Highly Interleaved 5-bit, 250-MSample/s, 1.2-mW ADC With Redundant Channels in 65-nm CMOS

Brian P. Ginsburg, *Member, IEEE*, and Anantha P. Chandrakasan, *Fellow, IEEE*

***Abstract*—This successive approximation register ADC uses time-interleaving to gain the energy advantage of slower circuits (reduced supply voltage and improved bias points) without sacrificing high speed operation. The drawbacks of interleaving are addressed through architectural solutions. Channel redundancy counteracts the severe yield loss that parallel circuits experience due to local variation. Clock partitioning restricts the distribution of the precise, high-speed sampling clock to three centrally located sampling networks. Only a low frequency clock is distributed across the majority of die area. The skew-resistant global top-plate sampling network is extended to allow overlapped sampling windows without introducing extra sources of crosstalk. The 36-way interleaved 5-bit ADC operates with a core voltage of 800 mV and consumes 1.20 mW total power at 250 MS/s. At Nyquist, the SNDR is 28.4 dB. The 6 redundant channels (17% overhead) increase the yield of the 24 measured chips from 42% to 88%.**

***Index Terms*—ADC, deep submicron CMOS, SAR, time-interleaving.**

I. Introduction

THE proliferation of portable and battery-powered devices places increasingly stringent requirements on power dissipation. In many wireless or communication systems high-speed analog-to-digital converters (ADCs) are a critical block for overall performance and power. Advanced CMOS processses, circuit innovations, and the increasing use of the successive approximation register (SAR) topology have led to dramatic improvements in ADC energy efficiency [1]–[6]. Here, we focus on the use of parallelism, which is compatible with many of these other advances, to improve energy efficiency of a mixed-signal system.

Time-interleaved [7] ADCs have multiple slower sub-ADCs, or channels, that sample at fixed offsets in time to achieve an overall higher sampling rate. This technique is traditionally used to push the speed/resolution boundary of an architecture. Eighty interleaved pipelined ADCs have enabled 20 GS/s, 8-bit operation [8]. Interleaved SAR converters achieve similar sampling rates to flash ADCs [1]–[3], [9], and, very recently, are even used to sampled at tens of gigahertz [10]. In this paper, however, we focus on a different motivation for time-interleaving: slower ADCs are fundamentally more energy efficient than their very high speed counterparts. Time-interleaving extends the conversion time of individual channels, which permits lower supply voltages and subthreshold biasing of analog circuits. These in turn improve the energy efficiency of the channels and the overall ADC if overhead (e.g., clocking) is limited.

Unfortunately, time-interleaving introduces problems not present in single-channel ADCs. Mismatches between channels' offsets, gains, and sampling instants produce distortion that reduce the overall signal-to-noise-plus-distortion ratio (SNDR) of the converter, as analyzed in [11]. These mismatches are due to variations between channels and can be considered local in nature. In fact, any local variations are detrimental to parallel systems, where their impact increases with the number of parallel elements. Most highly-interleaved converters use extensive calibration to correct for variations between channels [8], [10]. Several techniques have been developed to perform background calibration for offset, gain, and/or timing mismatches [12]–[14]. Calibration is effective but requires additional correction circuitry per-channel that must be carefully evaluated for its energy cost, particularly when, as in this work, interleaving is used to lower power consumption.

Here, we introduce the use of redundant channels in a time-interleaved converter to mitigate the effects of local variation. Redundancy becomes increasingly effective as the level of parallelism in a system increases, and, in this implementation, can be used to correct parametric errors and even defects within a channel. Sizing, calibration, and now redundancy, are the essential elements to improving yield in mixed-signal circuits.

This paper presents a 36-channel, 5-bit, 250-MS/s ADC [15]. Each channel is implemented using the SAR topology. The large number of channels permits operation at 800 mV with the analog circuits biased in the subthreshold region of operation. Architectural solutions are proposed to address the major limitations of interleaving. As mentioned above, redundancy is applied to the channels to address mismatch and local variation. A global passive top-plate sampling network reduces timing skew and has been extended to permit longer and overlapped sampling windows with minimal crosstalk. Finally, a block structure and hierarchical clock network simplify clock generation and distribution requirements.

II. Parallelism in ADCs

While interleaving has been used to lower ADC power consumption by enabling the choice of more energy efficient architectures [1], [9], this work further increases the number of channels to improve the energy efficiency without a change in the channel architecture. We will begin with a discussion of the relationship between energy and parallelism in ADCs for a fixed

Manuscript received April 04, 2008; revised July 16, 2008. Current version published December 10, 2008. The work was supported by the Defense Research Advanced Projects Agency (DARPA). The work of Brian P. Ginsburg was supported by the NDSEG Fellowship.

B. P. Ginsberg is with Texas Instruments, Dallas, TX 75266 USA (e-mail: bginzz@ti.edu).

A. P. Chandrakasan is with the Massachusetts Institute of Technology, Cambridge, MA 02139 USA.

Digital Object Identifier 10.1109/JSSC.2008.2006334

0018-9200/$25.00 © 2008 IEEE

Authorized licensed use limited to: MIT Libraries. Downloaded on January 17, 2009 at 11:26 from IEEE Xplore. Restrictions apply.

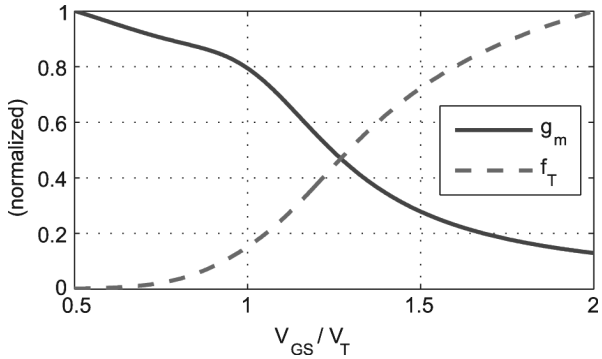

Fig. 1. Simulated nMOS g_m and f_T as a function of gate voltage.

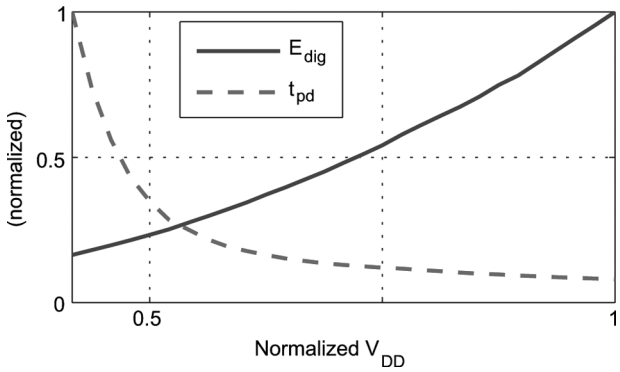

Fig. 2. Simulated inverter energy and delay.

sampling rate and then discuss the deleterious effects of local variation in highly parallel systems.

A. Mixed-Signal Optimum Energy Point

Both the analog and digital energy efficiencies of an ADC improve as the speed of operation decreases. For many linear analog circuits (such as the preamplifier shown in Fig. 11) the energy scales as

$$E_{\text{analog}} \propto \frac{V_{\text{DD}} C_L I_D}{g_m}, \tag{1}$$

where V_{DD} is the supply voltage, C_L is the total load capacitance, I_D is the quiescent bias current, and g_m is the transistor transconductance. In (1), energy consumption is inversely proportional to the transconductance efficiency, g_m/I_D, which is plotted in Fig. 1. Transconductance efficiency is maximized in the subthreshold region of operation where $V_{GS} < V_T$. Unfortunately, as shown, device f_T is severely degraded in this region, precluding high speed operation. For converters operating at hundreds of MS/s, the analog circuits typically must be biased in strong inversion with consequent reduction in energy efficiency.

