many ADC architectures have traditionally relied upon the use of op-amps, including pipeline and delta-sigma ADCs. In [73], a reconfigurable, op-amp based ADC is presented that can switch between pipeline and delta-sigma modes. The converter achieves a bandwidth range of 0–10 MHz and a resolution range of 6–16 bits. Switches are used to reconfigure the op-amps, comparators, and capacitors depending on the required resolution and yield. In delta-sigma mode, the clock frequency of the converter determines the ADC signal-to-noise ratio (SNR).

For analog circuits like op-amps, bias current scaling only has limited efficacy when reducing ADC sample rates, due to the difficulty in designing analog circuits that remain properly biased over the wide range of current. Instead, rapid power gating and duty cycling can be employed to achieve near linear scaling in power consumption with sampling frequency [74], [75]. In [75], a successive approximation register (SAR) ADC consumes energy only during a conversion, and then completely shuts down between samples, enabling a constant energy per conversion regardless of sampling frequency. Resolution scalability is particularly challenging to implement efficiently, because components like capacitors and transistors are often sized and/or biased to meet the highest resolution requirement. In [75], the ADC supports both 8-bit and 12-bit modes, switching between resolution modes by switching between two bit-resolution comparators. In the 8-bit mode, a lower power bit-resolution comparator can be used because more noise and offset are acceptable. However, the sampling capacitor is still sized for 12 bits of linearity, resulting in energy overhead.

B. Voltage Scalable ADC Techniques

Voltage scaling can offer significant energy savings to reconfigurable ADCs, for realizing both efficient frequency and resolution scalability. A key challenge, however, is to design ADC architectures and circuits that can tolerate a wide range of supply voltages, particularly low voltages. These challenges are similar to the ones posed by advanced CMOS processes.

A key technique to achieve voltage scalability is by employing highly digital circuits and architectures. ADC architectures that fit under this umbrella include SAR, delta-sigma and flash. For example, in [76], a frequency-to-digital delta-sigma modulator is presented that uses only inverters and digital logic gates, operating at a supply voltage of 0.2 V. Pipeline ADCs have also been shown to be tolerant of low voltages through techniques such as switched-op-amp and zero-crossing based ADCs [77].

For highly digital architectures, voltage scaling offers superior energy savings compared with duty cycling when reducing clock frequencies. In [78], a voltage-scalable SAR ADC is presented that operates from 0.5 V to 1 V, with sample rates scaling from 4.1 kS/s to 150 kS/s. When reducing the supply voltage from 1 V to 0.8 V, the sample rate reduces by 1.7× but the power reduces by 2.7×, at no significant degradation in resolution. At low voltages, however, sources of loss such as leakage, increased offsets, and a reduced dynamic range can degrade performance.

Voltage scaling can also be employed to improve energy efficiency when scaling resolution, particularly for low-resolution ADCs (< 75 dB SNR) where thermal noise does not limit performance [79]. Trading voltage for resolution can be achieved in highly digital delta-sigma converters [76], where the oversampling ratio and frequency can be adjusted in tandem with the supply voltage; however, for Nyquist-rate ADCs, efficient resolution scaling is more challenging. Ideally, power consumption should reduce exponentially with reduced bit resolution. In practice, most reconfigurable ADCs allow for only linear or slightly better power savings when reducing resolution. Voltage scaling offers much promise to further approach the ideal exponential relationship between power consumption and bit resolution.

C. Inverter-Based Flash ADC That Leverages Redundancy and Reconfigurability

A highly voltage scalable, flash ADC was recently demonstrated that can operate from 0.2 V to 0.9 V and addresses several of the challenges facing voltage scalable ADC design [80]. Fig. 14 shows a block diagram of the ADC. To ensure operation over such a wide range of supply voltage, the ADC extensively leverages digital circuit blocks combined with calibration. The flash ADC has no bias currents and thus energy is only dissipated through switching events (CV^2) and by leakage currents. By minimizing the use of analog circuits, the architecture is compatible with digital CMOS processes. As true differential architectures are not amenable to ultralow-voltage

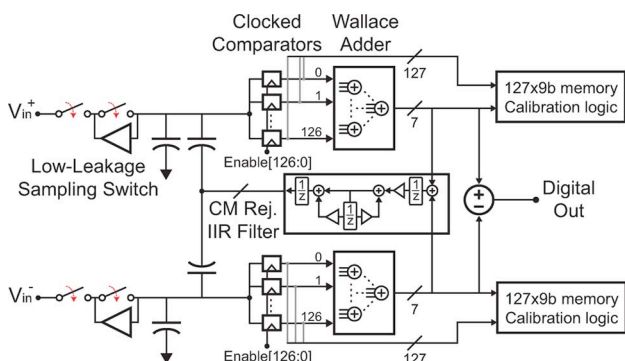


Fig. 14. Block diagram of the flash ADC.




Fig. 15. Maximum (a) sampling frequency and (b) FOM versus supply voltage, indicating presence of minimum FOM at $V_{DD} = 0.4$ V.

integration due to limited voltage headroom, the ADC attempts to mimic the advantages of differential circuits through digital signal processing. The ADC can be configured in either a single-ended or pseudo-differential configuration. It consists of a sampling network, two arrays of 127 dynamic digital comparators and a digital backend. Instead of a traditional reference ladder that draws static current, the ADC uses dynamic comparators with inherent static voltage offsets to generate comparator thresholds. As these thresholds vary die-to-die and with supply voltage, redundancy is exploited to improve linearity [81]. The ADC is calibrated with an off-chip ramp and an estimate of the cumulative distribution function (CDF) of comparator thresholds is generated in on-chip memory. This data is then processed off-chip to determine which subset of the 254 comparators to enable. The comparators include dummy input devices to reduce kickback noise [82].

For sampled ADCs, a key challenge at low voltages is how to mitigate leakage currents [83]. These currents can cause bit errors by discharging nodes such as the sampling capacitor in a flash or SAR ADC or charge-transfer capacitors in a pipelined ADC. This problem is fundamentally related to the ratio of on conductance to off current of a transistor, which significantly degrades at lower supply voltages. In [80], three circuit techniques were applied to improve this ratio: voltage boosting on the sampling switch, stacked sampling switches, and an analog feedback circuit to reduce leakage.

The ADC was fabricated in a 0.18 μm 5M2P CMOS process and occupies 2 mm^2 . The ADC operates from 2 kS/s at 0.2 V to 17.5 MS/s at 0.9 V, as shown in Fig. 15(a). The total power consumption of the ADC at 0.4 V, 400 kS/s is 2.84 μW and 1.66 μW in pseudo-differential and single-ended mode, respectively, of which 135 nW is leakage power. Shown in Fig. 15(b) is the figure of merit (FOM) of the ADC in single-ended mode versus supply voltage. In Fig. 15(b), the ADC is operated at the maximum possible frequency at each voltage. At low voltages, the leakage current degrades the FOM due to low

sampling rates, whereas at high voltages, CV^2 losses degrade the FOM, leading to the emergence of a minimum FOM supply voltage of 0.4 V [84]. At this voltage, the ADC achieves a FOM of 125 fJ/conversion-step in single-ended mode [5.05 effective number of bits (ENOB)] and 150 fJ/conversion-step in pseudo-differential mode (5.56 ENOB). Over the entire operating frequency and voltage range, the ADC achieves an ENOB greater than 4.

The comparator thresholds vary with both temperature and supply voltage, and ADC recalibration is required to maintain linearity. In single-ended 6-bit mode, the ADC ENOB degrades from 5.05 at 25 °C to 4.28 at 75 °C without recalibration. After recalibration the ENOB returns to 5.08.

VI. UDVS SYSTEM APPLICATIONS

A. H.264/AVC Video Decoder

As previously discussed, performance requirements of video decoding can vary widely depending on the application and video content. The prevalence of video on portable battery-operated devices makes it necessary for the video decoder to support this range in an energy efficiency manner using UDVS. A highly scalable low-power H.264/AVC video decoder [85], [86] was designed to meet the performance requirement from QCIF at 15 fps [380 kpixels/s] to 1080p at 30 fps [62 Mpixels/s], which is a range of $164\times$. The video decoder demonstrates that low voltage architectures and circuit optimizations are necessary to improve the power-performance trade-off and enable the effective use of UDVS.

1) *Architecture Optimizations*: Video decoding is a performance-constrained application where the frame rate and resolution dictate the throughput requirement. Increased concurrency is necessary in order to meet this requirement at lower voltages.

a) *Pipelining*: The top-level architecture of the decoder is shown in Fig. 16. At the system level of the decoder, FIFOs of varying depths connect the major processing units: entropy decoding (ED), inverse transform (IT), motion compensation (MC), spatial prediction (INTRA), delocking filter (DB), memory controller (MEM), and frame buffer (FB). The pipelined architecture allows the decoder to process several 4×4 blocks of pixels simultaneously, requiring fewer cycles to decode each frame. The number of cycles required to process each 4×4 block varies for each unit as shown in Table 1. The table describes the pipeline performance for decoding P-frames (temporal prediction). Most of the optimization efforts were focused on P-frame performance, since they occur more frequently than I-frames (spatial prediction) in highly compressed videos.

The luma and chroma components have separate pipelines that share minimal hardware and are mostly

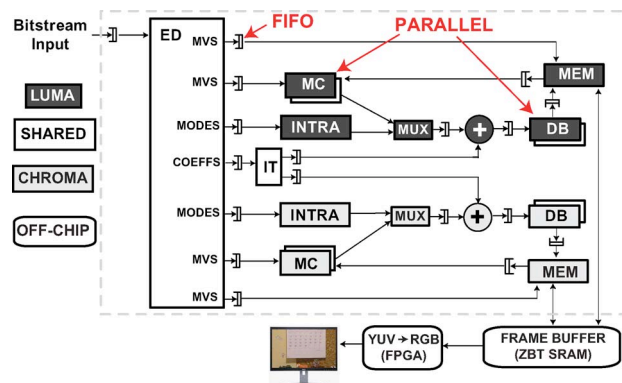


Fig. 16. H.264/AVC decoder architecture.

decoupled from each other. In most cases, the luma and chroma components of each 4×4 block are processed at the same time which enables further cycle count reduction.

The workload of each unit varies depending on the prediction mode as well as the content of the video sequence. To adapt for the workload variation of each unit, FIFOs were inserted between each unit. The FIFOs help to average out the cycle variations which increases the throughput of the decoder by reducing the number of stalls, as described in [87].

b) *Parallelism*: Parallelism can be used within each processing unit to reduce the number of cycles required to process each 4×4 block and balance out the cycles in each stage of the pipeline. This is particularly applicable to the MC and DB units, which were found to be key bottlenecks in the system pipeline. The cost of parallelism is primarily an increase in area. The increase in logic area due to the parallelism in the MC and DB units is about 12%. When compared to the entire decoder area (including on-chip memory) the area overhead is less than 3%.

The motion compensation (MC) unit uses interpolation to predict a 4×4 block in the current frame from pixels in a previous frame to exploit temporal redundancy. To improve the throughput of MC, two identical interpolators operate in parallel as shown in Fig. 17. The

Table 1 Cycles per 4×4 Luma Block for Each Unit in P-Frame Pipeline of Fig. 16, Assuming no Stalling. [] is Performance After Optimizations in Section VI-A1b

Pipeline Unit	Min Cycles	Max Cycles	Average Cycles
ED	0	33	4.6
IT	0	4	1.6
MC	4 [2]	9 [4.5]	4.6 [2.3]
DB	8 [2]	12 [6]	8.9 [2.9]
MEM	8	31	18




Fig. 17. Parallel MC interpolators. Note: the numbered 4×4 blocks reflect the processing order within a macroblock.

first interpolator predicts 4×4 blocks on the even rows of a macroblock, while the second predicts 4×4 blocks on the odd rows. This parallel structure can double the throughput of the MC unit if during each cycle, both motion vectors are available and two new columns of 9 pixels from the frame buffer are available at the inputs of both interpolators. The chroma interpolator is also replicated four times such that it can predict a 2×2 block every cycle.

The boundary strength information of the adaptive delocking filter (DB) is the same for all edges on a given side of a 4×4 block. Accordingly, the DB is designed to have 4 luma and 2 chroma filters running in parallel which share the same boundary strength calculation, and filter an edge of a 4×4 block every cycle. For additional cycle reduction, the luma and chroma filters operate at the same time, assuming the input data and configuration parameters are available.

A luma macroblock (sixteen 4×4 blocks) has 128 pixel edges that need to be filtered; with 4 luma filters, a macroblock takes 32 clock cycles to complete. Unlike previous implementations [88], filtering on 4×4 blocks begins before the entire macroblock is reconstructed. To minimize stalls, the edge filtering order was carefully chosen to account for the 4×4 block processing order (shown in Fig. 17) while adhering to the left-right-top-bottom edge order specified in H.264/AVC [7].

c) *Multiple voltage/frequency domains:* Table 1 shows that MEM has a significantly higher cycle count than the other processing units. The decoder is partitioned into two domains (memory controller and core domain) to enable the frequency and voltage to be independently tailored for each domain. The two domains are completely independent, and are separated by asynchronous FIFOs and voltage level-shifters. Table 1 shows that there could be

further benefit to also placing the ED unit on a separate third domain. The ED is difficult to speed up with parallelism because it uses variable length coding which is inherently serial. For P-frames, the core domain can run at a lower frequency than the memory controller, which saves 25% in power [85].

Having two separate domains allows the frequency and voltage to be dynamically adjusted for the different workloads independently. This is important since workloads for these two domains will vary widely and differently depending on whether the decoder is working on an I-frame or P-frame. For example, the MEM domain requires a higher frequency for P-frames versus I-frames. Conversely, the core domain requires a higher frequency during I-frames since more residuals are present, which are processed by the ED unit.

2) *Circuit Optimization:* Video processing involves movement of a large amount of data. For high definition, each frame is on the order of megabytes which may be too large to place on-chip. Consequently, the frame buffer (FB) used to store the previous frames required for motion compensation is located off-chip. It is important to minimize the off-chip memory bandwidth in order to reduce the overall system power.

On-chip caches (shown in Fig. 18) can be used to reduce both reads and writes such that only the DB unit writes to the frame buffer and only the MC unit reads from it. This caching scheme reduces total off-chip memory bandwidth by almost 26% relative to the case where no caches are used. These caches were implemented using the custom low-voltage 8T SRAMs described in Section III to achieve low voltage operation.

3) *Results and Measurements:* The H.264/AVC video decoder, shown in Fig. 19 was implemented in 65-nm CMOS and the power was measured when performing real-time decoding of several 720p video streams at 30 fps (Table 2). Fig. 20 shows a comparison of this ASIC with other decoders. The power of the I/O pads was not included in the comparison of power measurements. The

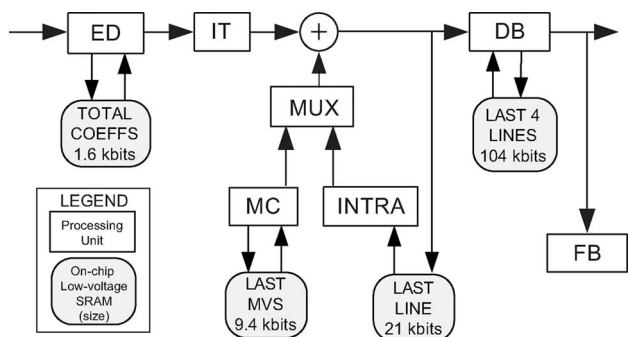


Fig. 18. On-chip caches reduce off-chip memory bandwidth.