Similarly, at slower sampling speeds, digital propagation delays can increase. Therefore, the supply voltage can be reduced to achieve V_{DD}^2 energy savings (Fig. 2). This is a driving motivation for the use of parallelism and/or pipelining in low power digital circuits [16]. While not usually discussed in ADC literature, digital power is a significant problem, particularly in low-to-medium resolution converters [17].

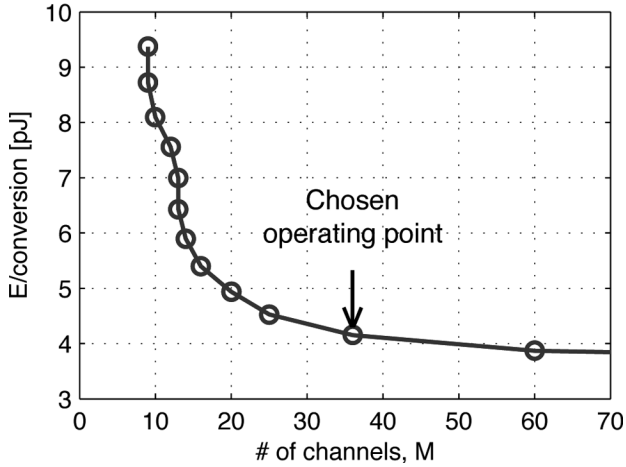

Fig. 3. Optimization results showing energy per conversion as the number of interleaved channels, M, increases.

As the number of channels increases, overheads due to synchronization, such as clocking and muxing of the ADC outputs, begin to counteract the benefits of slower channel operation. A mixed-signal energy optimization has been performed on an interleaved SAR ADC, as described in [17], and the energy per conversion is plotted versus the number of channels M in Fig. 3. The largest source of energy savings towards the right of the plot is the digital supply voltage scaling. After 25–30 channels, the curve exhibits a broad minimum. (The actual minimum value of 121 channels is not shown on the graph for the sake of clarity.) In this region, further savings from slower channel operation are roughly balanced by overhead of interleaving. Because the minimum is so broad, other design factors (e.g., area and design complexity) have led to the choice of 36 channels for this design, with a predicted energy consumption within 10% of optimal.

Along with the top-level energy results of Fig. 3, the optimization gives initial values for most analog design parameters (e.g., preamplifier transistor widths, bias currents, and sampling switch sizes) and buffering and drive strengths at the output of digital blocks [17]. These values have been used to shorten the design cycle time, similar to other ADC optimizations [18].

B. Impact of Local Variation

While offset, gain, and timing mismatches between channels are the most discussed impairments of interleaved ADCs [11]–[14], they represent only special cases of channel variation. In fact, all local variations produce a greater impact on yield in interleaved versus single-channel ADCs. The performance degradation may not be apparent by looking at statistical quantities, such as SNDR or integral nonlinearity (INL), calculated from the interleaved output of an ADC because these can average out uncorrelated errors between channels. In most practical systems, every sample should meet a similar level of precision. For this reason, a chip will be considered functional only if every individual channel, as well as the overall interleaved ADC meet target performance metrics (e.g., INL, SNDR, etc.).

Assuming the yield of an individual channel Y_C can be independently attributed as Y_G or Y_L from global or local variations, respectively, $(Y_C = Y_G Y_L)$ the yield of an M-way interleaved

Authorized licensed use limited to: MIT Libraries. Downloaded on January 17, 2009 at 11:26 from IEEE Xplore. Restrictions apply.

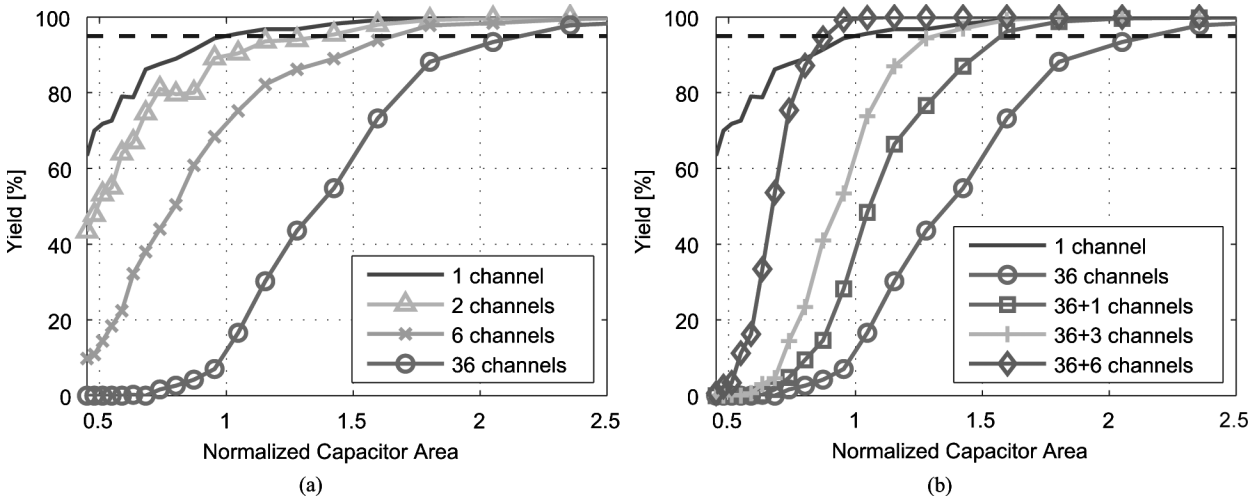

Fig. 4. Behavioral simulations showing the INL yield versus the unit capacitor size within each channel. (a) Number of parallel channels is varied. (b) Number of redundant channels is varied with 36 nominal channels.

ADC is $Y = Y_G Y_L^M$. Thus, even if global variations may have greater impact in a non-interleaved converter ($Y_G < Y_L$), as the number of channels increases local variations severely impact yield. For this reason, the remaining analysis assumes that total yield is due to local variations only, $Y_C = Y_L$. A breakdown of the measured variation mechanisms is presented in Section V-B.

From a design perspective, this effect can be clearly seen by looking at the static linearity of a channel. In a charge redistribution SAR ADC [19], INL is ultimately limited by capacitor array matching. As matching improves with larger capacitor sizes, the INL yield of a single SAR ADC improves with a bigger array, as shown in the top curve of Fig. 4(a). With good layout techniques, matching is inherently a local effect. Therefore, each channel's INL varies independently, and yield drops with increased M. As shown, for 36 channels and the same 95% yield, the unit capacitor size must increase to 2.2 × the original value. This is not only a large overall area penalty because capacitors are the plurality of channel area, but it also is a large reference voltage and digital switching energy penalty, which was not accounted for in the energy optimization. By using sizing to recover yield, much of the energy savings from interleaving would be lost.

As an alternative approach, calibration can be used to correct for static nonlinearities and offsets in SAR ADCs [20] at the cost of additional channel complexity. While calibration is often an efficient mechanism to reduce the required raw (uncalibrated) yield, it does not change the requirement of post-calibrated yield. It is therefore more difficult to calibrate an interleaved ADC with required channel yield $Y_C = Y^{1/M}$ than to calibrate a non-interleaved ADC. In addition, fault coverage is inherently limited by the particular calibration implementation. For instance, the technique in [20] can correct for INL errors but not for dynamic nonlinearities.

Rather than trying to use sizing or calibration to improve channel performance, the proposed ADC uses redundant channels to correct for yield loss due to local variation. Memories routinely use redundancy to correct failures in large arrays [21]. ADCs have also used redundancy to relax comparator requirements in a flash converter [22]. The proposed implementation, however, uses a much larger redundant element (an entire channel) but with fewer total redundant elements. By adding extra channels and choosing the "best" ones, the statistical outliers with the worst performance can be eliminated.

TABLE I
UNIT CAPACITOR SIZES FOR 95% OVERALL YIELD

# nominal channels	# redundant channels	Unit capacitor size
1	0	1×
36	0	2.2×
36	6	0.9×

The advantage of redundancy can be seen by revisiting the capacitor array sizing simulations discussed earlier. As shown in Fig. 4(b), every redundant channel beyond the 36 nominal ones shifts the yield curve to the left, and using six redundant channels recovers the yield lost due to parallelism. The required unit capacitor sizes for 95% yield are summarized in Table I. For the same 36 nominal channels, only a 17% area overhead in number of channels permits more than a 2 × reduction in capacitor size per channel. This redundancy implementation therefore only requires a minimal increase in total chip area but with significantly smaller devices for lower power consumption.

The overall yield of the ADC due to identical independently distributed local variations among the channels is a Bernoulli random process. If the probability that a single channel works is Y_C, then the probability that at least M out of $M_{\text{total}} = M + M_{\text{red}}$ channels meet the specification is

$$Y_{\text{red}} = \sum_{j=M}^{M+M_{\text{red}}} \binom{M_{\text{total}}}{j} Y_C^j (1 - Y_C)^{M_{\text{total}} - j}. \quad (2)$$

For a target yield of 95% with $M = 36$, increasing the number of redundant channels M_{red} from 0 to 6 reduces the required channel yield from the stringent specification of 99.86% to the much more relaxed 91.9%.

Channel redundancy is not limited to improving static nonlinearities. It can correct for any parameters that vary between

Authorized licensed use limited to: MIT Libraries. Downloaded on January 17, 2009 at 11:26 from IEEE Xplore. Restrictions apply.

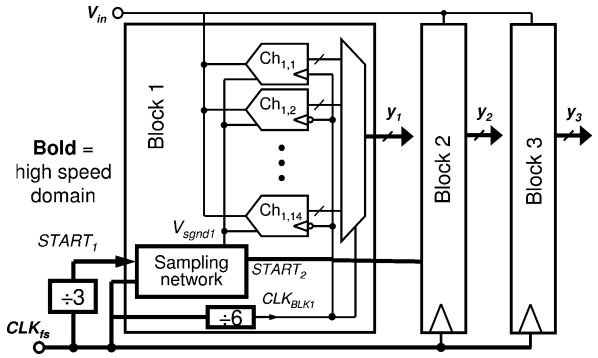

Fig. 5. Top-level ADC block diagram.

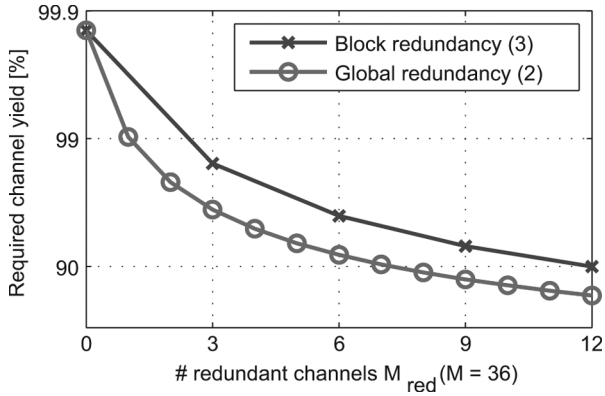

Fig. 6. Comparison of the implemented redundant channel structure to one where all channels are interchangeable.

channels, including dynamic nonlinearities, gain and offset mismatch, timing skew, and even localized defects. The redundancy implementation details and top-level ADC architecture are presented in the next section. The channel selection procedure, which controls for what faults redundancy corrects, will be discussed in Section V-A.

Calibration and redundancy are fully complementary techniques and have different design strengths. Redundancy becomes increasingly effective at higher levels of parallelism, where the same number of redundant elements give a greater improvement in yield. At higher resolutions, the overhead from complex calibration schemes is reduced in comparison to overall converter power. As this ADC is low resolution but highly parallel, the relatively simple digital offset correction is the only calibration technique used (see Section IV-B), whereas redundancy is the primary mechanism to improve most parameters.

III. Highly-Interleaved ADC Architecture

The top-level architecture of the ADC is shown in Fig. 5. Three blocks each contain 12 nominal and two redundant channels. The block also has a sampling network, clock generation block, and output mux. All the channels within a block share the same $\mathrm{CLK_{BLK}}$ with frequency $f_S/6$, where f_S is the overall sampling frequency. Half of the channels align to the positive edge of this clock, and half align to the negative edge. All critical sampling edges are confined to the block sampling network, described later in Section IV-A, easing the requirements of the jitter and duty cycle distortion of $\mathrm{CLK_{BLK}}$. The channels within a block, as well as the blocks themselves, synchronize by passing a $START$ token that, when received, initiates sampling [9]. The path of this $START$ token fixes the channel order, but any individual channel can be skipped, as described in Section IV-B.

The block sampling network receives the precise, high-speed sampling clock CLK_{fS}. This clock is only distributed to the three block sampling networks, which are placed near each other in the layout. The block clock $\mathrm{CLK_{BLK}}$ is distributed across a much larger area, but it switches at a lower frequency. In addition, the ADC performance is not sensitive to even moderate amounts (up to 1 ns) of skew on $\mathrm{CLK_{BLK}}$, so it is distributed in an imbalanced network (different path lengths to the various channels) to minimize total routing capacitance and further reduce power dissipation. This clock partitioning greatly reduces clock generation requirements (only one $f_S/6$ clock is derived per-block) and distribution power.

With this implementation, however, the assumptions underlying the redundancy analysis in Section II-B, are no longer met. The derivation of (2) assumes that the best M channels can always be chosen, but here channels are grouped within blocks. Only the best 12 out of 14 channels within a block can be chosen. The total yield of an ADC with such a block structure is

$$Y_{\mathrm{red}} = \left[\sum_{j=M/N_{\mathrm{block}}}^{M_{bl+\mathrm{red}}} \binom{M_{bl+\mathrm{red}}}{j} Y_C^j (1-Y_C)^{M_{bl+\mathrm{red}}-j} \right]^{N_{\mathrm{block}}}. \quad (3)$$

Here, N_{block} is the number of blocks, and $M_{bl+\mathrm{red}} = (M + M_{\mathrm{red}})/N_{\mathrm{block}}$ is the total number of channels (nominal and redundant) within a block. Fig. 6 compares (2) with (3) (with $N_{\mathrm{block}} = 3$) as the number of total redundant channels is increased. As shown, the final implementation with blocks of 12 nominal and 2 redundant channels does increase the required channel yield to 96% (compared to 91.9% for global redundancy); however, the resultant decrease in complexity makes this a good tradeoff. To fully implement global redundancy, either the full frequency sampling clock needs to be distributed to every channel, increasing the clocking power significantly, or extensive synchronization circuitry must be added to permit swapping channels between blocks.