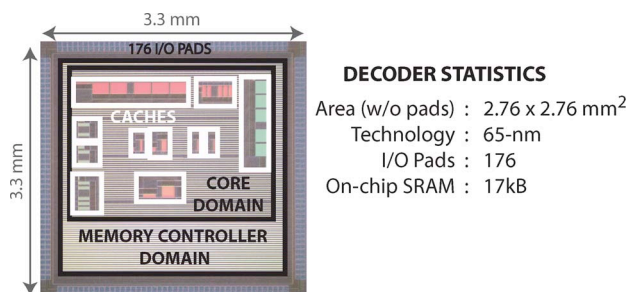


Fig. 19. Die photo of H.264/AVC baseline video decoder.

Table 2 Measured Performance Numbers for 720p at 30 fps

Video	mobcal	shields	parkrun
# of Frames	300	300	144
Bitrate (Mbps)	5.4	7.0	26
Off-chip BW (Gbps)	1.2	1.1	1.2
Core Freq (MHz)	14	14	25
Memory Control Freq (MHz)	50	50	50
Core Vdd [V]	0.7	0.7	0.8
Memory Control Vdd [V]	0.84	0.84	0.84
Power (mW)	1.8	1.8	3.2

decoder can operate down to 0.5 V for QCIF at 15 fps and achieves 1080p at 30 fps at a core voltage of 0.85 V. The power scales accordingly ranging from 29 μ W to 8 mW, a 280 \times range, to cover the 164 \times workload range. The order of magnitude reduction in power over the other reported decoders can be attributed to a combination of using the low-power techniques described in this work and a more advanced process.

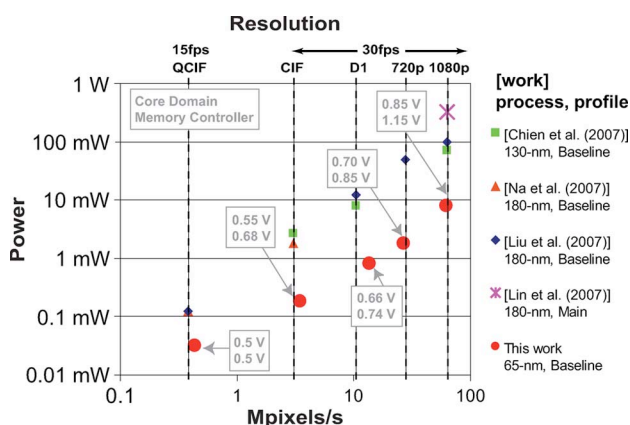


Fig. 20. Comparison with other H.264/AVC decoders: Chien et al. [89], Na et al. [90], Liu et al. [91], Lin et al. [92].

Table 3 Impact of Ultralow-Voltage Techniques on Video Decoder Test-Chip

Performance	QCIF@15fps	720p@30fps
Nominal Decoder Voltage	1.2 V	1.2 V
UDVS Core Decoder Voltage	0.5 V	0.7 V
UDVS Core Frequency	265 kHz	14 MHz
UDVS Total Power	29 μ W	1.8 mW
Power Reduction	5.7 \times	2.9 \times

4) Summary: The video decoder demonstrated three ultralow-voltage techniques described previously in this paper, namely:

- voltage scalable SRAM;
- parallelism and pipelining for increased concurrency;
- multidomain UDVS.

Table 3 illustrates the aggregate impact of these three techniques on the the total power of a video decoder. The power savings are shown for two different resolutions and frames rates requiring different decoder workloads.

B. An Ultralow-Voltage Sensor Processor

Previous sections have discussed techniques to enable ultradynamic voltage scaling in different components of a system. Incorporating many of these approaches, a 65 nm system-on-a-chip functioning down to 300 mV was demonstrated in [12]. The system, pictured in Fig. 21, includes a 16-b microcontroller core and a 128-kb SRAM; both are powered by an integrated dc–dc converter.

Fig. 22 shows a block diagram of the core logic, which is based on the MSP430 microcontroller architecture [14]. The core is synthesized using a custom sub- V_t standard cell

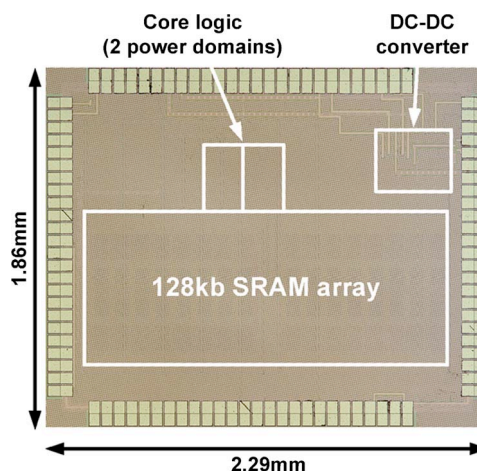


Fig. 21. Die micrograph of microcontroller test-chip.

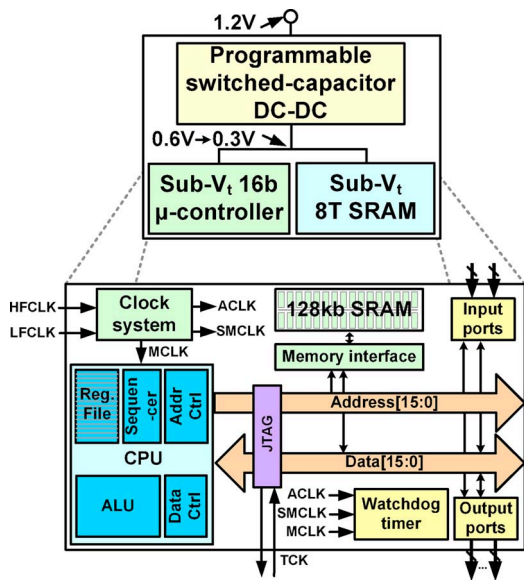


Fig. 22. Block diagram of the system and the core logic.

library designed to mitigate local variation and degraded I_{ON}/I_{OFF} , as described in Section II. In addition, the design employed a variation-aware timing approach based on Monte Carlo simulation to verify hold time constraints. The core logic interfaces to 128-kb of unified instruction and data memory implemented as a custom SRAM.

As detailed in Section III, conventional 6T SRAMs do not function in sub- V_t due to severely degraded read stability and read-current. The SRAM in this system employs an 8T bit-cell to remove the read stability limitation and allow the read-buffer devices to be optimized primarily for read-current. Further, peripheral circuit assists enforce the relative device strengths needed for read and write functionality, despite extreme variation in sub- V_t .

To realize the full energy savings of ultralow-voltage operation, a dc–dc converter supplying these voltages at high efficiencies is essential. Since the power consumption

of the logic and SRAM load circuits drops exponentially at sub- V_t voltages, the dc–dc converter was designed to deliver a maximum of 500 μW of load power. This reduced load power demand makes switched capacitor dc–dc conversion an ideal choice for this application. The switched capacitor (SC) dc–dc converter employs techniques described in Section IV and makes use of 600 pF of total on-chip charge transfer (flying) capacitance to provide scalable load voltages from 300 mV to 1.1 V. The logic and SRAM circuits in this system utilize voltages up to 600 mV.

1) *Measurement Results:* Fig. 23(a) plots the measured energy per cycle versus supply voltage for the microcontroller logic and SRAM at 0 °C, 25 °C, and 75 °C. The energy is measured while the system executes test code which cycles through the available instructions and addressing modes. Memory and logic together consume 27.2 pJ per clock cycle at 500 mV and 25 °C. The optimum energy does not vary much across 20 chips; the measurements have a σ/μ of 0.0897.