IV. Circuit Details

A. Sampling Network

Sampling remains a difficult problem in interleaved converters, both due to the high speed analog input signal and the distributed sampling network across a chip. Timing skew, which is any deviation from the ideal channel sampling instants, results in a degradation of performance as input signal frequency increases. Assuming the timing skew of each channel is a Gaussian random variable with standard deviation σ_{skew}, an M-way interleaved ADC sampling a sinusoidal signal of frequency f_{in} has a skew-limited SNDR of [23]

$$\mathrm{SNDR} = 20\log\left(\frac{1}{\sigma_{\mathrm{skew}} 2\pi f_{\mathrm{in}}}\right) - 10\log\left(\frac{M-1}{M}\right). \quad (4)$$

Authorized licensed use limited to: MIT Libraries. Downloaded on January 17, 2009 at 11:26 from IEEE Xplore. Restrictions apply.

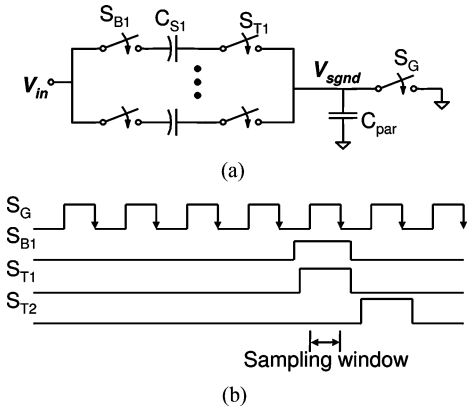

Fig. 7. (a) Schematic and (b) timing diagram of global passive top-plate sampling network.

For instance, a 250 MHz input requires $\sigma_{\rm skew} < 16\ {\rm ps}_{\rm rms}$ for 5-bit resolution. This level of timing precision is very difficult to achieve via clock distribution across a large die area. Phase mismatch due to bandwidth limitations of the sampling network produces a similar effect to timing skew.

One approach that completely eliminates timing skew and phase mismatch is to precede the time-interleaved ADC with a single high-speed active sample and hold, sampling at the maximum sampling rate [24], [25]. Each channel only sees a held signal, and timing skew does not produce any voltage error if sufficient settling time is used. The active S/H, however, would lead to a dramatic increase in this ADC's power consumption [26].

Gustavsson and Tan proposed a global passive sampling technique that can reduce the timing skew related distortion by 10–20 dB [27], [28], extending a technique introduced for double sampling ADCs [29]. They use a global top-plate sampling network, as shown in Fig. 7(a). The top plate switch of a conventional sampling network is split into two series switches, a global switch S_G shared among all the channels, and a per-channel switch S_{Ti}. The timing diagram of the switches is shown in Fig. 7(b). The sampling period for channel i occurs when both S_G and S_{Ti} are closed, with only one channel sampling at a time. Switch S_G opens first, defining the sampling instant; as this switch is shared among all the channels, skew is ideally eliminated.

The presence of the parasitic capacitor $C_{\rm par}$, however, leads to some residual voltage error due to timing skew between the per-channel switches S_{Ti}. Defining the capacitor divider ratio

$$a = \frac{C_{\rm par}}{C_{\rm par} + C_{Si}} \tag{5}$$

the timing-skew related SNDR is [27]

$$\text{SNDR} = 20\log\left(\frac{1}{\sigma_{\rm skew} 2\pi f_{\rm in}}\right) - 10\log\left(\frac{M-1}{M}\right) - 20\log(a). \tag{6}$$

Compared to (4), the SNDR improves by $20\log(a)$.

To avoid this sensitivity to parasitic capacitance, the global switch can also be placed on the bottom plate in series with $V_{\rm in}$ [30]. As the input signal is disconnected from the channel after S_G opens, skew is completely eliminated. Bottom plate switches, however, are floating and therefore tend to be larger than top-plate ones that are referenced to ground. By placing two switches in series, the total switch area must increase by $4\times$ to achieve the same $R_{\rm on}$. To minimize the total increase in switching power, we placed S_G on the top plate. In addition, this reduces signal dependent charge injection when S_G opens.

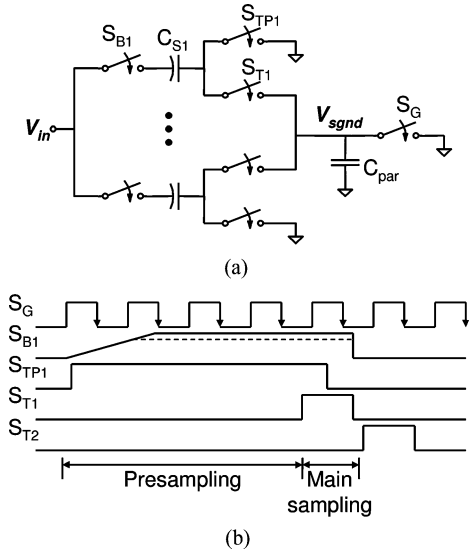

Fig. 8. Sampling network extended for overlapped sampling windows. (a) Schematic and (b) timing diagram.

By using one sampling network per-block (i.e., 3 S_G's), wiring and diffusion capacitances of only 14 channels, instead of all 42, contribute to $C_{\rm par}$, and a is therefore small enough to keep the residual skew low enough for the target resolution. At resolutions over 8 bits, however, it will be practically impossible to maintain a at a sufficiently low level, necessitating the use of the bottom plate global switch [30] or other techniques.

The total sampling window in the network of Fig. 7(a) is limited to one period of the full speed sampling clock, which is much shorter than the period of the channel clock $\rm CLK_{BLK}$. The small sampling window can lead to incomplete settling of the step response that occurs at the onset of sampling and causes the periodic behavior of SNDR with respect to the channel sampling frequency observed in other interleaved converters [9], [10]. We have extended the global top-plate sampling network to allow for longer sampling windows that overlap between channels but in such a way as to minimize crosstalk. As multiple channels sample the input simultaneously, any charge injected onto shared nodes by opening or closing the sampling switches of one channel must either produce an error voltage much smaller than 1 LSB or must be timed such that the transient of the entire sampling network will die out before the next critical sampling instant. To reduce reliance on precise timing, our sampling network minimizes the number of shared nodes and keeps the maximum level of voltage transients at an acceptable level. Overlapped channel sampling also increases the total load seen by the input driver. The total capacitance presented by all channels sampling concurrently is kept to a level that can be driven with a 50 Ω with sufficient input bandwidth.

The extended sampling network is shown in Fig. 8. A separate top-plate switch S_{TPi} is closed during a pre-sampling period to directly ground the top plate of C_{Si}; only one channel

Authorized licensed use limited to: MIT Libraries. Downloaded on January 17, 2009 at 11:26 from IEEE Xplore. Restrictions apply.

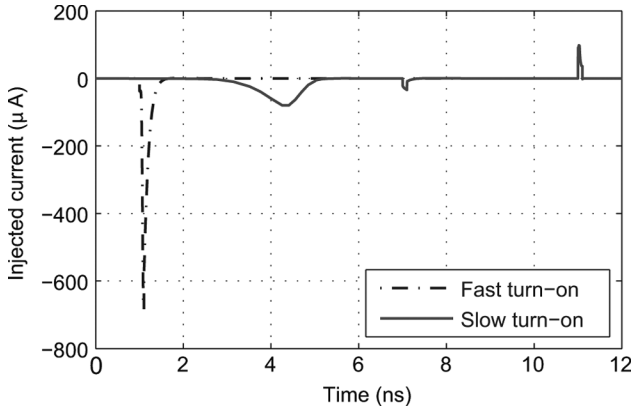

Fig. 9. Simulation comparing the current injected onto $V_{\rm in}$ when S_B is driven with step and ramp drives.

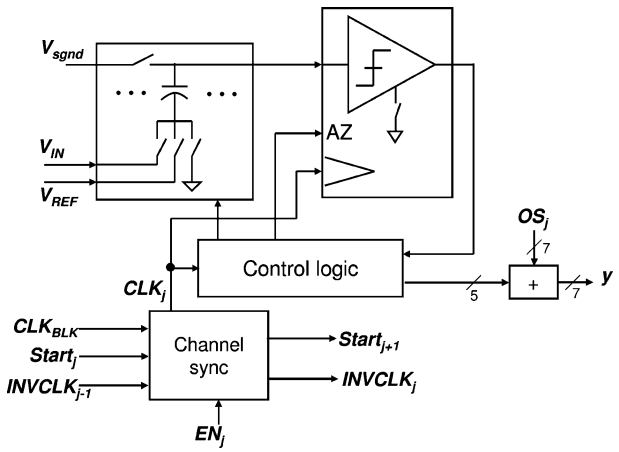

Fig. 10. Channel block diagram. $\rm OS_{\it j}$ is the digital offset correction data.

is connected to $V_{\rm sgnd}$ at a time, avoiding coupling on this node. The bottom plate switch S_{Bi} is driven with a slow ramp drive to a boosted supply voltage. As compared to a step drive, simulations show that the slow turn on reduces the peak current injected onto the input node by more than 8.5 × (Fig. 9). The kickback noise is sufficiently low that no measures are taken to further align one channel's start of sampling away from critical sampling edges of other channels. The bottom plate switch is also driven to a higher supply voltage of 1.2 V, which is within the process limits but above the core operating voltage of the rest of the ADC, which aids in tracking the high frequency analog input signal. As shown in [17], the additional supply permits more aggressive scaling of the core supply for overall energy savings.

B. Channel Circuits

In addition to the standard SAR components of a capacitive digital-to-analog converter (DAC), comparator, and control logic, the channel (Fig. 10) also includes synchronization and digital offset correction blocks. The DAC is implemented using the split capacitor array [3] for its good energy and linearity performance. To reduce sensitivity to systematic mismatches, the layouts of the sub-arrays are swapped between the two common centroid differential capacitor arrays [26].

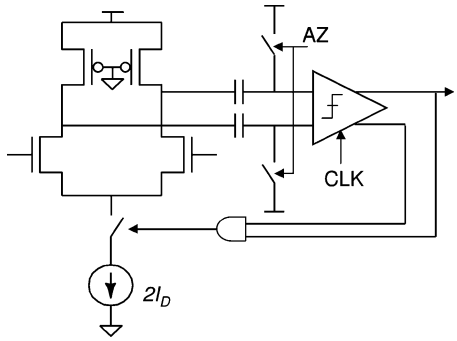

Fig. 11. Comparator with duty-cycled preamplifier and regenerative latch.

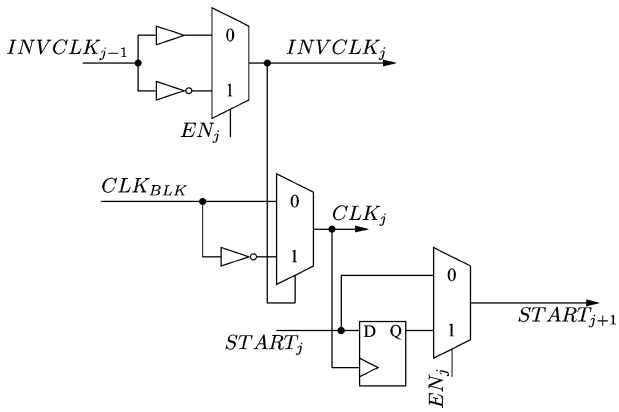

Fig. 12. Synchronization block of channel i.

The comparator (Fig. 11) consists of a single low gain preamplifier with output offset storage and a regenerative latch. With digital offset correction, implemented as a saturating adder embedded within the channel, the input referred offset of the comparator can increase by 3×[17]. The offset tolerance is now sufficiently high that the second preamplifier in [9] is not needed, leading to significant analog power savings. The preamplifier current is disabled after the latch has finished settling and before the start of the next bit decision phase.

Fig. 12 shows the channel synchronization block. In normal operation, the channel selects the opposite phase of the $\rm CLK_{BLK}$ (controlled by $INVCLK_j$) than the previous channel and delays the $START_j$ token by half of a clock cycle. If the channel is disabled ($EN_j = 0$) due to redundancy, $INVCLK_{j-1}$ and $START_j$ are passed transparently to the next channel; further, the channel clock is disabled and the analog bias current is shut down. This simple digital block is the only channel level overhead of the redundancy implementation.

A shared bus implements the output mux. When its outputs are ready, the channel drives the bus for half of a $\rm CLK_{BLK}$ period, and the bus is sampled by a double edge triggered flip flop. With this structure, the output mux requires no knowledge of which channels are enabled, leading to a very simple implementation. As the offset is correctable digitally to 1/4 LSB [17], the output mux is 7 bits wide.

V. Measurements

This 36 + 6-channel ADC has been fabricated in a 65-nm CMOS process. A serial interface programs the chip's configu-

Authorized licensed use limited to: MIT Libraries. Downloaded on January 17, 2009 at 11:26 from IEEE Xplore. Restrictions apply.

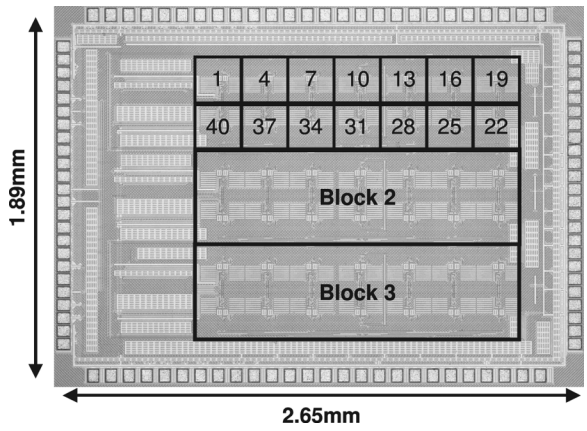

Fig. 13. Photograph of 5 mm^2 die.

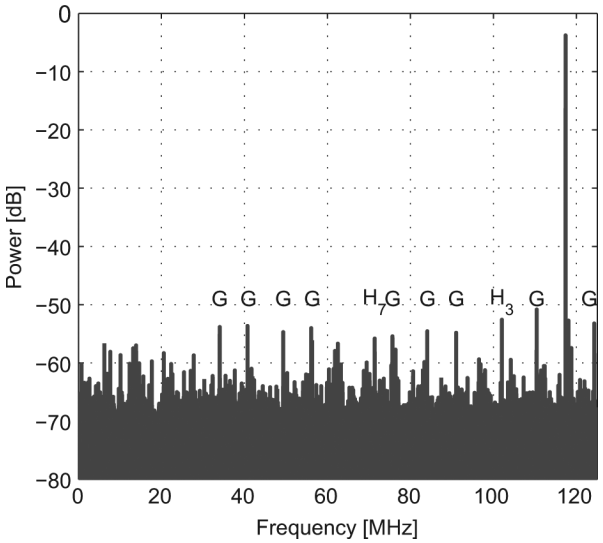

Fig. 14. FFT of a near-Nyquist 117-MHz input sampled at 250 MS/s. Harmonics and spurs from gain/timing mismatch (G) are labeled.