Fig. 23(b) plots the microcontroller performance versus supply voltage at 0 °C, 25 °C, and 75 °C. The frequency at 25 °C ranges from 8.7 kHz to 1 MHz across the operating range of 0.3 V to 0.6 V. The σ/μ of measurements across 20 chips at 500 mV is 0.133.

The efficiency of the dc–dc converter delivering 500 mV is shown in Fig. 24. The converter achieves more than 75% efficiency with an order of magnitude change in load power, between 10 μW to 250 μW . With the microcontroller as a load, the converter provides 75% efficiency at 12 μW . When measured standalone, the converter reaches a peak efficiency of 78%.

The inclusion of a dc–dc converter enables the system to dynamically scale V_{DD} to 300 mV during standby mode, where memory and logic together consume less than 1 μW , as shown in Fig. 25. Accounting for the dc–dc converter efficiency loss at such low power levels, this represents a 2.1 \times reduction in leakage power compared to keeping V_{DD} constant at 500 mV during standby.

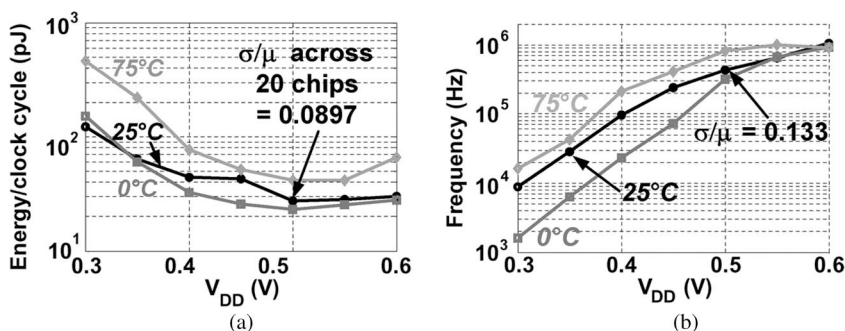


Fig. 23. (a) Energy versus V_{DD} of logic and memory over temperature. (b) Frequency versus V_{DD} across temperature. The σ/μ of measurements across 20 chips at 500 mV is shown.




Fig. 24. DC-DC converter efficiency while delivering 500 mV. The dc-dc converter is powered by a 1.2 V supply. Arrows mark efficiency gain from scalable switch width design as discussed in [29].




Fig. 25. Standby power versus V_{DD} across temperature. The σ/μ of measurements across 20 chips at 300 mV is shown.

2) *Impact of Voltage Scaling:* The microcontroller demonstrated three ultralow-voltage techniques described previously in this paper: low-voltage standard cell design (Section II), voltage scalable SRAM (Section III), and switched capacitor dc-dc converter (Section IV). Although

voltage scaling necessitates certain overhead, such as standard cell upscaling or SRAM peripheral circuit assists, it nevertheless allows the microcontroller to better optimize power at relaxed performance scenarios, for instance when performing basic pulse oximetry, or when executing arrhythmia analysis on an ECG. Further, when the microcontroller must be constantly powered to respond to external events, it is advantageous to operate at the minimum possible V_{DD} to minimize standby power. Table 4 illustrates the net impact of the three techniques on the active and standby power of the sub- V_T microcontroller, relative to a baseline microcontroller with the same functionality but designed for nominal V_{DD} operation.

VII. FUTURE CHALLENGES

The mainstream adoption of the UDVS techniques described in this paper faces some key engineering challenges, which need to be solved by future research. The circuit verification time is expected to greatly increase, as each design must be verified across a wider range of supply voltages and clock frequencies, which are critical for UDVS systems. To make this verification time manageable, new computer-aided design (CAD) tools are needed to handle statistical variation and guarantee chip yields at low voltages. Similarly, new postproduction silicon testing techniques are needed to handle the large increase in the test space with the wider range of voltages and frequencies.

In order to find the energy-optimal level of fine-grained UDVS, some optimization techniques need to be developed that take into account incremental costs such as power grids, voltage regulators, clock trees, PLLs, asynchronous FIFOs, and level-converters. Some high-level workload-prediction algorithms are needed to take advantage of the hardware scalability and correctly set the minimum required frequency for a given throughput constraint. Since the replica ring oscillator used in typical DVS might not match the critical path at low voltages due to transistor variation, some new circuits are needed to

Table 4 Impact of Ultralow-Voltage Techniques in Low-Voltage Microcontroller Test-Chip

	Active Mode		Standby Mode
	Pulse Oximetry 500 kHz	ECG Analysis 1 MHz	
Minimum V_{DD}			
Baseline μ -controller	1.2 V	1.2 V	1.2 V
Low-voltage μ -controller	0.5 V	0.6 V	0.3 V
Power Reduction From Voltage Scaling			
Logic	5.14×	3.57×	83.6×
SRAM	3.35×	2.33×	54.2×

dynamically determine the minimum voltage for a given frequency. For example, one proposed solution is the work of [93], which dynamically tunes processor supply voltages to the minimum allowable by each chip.

VIII. CONCLUSION

This paper presented enabling technologies for ultradynamic voltage scalable systems, where aggressive voltage scaling is employed to optimize system power in response to order-of-magnitude variations in workload. To enable UDVS, system components must address the challenges of low-voltage operation and at the same time, support bursts of high performance demanded by the application.

Specifically, an SRAM functioning between 0.25 V to 1.2 V employs an 8T bit-cell to improve read stability at low voltages while allowing read-buffer devices to be optimized primarily for speed. Further, peripheral circuits to assist reading and writing in sub- V_t can be disabled at nominal V_{DD} to minimize their power and speed overhead. Logic circuits must similarly contend with increased process variation; cell libraries must be designed for functionality under variation, while statistical timing analysis should account for the exponential delay- V_t relationship in sub- V_t .

Highly reconfigurable ADCs are required in UDVS systems, as the necessary signals to digitize can have widely varying resolution requirements and bandwidth. To minimize energy consumption over this wide range requires a multitude of techniques including voltage scaling, duty cycling, reconfigurable architectures, and varying bias currents. A key technique to achieve voltage scalability in ADCs is by employing highly digital circuits and architectures. For instance, a flash ADC operating from 0.2 to 0.9 V employs dynamic digital comparators and

exploits redundancy and reconfigurability to tolerate large comparator offsets and improve linearity.

To realize the full energy savings of UDVS, a dc–dc converter achieving high efficiency across a wide range of load voltages and power levels is essential. In particular, it must employ extremely simple control circuitry to minimize power overhead while supplying microwatts at low voltages. This paper presented an on-chip, switched capacitor converter providing good efficiency from 0.3 V to 1.1 V, and a buck converter with off-chip inductor achieving > 80% efficiency when supplying between 1 μ W to 100 μ W. Further, a simple control loop allows the dc–dc converter to dynamically track the system minimum energy point as it varies with workload and temperature.

The above techniques are demonstrated in two UDVS system applications. An H.264/AVC video decoder operates between 0.5 V to 0.85 V to support a 164 \times range in workload with 280 \times scaling in power. Parallelism and pipelining greatly reduce the number of clock cycles required to decode each frame, enabling throughput requirements to be met at low voltages. A 65 nm, 16-bit microcontroller demonstrates ultralow-voltage operation despite process variation, incorporating logic and SRAM which operate down to 300 mV with an on-chip, switched-capacitor dc–dc converter. The circuit components and techniques we have discussed here represent the first critical steps towards the implementation of practical UDVS systems. ■

Acknowledgment

The authors would like to acknowledge Dimitri Antoniadis, Yu Cao, Eric Wang, and Wei Zhao for help with predictive technology models.