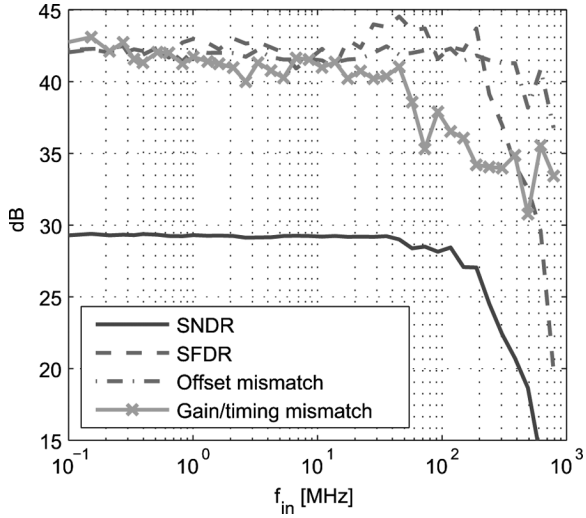

Fig. 15. Dynamic performance of the ADC versus input frequency sampled at 250 MS/s.

TABLE II
ADC SUMMARY

Technology	65-nm CMOS 1P6M
Resolution	5 bit
Input range	700 mV pp differential
Sampling rate	250 MS/s
Core supply voltage	800 mV
Sampling supply voltage	1.2 V
f in	117 MHz
SNDR	28.4 dB
SFDR	47.5 dB
THD	− 44.6 dB
Analog power	270 μW
Digital power	930 μW
Total power	1.20 mW
Active area	2.10 mm × 1.45 mm
Yield without redundancy	42%
Yield with redundancy	88%

ration registers, which include channel selection bits and the coefficients for digital offset correction. The output drivers use a differential, low-swing architecture to facilitate the high-speed interface with a logic analyzer. In addition, the first and third blocks each have a debug bus that can drive a single channel's outputs off chip. The photograph of the 5-mm^2 die is shown in Fig. 13, where the blocks are shown from top to bottom. The sampling networks are located in the center of each block. The analog input and the switched ground nodes V_{sgnd} (see Fig. 8) are routed to the channels via balanced trees to reduce systematic timing errors.

A brief summary of single chip ADC performance will be presented here, followed by a discussion of the redundancy, channel selection, and yield in the succeeding sections. At 250 MS/s, the ADC draws 1.17 mA from a 0.8-V core supply and 0.2 mA from the 1.2-V sampling supply, for a total (excluding I/O) 1.20-mW power consumption. An FFT of a 117-MHz near Nyquist input frequency is shown in Fig. 14. The largest spurs are labeled in the output spectrum. In addition to the third and fifth harmonics, several spurs are located at frequencies $k/M \cdot f_S \pm f_{\mathrm{in}}$ corresponding to gain and timing errors. These limit the spurious free dynamic range (SFDR) to 47.5 dB, but their total power has negligible impact on the SNDR of 28.4 dB. The THD is −44.6 dBc.

Fig. 15 shows the dynamic performance as the input frequency is swept. The separate contributions of the mismatches between channels are also plotted. The flat portion at low frequencies of the gain/timing curve is due to gain mismatch, and the roll off at high frequencies is due to timing skew, which is less than 15 $\mathrm{ps}_{\mathrm{rms}}$. The ENOB drops from 4.6 at DC to 4.4 at Nyquist, corresponding to a figure of merit, $P/(2f_{\mathrm{in}}2^{\mathrm{ENOB}})$, of 240 fJ/conversion step. The ADC is summarized in Table II.

A. Channel Selection

Redundancy relies upon the ability to choose the "best" channels or, alternatively, eliminate the "worst" ones; this is the task of channel selection, which is performed with an external signal source, logic analyzer, pattern generator, and a computer running Matlab. Initially, all 42 channels are characterized with a low and high frequency input signal, and a set of metrics are calculated per channel: INL, THD, SNDR, offset, gain, and timing

Authorized licensed use limited to: MIT Libraries. Downloaded on January 17, 2009 at 11:26 from IEEE Xplore. Restrictions apply.

TABLE III
CHANNEL SELECTION WEIGHTING COEFFICIENTS AND TARGET PARAMETERS

Weighting Coefficient		Target	
k_{INL}	1	T_{INL}	0.4 LSB
k_{VOS}	0.5	T_{VOS}	1 LSB
k_{SNDR_H}	3	$T_{SNDR\ H}$	28 dB
k_{THD}	1	T_{THD}	36 dB
$k_{t_{skew}}$	$(1/30)\,\mathrm{ps}^{-1}$	$T_{t_{skew}}$	30 ps

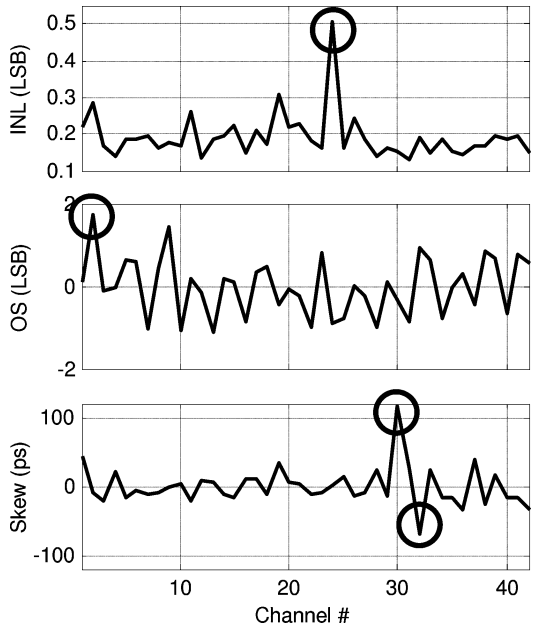

Fig. 16. Measured variation of INL, analog V_{OS}, and timing skew across a single chip.

skew. For each metric, a per-channel penalty is assessed for any deviation from a target, as calculated in (7) for INL.

$$p_{\mathrm{INL,i}} = \begin{cases} \mathrm{INL}_i - T_{\mathrm{INL}} & \mathrm{INL}_i > T_{\mathrm{INL}} \\ 0 & \text{o.w.} \end{cases} \tag{7}$$

This target should be more conservative than the actual specification in order to eliminate marginal channels when possible (i.e., if there are fewer than 2 channels per block that fail a specification). A total penalty for channel i is a weighted sum of these per-metric penalties,

$$P_i = \sum_j k_j p_{j,i} \quad j \in \{\mathrm{INL},\ V_{\mathrm{OS}},\ \mathrm{SNDR},\ \mathrm{THD},\ t_{\mathrm{skew}}\}. \tag{8}$$

A summary of targets and coefficients for the final channel selection is shown in Table III.

The channels with the highest penalty per block are eliminated until 36 channels remain. At that point, the chip is programmed with the chosen channel configuration, digital offset is measured, correction coefficients are programmed, and the final performance is evaluated.

As an example of channel selection, Fig. 16 shows a typical variation of three metrics (INL, analog offset voltage, and timing skew) across the 42 channels on a single chip. When the four circled outliers are removed, all three metrics show more than 30% improvement. Two remaining channels can be eliminated to further improve these or other metrics.