REFERENCES

- [1] B. Davari, R. H. Dennard, and G. G. Shahidi, "CMOS scaling for high performance and low-power-the next ten years," *Proc. IEEE*, vol. 83, no. 4, pp. 595–606, Jan. 1996.
- [2] M. Horowitz, E. Alon, D. Patil, S. Naffziger, R. Kumar, and K. Bernstein, "Scaling, power, and the future of CMOS," in *IEEE Int. Electron Devices Meeting*, Dec. 2005, pp. 7–15.
- [3] K. Itoh, "Adaptive circuits for the 0.5-V nanoscale CMOS era," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2009, pp. 14–20.
- [4] T. D. Burd and R. W. Brodersen, "Design issues for dynamic voltage scaling," in *Proc. IEEE Int. Symp. Low Power Electronics and Design*, Aug. 2000, pp. 9–14.
- [5] V. Gutnik and A. Chandrakasan, "Embedded power supply for low-power DSP," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 5, no. 4, pp. 425–435, Dec. 1997.
- [6] *Technical Note TN2188 Exporting Movies for iPod, Apple TV and iPhone*. [Online]. Available: <http://www.developer.apple.com/technotes/tm2007/tn2188.html>
- [7] "Recommendation ITU-T H.264: Advanced Video Coding for Generic Audiovisual Services," ITU-T, Tech. Rep., 2003.
- [8] *N10176: Call for Test Materials for Future High-Performance Video Coding Standards Development*, ISO/IEC JTC1/SC29/WG11, Oct. 2008.
- [9] *Draft "Request for Video Test Sequences for Enhanced Performance Video Coding Work" (for Approval)*, Jul. 2009, VCEG-L93.
- [10] B. Zhai, L. Nazhandali, J. Olson, A. Reeves, M. Minuth, R. Helfand, S. Pant, D. Blaauw, and T. Austin, "A 2.60 pJ/Inst subthreshold sensor processor for optimal energy efficiency," in *Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2006, pp. 154–155.
- [11] S. Hanson, B. Zhai, M. Seok, B. Cline, K. Zhou, M. Singhal, M. Minuth, J. Olson, L. Nazhandali, T. Austin, D. Sylvester, and D. S. Blaauw, "Performance and variability optimization strategies in a sub-200 mV, 3.5 pJ/inst, 11 nW subthreshold processor," in *Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2007, pp. 152–153.
- [12] J. Kwong, Y. Ramadass, N. Verma, M. Koesler, K. Huber, H. Moormann, and A. Chandrakasan, "A 65 nm sub- V_t microcontroller with integrated SRAM and switched-capacitor dc–dc converter," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2008, pp. 318–319.
- [13] A. C.-W. Wong, D. McDonagh, G. Kathiresan, O. C. Omeni, O. El-Jamaly, T. C.-K. Chan, P. Paddan, and A. J. Burdett, "A 1 V, micropower system-on-chip for vital-sign monitoring in wireless body sensor networks," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2008, pp. 138–139.
- [14] *MSP430x4xx Family User's Guide*. [Online]. Available: <http://www.focus.ti.com/docs/prod/folders/print/msp430fg4619.html>
- [15] V. Chan and S. Underwood, *A Single-Chip Pulsioximeter Design Using the MSP430*. [Online]. Available: <http://www.ti.com/lit/pdf/slaa274>
- [16] [Online]. Available: <http://www.eplimited.com>
- [17] J. L. Willems, C. Abreu-Lima, P. Arnaud, J. H. van Bommel, C. Brohet, R. Degani, B. Denis, J. Gehring, I. Graham, and

- G. van Herpen, "The diagnostic performance of computer programs for the interpretation of electrocardiograms," *New Eng. J. Med.*, vol. 325, no. 25, pp. 1767–1773, Dec. 1991.
- [18] V. Chudacek, M. Hupytch, D. Novak, and L. Lhotska, "Automatic QT interval measurement using rule-based gradient method," in *Comput. Cardiol.*, Sep. 2006, pp. 349–351.
- [19] D. M. Roden, "Drug-induced prolongation of the QT interval," *New Eng. J. Med.*, vol. 350, no. 10, pp. 1013–1022, Mar. 2004.
- [20] P. M. Okin, R. B. Devereux, B. V. Howard, R. R. Fabsitz, E. T. Lee, and T. K. Welty, "Assessment of QT interval and QT dispersion for prediction of all-cause and cardiovascular mortality in American Indians," *Circulation*, vol. 101, pp. 61–66, 2000.
- [21] B. Calhoun, A. Wang, and A. Chandrakasan, "Modeling and sizing for minimum energy operation in subthreshold circuits," *IEEE J. Solid-State Circuits*, vol. 40, no. 9, pp. 1778–1786, Sep. 2005.
- [22] P. Macken, M. Degrauwe, M. V. Paemel, and H. Oguey, "A voltage reduction technique for digital systems," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 1990, pp. 238–239.
- [23] K. J. Kuhn, "Reducing variation in advanced logic technologies: Approaches to process and design for manufacturability of nanoscale CMOS," in *IEEE Int. Electron Devices Meeting (IEDM)*, Dec. 2007, pp. 471–474.
- [24] B. Zhai, S. Hanson, D. Blaauw, and D. Sylvester, "Analysis and mitigation of variability in subthreshold design," in *Proc. IEEE Int. Symp. Low Power Electron. Design*, Aug. 2005, pp. 20–25.
- [25] A. Wang and A. Chandrakasan, "A 180-mV subthreshold FFT processor using a minimum energy design methodology," *IEEE J. Solid-State Circuits*, vol. 40, no. 1, pp. 310–319, Jan. 2005.
- [26] J. Chen, L. T. Clark, and Y. Cao, "Robust design of high fan-in/out subthreshold circuits," in *IEEE Int. Conf. Comput. Design: VLSI Comput. Processors*, Oct. 2005, pp. 405–410.
- [27] J. Chen, L. T. Clark, and T.-H. Che, "An ultra-low-power memory with a subthreshold power supply voltage," *IEEE J. Solid-State Circuits*, vol. 41, no. 10, pp. 2344–2353, Oct. 2006.
- [28] M. J. M. Pelgrom, A. C. J. Duinmaijer, and A. P. G. Welbers, "Matching properties of MOS transistors," *IEEE J. Solid-State Circuits*, vol. 24, no. 5, pp. 1433–1439, Oct. 1989.
- [29] J. Kwong, Y. K. Ramadass, N. Verma, and A. P. Chandrakasan, "A 65 nm Sub- V_t microcontroller with integrated SRAM and switched capacitor dc-dc converter," *IEEE J. Solid-State Circuits*, vol. 44, no. 1, pp. 115–126, Jan. 2009.
- [30] Y. Cao and L. T. Clark, "Mapping statistical process variations toward circuit performance variability: An analytical modeling approach," in *Proc. ACM IEEE Design Automation Conf.*, Jun. 2005, pp. 658–663.
- [31] W. Zhao and Y. Cao, "New generation of predictive technology model for sub-45 nm early design exploration," *IEEE Trans. Electron Devices*, vol. 53, no. 11, pp. 2816–2823, Nov. 2006, verified (IEEE Proc).
- [32] E. Vittoz and J. Fellrath, "CMOS analog integrated circuits based on weak-inversion operation," *IEEE J. Solid-State Circuits*, vol. SC-12, pp. 224–231, Jun. 1977.
- [33] R. M. Swanson and J. D. Meindl, "Ion-implanted complementary MOS transistors in low-voltage circuits," *IEEE J. Solid-State Circuits*, vol. SC-7, no. 2, pp. 146–153, Apr. 1972.
- [34] C. H. Kim, H. Soeleman, and K. Roy, "Ultra-low-power DLMS adaptive filter for hearing aid applications," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 11, no. 6, pp. 1058–1067, Dec. 2003.
- [35] B. Calhoun and A. Chandrakasan, "Ultra-dynamic voltage scaling (UDVS) using sub-threshold operation and local voltage dithering in 90 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2005, pp. 300–301.
- [36] M.-E. Hwang, A. Raychowdhury, K. Kim, and K. Roy, "A 85 mV 40 nW process-tolerant subthreshold 8×8 FIR filter in 130 nm technology," in *Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2007, pp. 154–155.
- [37] V. Sze and A. Chandrakasan, "A 0.4-V UWB baseband processor," in *Proc. IEEE Int. Symp. Low Power Electronics and Design*, Aug. 2007, pp. 262–267.
- [38] A. Raychowdhury, B. Paul, S. Bhunia, and K. Roy, "Computing with subthreshold leakage: Device/circuit/architecture co-design for ultra-low-power subthreshold operation," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 13, no. 11, pp. 1213–1224, Nov. 2005.
- [39] H. Kaul, M. Anders, S. Mathew, S. Hsu, A. Agarwal, R. Krishnamurthy, and S. Borkar, "A 320 mV 56 μ W 411 GOPS/Watt ultra-low voltage motion estimation accelerator in 65 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2008, pp. 316–317.
- [40] E. Seevinck, F. J. List, and J. Lohstroh, "Static-noise margin analysis of MOS SRAM cells," *IEEE J. Solid-State Circuits*, vol. SC-22, no. 5, pp. 748–754, Oct. 1987.
- [41] K. Takeda, Y. Hagihara, Y. Aimoto, M. Nomura, Y. Nakazawa, T. Ishii, and H. Kobatake, "A read-static-noise-margin-free SRAM cell for low-V_{DD} and high-speed applications," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2005, pp. 478–479.
- [42] A. Bhavnagarwala, X. Tang, and J. Meindl, "The impact of intrinsic device fluctuations on CMOS SRAM cell stability," *IEEE J. Solid-State Circuits*, vol. 36, no. 4, pp. 658–665, Apr. 2001.
- [43] H. Qin, Y. Cao, D. Markovic, A. Vladimirescu, and J. Rabaey, "SRAM leakage suppression by minimizing standby supply voltage," in *Proc. IEEE Int. Symp. Quality Electron. Design*, Mar. 2004, pp. 55–60.
- [44] K. Zhang, U. Bhattacharya, Z. Chen, F. Hamzaoglu, D. Murray, N. Vallepalli, Y. Wang, B. Zheng, and M. Bohr, "A SRAM design on 65 nm cmos technology with integrated leakage reduction scheme," in *Proc. IEEE Symp. VLSI Circuits*, Jun. 2004, pp. 294–295.
- [45] J. Wang and B. Calhoun, "Canary replica feedback for near-DRV standby VDD scaling in a 90 nm SRAM," in *Proc. IEEE Custom Integr. Circuits Conf.*, Sep. 2007, pp. 29–32.
- [46] K. Itoh, M. Horiguchi, and T. Kawahara, "Ultra-low voltage nano-scale embedded RAMs," in *Proc. IEEE Int. Symp. Circuits and Syst.*, May 2006, pp. 245–248.
- [47] K. Takeda, Y. Hagihara, Y. Aimoto, M. Nomura, Y. Nakazawa, T. Ishii, and H. Kobatake, "A read-static-noise-margin-free SRAM cell for low-V_{DD} and high-speed applications," *IEEE J. Solid-State Circuits*, vol. 41, no. 1, pp. 113–121, Jan. 2006.
- [48] A. Kawasumi, T. Yabe, Y. Takeyama, O. Hirabayashi, K. Kushida, A. Tohata, T. Sasaki, A. Katayama, G. Fukano, Y. Fujimura, and N. Otsuka, "A single-power-supply 0.7 V 1 GHz 45 nm sram with an asymmetrical unit- β -ratio memory cell," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2008, pp. 382–383.
- [49] L. Chang, D. M. Fried, J. Hergenrother, J. W. Sleight, R. H. Dennard, R. Montoye, L. Sekaric, S. J. McNab, A. W. Topol, C. D. Adams, K. W. Guarini, and W. Haensch, "Stable SRAM cell design for the 32 nm node and beyond," in *Proc. IEEE Symp. VLSI Circuits*, Jun. 2005, pp. 128–129.
- [50] J. P. Kulkarni, K. Kim, and K. Roy, "A 160 mV robust schmitt trigger based subthreshold SRAM," *IEEE J. Solid-State Circuits*, vol. 42, no. 10, pp. 2303–2313, Oct. 2007.
- [51] B. Calhoun and A. Chandrakasan, "A 256 kb sub-threshold SRAM in 65 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2006, pp. 480–481.
- [52] T.-H. Kim, J. Liu, J. Kean, and C. H. Kim, "A 0.2 V, 480 kb subthreshold SRAM with 1 k cells per bitline for ultra-low-voltage computing," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2007, pp. 330–331.
- [53] Z. Liu and V. Kursun, "High read stability and low leakage SRAM cell based on data/bitline decoupling," in *Proc. IEEE Int. Syst. Chip Conf.*, Sep. 2006, pp. 115–116.
- [54] I. J. Chang, J.-J. Kim, S. P. Park, and K. Roy, "A 32 kb 10 T subthreshold SRAM array with bit-interleaving and differential read-scheme in 90 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2008, pp. 388–389.
- [55] N. Verma and A. Chandrakasan, "A 65 nm 8 T sub- V_t SRAM employing sense-amplifier redundancy," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2007, pp. 328–329.
- [56] K. Zhang, U. Bhattacharya, Z. Chen, F. Hamzaoglu, D. Murray, N. Vallepath, Y. Wang, B. Zheng, and M. Bohr, "A 3-GHz 70 Mb SRAM in 65 nm cmos technology with integrated column-based dynamic power supply," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2005, pp. 474–475.
- [57] Z. Liu and V. Kursun, "High read stability and low leakage SRAM cell based on data/bitline decoupling," in *Proc. IEEE Int. Syst.-on-Chip Conf.*, Sep. 2006, pp. 115–116.
- [58] M. E. Sinangil, N. Verma, and A. P. Chandrakasan, "A reconfigurable 65 nm SRAM achieving voltage scalability from 0.25 v to 1.2 v and performance scalability from 20 khz to 200 MHz," in *Proc. IEEE ESSCIRC*, Sep. 2008, pp. 282–285.
- [59] L. Chang, Y. Nakamura, R. K. Montoye, J. Sawada, A. K. Martin, K. Kinoshita, F. H. Gebara, K. B. Agarwal, D. J. Acharyya, W. Haensch, K. Hosokawa, and D. Jamsek, "A 5.3 GHz 8 T-SRAM with operation down to 0.41 v in 65 nm CMOS," in