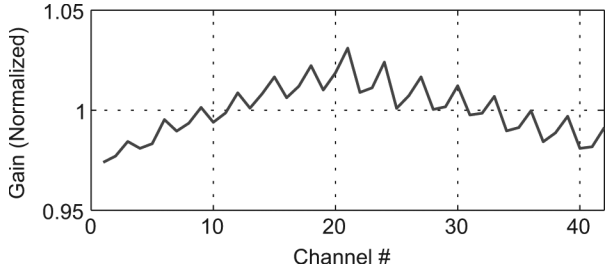

Fig. 17. Systematic variation of channel gain.

As currently implemented, channel selection would have to be done at test time. Ultimately, it is advantageous to integrate a channel selection engine on chip, both to reduce test time requirements and to be able to periodically re-perform channel selection to counteract dynamic changes in performance. For instance, in the implemented testchip channels that have the best THD at room temperature tend to have lower threshold voltage in their sampling network switches. At 85 °C, some of these channels have increased leakage from the sampling capacitor through these switches, degrading INL. The required test signal source and digital engine would be shared by all 36 channels on this chip for only limited overhead. Also, background channel characterization can be performed because there are always unused channels available.

B. Measurements Across Multiple Chips

In measurements, different types of variation can be observed that will be classified as follows. Systematic variations, which should be minimized through proper layout, are correlated for the same channel across multiple chips. Random variations can be global, which is the variation in chip-to-chip mean values. They can also be local variations that are correlated within a block (e.g., variations in the global sampling switch) or are completely independent between channels. The implemented channel redundancy is best at correcting independent channel errors and is relatively ineffective at correcting global variation or variations that are clustered within a block.

For each of the metrics, Table IV shows an estimate of the relative contributions of these different types to the total variance. Due to the limited number of chips tested and correlations between some of the variations, this is only an approximation. From the table, it is shown that most of the metrics are dominated by random channel variation and therefore correctable with the redundant channels. Gain, however, is dominated by the systematic variation plotted in Fig. 17. The evident bowing is due to $I \cdot R$ drop on the reference voltage lines, which run from left to right in Fig. 13. The channels in the middle (e.g., 19–24) are located furthest to the right and see the most attenuation on V_{ref}, which is equivalent to signal gain. Still, the total distortion from gain mismatch is small enough to negligibly affect performance (cf. Fig. 15).

Authorized licensed use limited to: MIT Libraries. Downloaded on January 17, 2009 at 11:26 from IEEE Xplore. Restrictions apply.

TABLE IV
TYPES OF VARIATION

	Total variance	Systematic	Global	Inter-block	Channel
INL	$(0.184\ \text{LSB})^2$	6%	0%	0%	94%
V_{OS}	$(0.764\ \text{LSB})^2$	33%	0%	0%	67%
Gain	$(0.019)^2$	91%	0%	0%	9%
SNDR [a]	$(1.189\ \text{dB})^2$	37%	0%	0%	63%
THD [a]	$(3.46\ \text{dB})^2$	39%	2%	0%	59%
t_{skew} (ps)	$(23.609)^2$	49%	0%	1%	50%

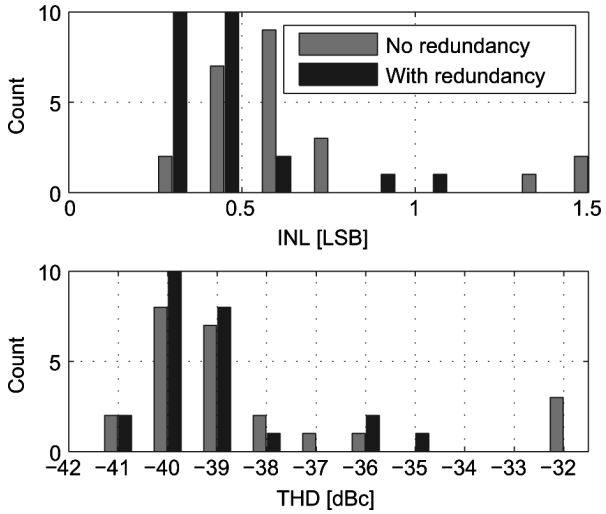

Fig. 18. Histogram showing improvement of worst channel's INL and THD after channel selection.

TABLE V
SUMMARY OF YIELD DATA

Metric	Specification	Yielding Chips out of 24	
		No Redundancy	With Redundancy
INL	≤ 0.6 LSB	14	21
V_{OS}	≤ 2 LSB	19	24
SNDR	≥ 27 dB	19	22
THD	≥ 33 dB	21	24
SNDR [a]	≥ 27 dB	17	23
All specs		10	21

The improvement of INL and THD after channel selection has been performed is shown in Fig. 18. The INL or THD are for the worst active channel on a chip. Redundancy compresses both histograms and completely removes the distant outliers. A summary of the yield data for the 24 measured chips is given in Table V. As shown, all metrics show improvement, and the percentage of chips that meet all the specifications increases from 42% to 88%. Looked at another way, redundancy reduces the number of failing chips by more than 4×.

VI. CONCLUSION AND DISCUSSION

This paper has demonstrated the use of parallelism in a mixed-signal context to lower supply voltages and to re-bias analog circuits into the subthreshold region for overall energy savings. The disadvantages of interleaving have been addressed through architectural solutions. While these architectural techniques have been presented in the context of this low power ADC, they are also applicable to more conventional interleaved ADCs that push the sampling speed boundary. A strict clock partitioning restricts the high-speed, precise domain to only a small chip area, with a lower speed clock distributed in an unbalanced network across all the channels. The global passive sampling network that mitigates timing skew has been extended for longer, overlapped sampling windows between channels without introducing significant crosstalk.

Finally, redundancy has been proposed as an efficient mechanism to counter yield loss due to local variation. With only a 17% area overhead, the yield is increased from 42% to 88%. Redundancy, along with sizing and calibration, are the essential elements to improve yield in a system. While this work has focused almost solely on redudancy, calibration and redundancy are fully compatible techniques; this work combined channel redundancy with offset calibration, but much more extensive combinations are possible. At higher resolutions, increasingly complex calibration techniques can be used without decreasing energy efficiency. Calibration is very powerful when all error sources are known at the design time and if they are easily corrected, but calibration suffers due to the underlying statistics in parallel systems. In constrast, redundancy relies on replacing, instead of fixing, failing elements and can even correct for isolated manufacturing defects. In addition, its effectiveness improves with increased levels of parallelism. While relatively unexplored to date, redudancy has great potential for mixed signal circuits.

ACKNOWLEDGMENT

The authors would like to thank Texas Instruments for fabricating the chip. They would also like to thank N. Verma and H.-S. Lee for discussions throughout the design process.

REFERENCES

[1] D. Draxelmayr, "A 6 b 600 MHz 10 mW ADC array in digital 90 nm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2004, pp. 264–265.
[2] S.-W. M. Chen and R. W. Brodersen, "A 6-bit 600-MS/s 5.3-mW asynchronous ADC in 0.13-μm CMOS," *IEEE J. Solid-State Circuits*, vol. 41, no. 12, pp. 2669–2680, Dec. 2006.
[3] B. P. Ginsburg and A. P. Chandrakasan, "500-MS/s 5-b ADC in 65-nm CMOS with split capacitor array DAC," *IEEE J. Solid-State Circuits*, vol. 42, no. 4, pp. 739–747, Apr. 2007.
[4] G. Van der Plas, S. Decoutere, and S. Donnay, "A 0.16 pJ/conversion-step 2.5 mW 1.25 GS/s 4 b ADC in a 90 nm digital CMOS process," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2006, pp. 566–567.
[5] G. Van der Plas and B. Verbruggen, "A 150 MS/s 133 μW 7 b ADC in 90 nm digital CMOS using a comparator-based asynchronous binary-search sub-ADC," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2008, pp. 242–243.