- Proc. IEEE Symp. VLSI Circuits, Jun. 2007, pp. 252–253.
- [60] A. Wang and A. Chandrakasan, “A 180-mV subthreshold FFT processor using a minimum energy design methodology,” *IEEE J. Solid-State Circuits*, vol. 40, no. 1, pp. 310–319, Jan. 2005.
- [61] Y. Ramadass and A. Chandrakasan, “Minimum energy tracking loop with embedded dc–dc converter delivering voltages down to 250 mV in 65 nm CMOS,” in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2007, pp. 64–587.
- [62] T. Pering, T. Burd, and R. Brodersen, “The simulation and evaluation of dynamic voltage scaling algorithms,” in *Proc. Int. Symp. Low Power Electron. Design*, Aug. 1998, pp. 76–81.
- [63] G.-Y. Wei and M. Horowitz, “A fully digital, energy-efficient, adaptive power-supply regulator,” *IEEE J. Solid-State Circuits*, vol. 34, no. 4, pp. 520–528, Apr. 1999.
- [64] Y. Ramadass and A. Chandrakasan, “Minimum energy tracking loop with embedded dc–dc converter enabling ultra-low-voltage operation down to 250 mV in 65 nm CMOS,” *IEEE J. Solid-State Circuits*, vol. 43, no. 1, pp. 256–265, Jan. 2008.
- [65] A. Stratakos, S. Sanders, and R. Brodersen, “A low-voltage CMOS dc–dc converter for a portable battery-operated system,” in *IEEE Power Electron. Specialists Conf.*, Jun. 1994, vol. 1, pp. 619–626.
- [66] Y. Ramadass and A. Chandrakasan, “Voltage scalable switched capacitor dc–dc converter for ultra-low-power on-chip applications,” in *IEEE Power Electron. Specialists Conf.*, Jun. 2007, pp. 2353–2359.
- [67] D. Maksimovic and S. Dhar, “Switched-capacitor dc–dc converters for low-power on-chip applications,” in *IEEE Power Electron. Specialists Conf.*, Aug. 1999, vol. 1, pp. 54–59.
- [68] G. Patounakis, Y. Li, and K. Shepard, “A fully integrated on-chip dc–dc conversion and power management system,” *IEEE J. Solid-State Circuits*, vol. 39, no. 3, pp. 443–451, Mar. 2004.
- [69] H. Qin, K. Cao, D. Markovic, A. Vladimirescu, and J. Rabaey. (2005, Jan.). Standby supply voltage minimization for deep sub-micron SRAM. *Microelectron. J.* [Online], pp. 789–800. Available: <http://www.gigascale.org/pubs/822.html>
- [70] B. A. Fette and S. D. DaBell, “Dynamically configurable switched capacitor power supply and method,” U.S. Patent 5 414 614, May 1995.
- [71] A. Rao, W. McIntyre, U.-K. Moon, and G. Temes, “Noise-shaping techniques applied to switched-capacitor voltage regulators,” *IEEE J. Solid-State Circuits*, vol. 40, no. 2, pp. 422–429, Feb. 2005.
- [72] LM2797/LM2798 120 mA High Efficiency Step-Down Switched Capacitor Voltage Converter With Voltage Monitoring, National Semiconductor Corp.
- [73] K. Gulati and H.-S. Lee, “A low-power reconfigurable analog-to-digital converter,” *IEEE J. Solid-State Circuits*, vol. 36, no. 12, pp. 1900–1911, Dec. 2001.
- [74] I. Ahmed and D. Johns, “A 50-MS/s (35 mW) to 1-kS/s (15 μ W) power scalable 10-bit pipelined ADC using rapid power-on ramps and minimal bias current variation,” *IEEE J. Solid-State Circuits*, vol. 40, no. 12, pp. 2446–2455, Dec. 2005.
- [75] N. Verma and A. Chandrakasan, “An ultra low energy 12-bit rate-resolution scalable SAR ADC for wireless sensor nodes,” *IEEE J. Solid-State Circuits*, vol. 42, no. 6, pp. 1196–1205, Dec. 2007.
- [76] U. Wismar, D. Wisland, and P. Andreani, “A 0.2 V 0.44 μ W 20 kHz analog to digital $\Sigma\Delta$ modulator with 57 fJ/conversion FoM,” in *Proc. IEEE Eur. Solid State Circuits Conf.*, Sep. 2006, pp. 187–190.
- [77] H.-S. Lee and C. Sodini, “Analog-to-digital converters: Digitizing the analog world,” *Proc. IEEE*, vol. 96, no. 2, pp. 323–334, Feb. 2008.
- [78] J. Sauerbrey, D. Schmitt-Landsiedel, and R. Thewes, “A 0.5-V 1- μ W successive approximation ADC,” *IEEE J. Solid-State Circuits*, vol. 38, no. 7, pp. 1261–1265, Jul. 2003.
- [79] B. Murmann, “A/D converter trends: Power dissipation, scaling and digitally assisted architectures,” in *Proc. IEEE Custom Integr. Circuits Conf.*, Sep. 2008, pp. 105–112.
- [80] D. C. Daly and A. P. Chandrakasan, “A 6-bit, 0.2 V to 0.9 V highly digital flash ADC with comparator redundancy,” in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2008, pp. 554–555.
- [81] M. Flynn, C. Donovan, and L. Sattler, “Digital calibration incorporating redundancy of flash ADCs,” *IEEE Trans. Circuits Syst.—Part II, Analog Digit. Signal Process.*, vol. 50, no. 5, pp. 205–213, 2003.
- [82] D. C. Daly and A. P. Chandrakasan, “A 6-bit, 0.2 V to 0.9 V highly digital flash ADC with comparator redundancy,” *IEEE J. Solid-State Circuits*, to be published.
- [83] S. Gambini and J. Rabaey, “Low-power successive approximation converter with 0.5 V supply in 90 nm CMOS,” *IEEE J. Solid-State Circuits*, vol. 42, no. 11, pp. 2348–2356, Nov. 2007.
- [84] A. Wang and A. Chandrakasan, “A 180 mV FFT processor using subthreshold circuit techniques,” in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, 2004, vol. 1, pp. 292–529.
- [85] V. Sze, D. F. Finchelstein, M. E. Sinangil, and A. P. Chandrakasan, “A 0.7-V 1.8 mW H.264/AVC 720p video decoder,” *IEEE J. Solid-State Circuits*, vol. 44, no. 11, Nov. 2009.
- [86] D. Finchelstein, V. Sze, M. Sinangil, Y. Koken, and A. Chandrakasan, “A low-power 0.7-V H.264 720p video decoder,” in *Proc. IEEE Asian Solid State Circuits Conf.*, Nov. 2008, pp. 173–176.
- [87] E. Fleming, C.-C. Lin, N. D. Arvind, G. Raghavan, and J. Hicks, “H.264 decoder: A case study in multiple design points,” in *Proc. Formal Methods and Models for Co-Design (MEMOCODE)*, Jun. 2008, pp. 165–174.
- [88] K. Xu and C.-S. Choy, “A five-stage pipeline, 204 cycles/mb, single-port SRAM-based deblocking filter for H.264/AVC,” *IEEE Trans. Circuits Syst. Video Technol.*, vol. 18, no. 3, pp. 363–374, Mar. 2008.
- [89] C.-D. Chien, C.-C. Lin, Y.-H. Shih, H.-C. Chen, C.-J. Huang, C.-Y. Yu, C.-L. Chen, C.-H. Cheng, and J.-I. Guo, “A 252 kgate/71 mW multi-standard multi-channel video decoder for high definition video applications,” in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, 2007.
- [90] S. Na, W. Hwangbo, J. Kim, S. Lee, and C.-M. Kyung, “1.8 mW, hybrid-pipelined H.264/AVC decoder for mobile devices,” in *Proc. IEEE Asian Solid State Circuits Conf.*, Nov. 2007.
- [91] T.-M. Liu, T.-A. Lin, S.-Z. Wang, W.-P. Lee, K.-C. Hou, J.-Y. Yang, and C.-Y. Lee, “A 125 μ W, fully scalable MPEG-2 and H.264/AVC video decoder for mobile applications,” *IEEE J. Solid-State Circuits*, vol. 42, no. 1, pp. 161–169, Jan. 2007.
- [92] C.-C. Lin, J.-W. Chen, H.-C. Chang, Y.-C. Yang, Y.-H. O. Yang, M.-C. Tsai, J.-I. Guo, and J.-S. Wang, “A 160 K Gates/4.5 KB SRAM H.264 video decoder for HDTV applications,” *IEEE J. Solid-State Circuits*, vol. 42, no. 1, pp. 170–182, Jan. 2007.
- [93] D. Ernst, S. Das, S. Lee, D. Blaauw, T. Austin, T. Mudge, N. S. Kim, and K. Flautner. (2004, Nov.). Razor: Circuit-level correction of timing errors for low-power operation. *IEEE Micro*. [Online], 24(6), pp. 10–20. Available: <http://www.gigascale.org/pubs/590.html>

ABOUT THE AUTHORS

Anantha P. Chandrakasan (Fellow, IEEE) 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. He is a coauthor of *Low Power Digital CMOS Design* (Kluwer Academic Publishers, 1995), *Digital Integrated Circuits* (Pearson Prentice-Hall, 2003, 2nd edition), and *Sub-Threshold Design for Ultra-Low Power Systems* (Springer, 2006). He is also a coeditor of *Low Power CMOS Design* (IEEE Press, 1998), *Design of High-Performance Microprocessor Circuits* (IEEE Press, 2000), and *Leakage in Nanometer CMOS Technologies* (Springer, 2005). His research interests include low-power digital integrated circuit design, wireless microsensors, ultrawideband radios, and emerging technologies.

Prof. Chandrakasan was a corecipient of 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, the 1999 DAC Design Contest Award, the 2004 DAC/ISSCC Student Design Contest Award, the 2007 ISSCC Beatrice Winner Award for Editorial Excellence and the 2007 ISSCC Jack Kilby Award for Outstanding Student Paper. He has served as a technical program cochair 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 Subcommittee Chair for ISSCC 1999-2001, the Program Vice-Chair for ISSCC 2002, the Program Chair for ISSCC 2003, and the Technology Directions Subcommittee Chair for ISSCC 2004-2009. He was an Associate Editor for the IEEE JOURNAL OF SOLID-STATE CIRCUITS from 1998 to 2001. He served on SSCS AdCom from 2000 to 2007 and he was the meetings committee chair from 2004 to 2007. He is the Conference Chair for ISSCC 2010. He is the Director of the MIT Microsystems Technology Laboratories.



Daniel Frederic Finchelstein (Member, IEEE) received the B.A.Sc. degree from the University of Waterloo, Canada, in 2003 and the M.S. and Ph.D. degrees from the Massachusetts Institute of Technology, Cambridge, in 2005 and 2009, respectively. His doctoral thesis focused on energy and memory efficient parallel video processing.

He has been an engineering intern on eight different occasions at companies like IBM, ATI, Sun, and Nvidia. He is currently working in the 3D graphics performance group at Nvidia Corporation, Santa Clara, CA. He holds one U.S. patent related to cryptographic hardware. His research interests include energy-efficient and high-performance digital circuits and systems. His recent focus has been on parallel processors and memory architectures for video and graphics processing.

Dr. Finchelstein received student design contest awards at both A-SSCC (2008) and DAC/ISSCC (2006), and has several conference and journal publications. He was awarded the Natural Sciences and Engineering Research Council of Canada (NSERC) graduate scholarship from 2003 to 2008. He also received the University of Waterloo Sanford Fleming Medal for ranking first in the graduating class.



Joyce Kwong (Student Member, IEEE) received the B.A.Sc. degree at the University of Waterloo, Canada, in 2004, where she was awarded the Governor General's Silver Medal, and the M.S. degree in electrical engineering from the Massachusetts Institute of Technology (MIT), Cambridge, in 2006. She is currently working towards the Ph.D. degree at MIT.

Ms. Kwong was the recipient of the Texas Instruments Graduate Woman's Fellowship for Leadership in Microelectronics in 2007 and the NSERC Postgraduate Fellowship. She was a corecipient of the 2008 ISSCC Jack Kilby Award for Outstanding Student Paper. Her research interests include low-voltage design methodology and system implementation.



Denis C. Daly (Member, IEEE) received the B.A.Sc. degree in engineering science from the University of Toronto, Canada, in 2003, and the S.M. and Ph.D. degrees from the Massachusetts Institute of Technology, Cambridge, in 2005 and 2009, respectively.

From May 2005 to August 2005 he worked in the Wireless Analog Technology Center at Texas Instruments, Dallas, TX, designing transceiver modules. From May 2003 to August 2003, he worked on high-speed signaling systems at Intel Laboratories, Hillsboro, OR. His research interests include low-power wireless transceivers, ultralow-power systems, and highly digital RF and analog circuits.

Dr. Daly received Natural Sciences and Engineering Research Council of Canada (NSERC) Postgraduate Scholarships in 2003 and 2005. He was awarded a Student Paper Prize at the 2006 RFIC Symposium and won third place in the 2006 DAC/ISSCC Student Design Contest (Operational System Category).



Yogesh Kumar Ramadass (Member, IEEE) received the B.Tech. degree in electronics and electrical communication engineering from the Indian Institute of Technology, Kharagpur, in 2004 and the S.M. and Ph.D. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, in 2006 and 2009, respectively.

From May 2007 to August 2007, he worked in the Wireless Analog Technology Center at Texas Instruments, Dallas, TX, designing power converters. His research interests include low-power circuit design, dc-dc converters and energy harvesting/processing circuits.

Dr. Ramadass was awarded the President of India Gold Medal in 2004, was a corecipient of the 2008 ISSCC Jack Kilby Award for Outstanding Student Paper, the 2007 ISSCC Beatrice Winner Award for Editorial Excellence, and the 2007 ISLPED Low Power Design Contest Award. He was a recipient of the 2008-2009 Intel Foundation Ph.D. Fellowship.




Mahmut Ersin Sinangil (Student Member, IEEE) received the B.Sc. degree in electrical and electronics engineering from Bogazici University, Istanbul, Turkey in 2006 and the S.M. degree in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, MA, in 2008. He is currently working toward the Ph.D. degree at MIT, where his research interests include low-power digital circuit design in the areas of SRAMs and video coding.



Mr. Sinangil received the Ernst A. Guillemin Thesis Award at MIT for his Master's thesis in 2008 and the Bogazici University Faculty of Engineering Special Award in 2006.

Vivienne Sze (Student Member, IEEE) received the B.A.Sc. (Hons.) degree in electrical engineering from the University of Toronto, Canada, in 2004, and the S.M. degree from the Massachusetts Institute of Technology (MIT), Cambridge, MA, in 2006, where she is currently working toward the Ph.D. degree.



From May 2007 to August 2007, she worked in the DSP Solutions R&D Center at Texas Instruments, Dallas, TX, designing low-power video coding algorithms. From May 2002 to August 2003, she worked at Snowbush Microelectronics, Toronto, ON, Canada, as an IC Design Engineer. Her research interests include low-power circuit and system design, and low-power algorithms for video compression. Her work has resulted in several contributions to the ITU-T Video Coding Experts Group (VCEG) for the next-generation video coding standard.

Ms. Sze was a recipient of the 2007 DAC/ISSCC Student Design Contest Award and a corecipient of the 2008 A-SSCC Outstanding Design Award. She received the Julie Payette—Natural Sciences and Engineering Research Council of Canada (NSERC) Research Scholarship in 2004, the NSERC Postgraduate Scholarship in 2007, and the Texas Instruments Graduate Woman's Fellowship for Leadership in Microelectronics in 2008. She was awarded the University of Toronto Adel S. Sedra Gold Medal and W.S. Wilson Medal in 2004.

Naveen Verma (Member, IEEE) received the B.A.Sc. degree in electrical and computer engineering from the University of British Columbia, Canada, in 2003 and the M.S. and Ph.D. degrees in electrical engineering from Massachusetts Institute of Technology, Cambridge, in 2005 and 2009, respectively.



Since July 2009 he has been an Assistant Professor of Electrical Engineering at Princeton University, Princeton, NJ. He has coauthored book-chapters in *Embedded Memory for Nano-Scale VLSI* (Springer, 2009) and *Adaptive Techniques for Processor Optimization* (Springer, 2008). His research focuses on ultralow-power integrated circuits including low-voltage digital logic and SRAMs, low-noise analog instrumentation and data-conversion, and energy-efficient processing algorithms especially for biomedical applications.

Prof. Verma is corecipient of 2008 ISSCC Jack Kilby Award for Outstanding Student Paper, and 2006 DAC/ISSCC Student Design Contest Award. During his doctoral research he was an Intel Foundation Ph.D. Fellow and an NSERC Fellow.