Authorized licensed use limited to: MIT Libraries. Downloaded on January 17, 2009 at 11:26 from IEEE Xplore. Restrictions apply.

[6] M. Boulemnakher *et al.*, "A 1.2 V 4.5 mW 10 b 100 MS/s pipelined ADC in a 65 nm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2008, pp. 250–251.
[7] W. Black and D. Hodges, "Time interleaved converter arrays," *IEEE J. Solid-State Circuits*, vol. SSC-15, pp. 929–938, Dec. 1980.
[8] K. Poulton *et al.*, "A 20 GS/s 8 b ADC with a 1 MB memory in 0.18 μm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, 2003, pp. 318–496.
[9] B. P. Ginsburg and A. P. Chandrakasan, "Dual time-interleaved successive approximation register ADCs for an ultra-wideband receiver," *IEEE J. Solid-State Circuits*, vol. 42, no. 2, pp. 247–257, Feb. 2007.
[10] P. Schvan *et al.*, "A 24 GS/s 6 b ADC in 90 nm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2008, pp. 544–545.
[11] N. Kurosawa *et al.*, "Explicit analysis of channel mismatch effects in time-interleaved ADC systems," *IEEE Trans. Circuits Syst. I, Fundam. Theory Appl.*, vol. 48, no. 3, pp. 261–271, Mar. 2001.
[12] K. Dyer *et al.*, "An analog background calibration technique for time-interleaved analog-to-digital converters," *IEEE J. Solid-State Circuits*, vol. 33, no. 12, pp. 1912–1919, Dec. 1998.
[13] D. Fu *et al.*, "A digital background calibration technique for time-interleaved analog-to-digital converters," *IEEE J. Solid-State Circuits*, vol. 33, no. 12, pp. 1904–1911, Dec. 1998.
[14] S. Jamal, "A 10-b 120-Msample/s time-interleaved analog-to-digital converter with digital background calibration," *IEEE J. Solid-State Circuits*, vol. 37, no. 12, pp. 1618–1627, Dec. 2002.
[15] B. P. Ginsburg and A. P. Chandrakasan, "Highly-interleaved 5 b 250 MS/s ADC with redundant channels in 65 nm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2008, pp. 240–241.
[16] A. Chandrakasan, S. Sheng, and R. Brodersen, "Low-power CMOS digital design," *IEEE J. Solid-State Circuits*, vol. 27, no. 4, pp. 473–484, Apr. 1992.
[17] B. P. Ginsburg and A. P. Chandrakasan, "The mixed signal optimum energy point: Voltage and parallelism," in *Proc. 45th Design Automation Conf.*, Jun. 2008, pp. 244–249.
[18] M. del Mar Hershenson, "Design of pipeline analog-to-digital converters via geometric programming," in *Proc. IEEE/ACM Int. Conf. Computer Aided Design*, Nov. 10–14, 2002, pp. 317–324.
[19] J. McCreary and P. Gray, "All-MOS charge redistribution analog-to-digital conversion techniques. I," *IEEE J. Solid-State Circuits*, vol. 10, no. SSC-12, pp. 371–379, Dec. 1975.
[20] H.-S. Lee, D. Hodges, and P. Gray, "A self-calibrating 15 bit CMOS A/D converter," *IEEE J. Solid-State Circuits*, vol. SSC-19, no. 12, pp. 813–819, Dec. 1984.
[21] S. Schuster, "Multiple word/bit line redundancy for semiconductor memories," *IEEE J. Solid-State Circuits*, vol. SSC-13, no. 10, pp. 698–703, Oct. 1978.
[22] M. Flynn, C. Donovan, and L. Sattler, "Digital calibration incorporating redundancy of flash ADCs," *IEEE Trans. Circuits Syst. II, Analog. Digit. Signal Process.*, vol. 50, no. 5, pp. 205–213, May 2003.
[23] Y.-C. Jenq, "Digital spectra of nonuniformly sampled signals: Fundamentals and high-speed waveform digitizers," *IEEE Trans. Instrum. Meas.*, vol. IM-37, no. 6, pp. 245–251, Jun. 1988.
[24] K. Poulton, J. Corcoran, and T. Hornak, "A 1-GHz 6-bit ADC system," *IEEE J. Solid-State Circuits*, vol. SSC-22, no. 12, pp. 962–970, Dec. 1987.
[25] K. Y. Kim, N. Kusayanagi, and A. Abidi, "A 10-b, 100-MS/s CMOS A/D converter," *IEEE J. Solid-State Circuits*, vol. 32, no. 3, pp. 302–311, Mar. 1997.
[26] B. P. Ginsburg, "Energy-efficient analog-to-digital conversion for ultra-wideband radio," Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, MA, 2007.
[27] M. Gustavsson and N. Tan, "A global passive sampling technique for high-speed switched-capacitor time-interleaved ADCs," *IEEE Trans. Circuits Syst. II, Analog. Digit. Signal Process.*, vol. 47, no. 9, pp. 821–831, Sep. 2000.
[28] M. Gustavsson and N. Tan, "Method and device for analogue to digital conversion," U.S. Patent 6,160,508, Dec. 12, 2000.
[29] M. Waltari and K. Halonen, "Timing skew insensitive switching for double sampled circuits," in *Proc. IEEE Int. Symp. Circuits and Systems (ISCAS)*, 1999, pp. 61–64.
[30] S. Gupta *et al.*, "A 1 GS/s 11 b time-interleaved ADC in 0.13 μm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2006, pp. 576–577.

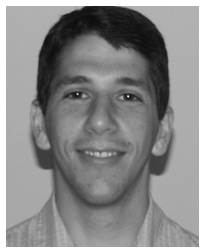

Brian P. Ginsburg (S'04) received the B.S. and M.Eng. degrees in electrical engineering and computer science and the Ph.D. degree from the Massachusetts Institute of Technology (MIT), Cambridge, MA, in 2003 and 2007, respectively.

He is currently a mixed signal designer within the wireless unit of Texas Instruments in Dallas, TX. His research interests include low power mixed-signal and RF circuit design, analog-to-digital converters, and system energy optimization.

Dr. Ginsburg was named a Siebel Scholar in 2003, received the NDSEG Fellowship in 2004, and was a winner of the 2006 ISLPED International Low Power Design Contest.

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

Since September 1994, he has been with the Massachusetts Institute of Technology, Cambridge, where he is currently the Joseph F. and Nancy P. Keithley Professor of Electrical Engineering. His research interests include low-power digital integrated circuit design, wireless microsensors, ultra-wideband radios, and emerging technologies. He is a coauthor of *Low Power Digital CMOS Design* (Kluwer, 1995) and *Digital Integrated Circuits* (Pearson Prentice-Hall, 2003, 2nd edition). He is also a co-editor of *Low Power CMOS Design* (IEEE Press, 1998), *Design of High-Performance Microprocessor Circuits* (IEEE Press, 2000), and *Leakage in Nanometer CMOS Technologies* (Springer, 2005).

Dr. Chandrakasan has received several awards including the 1993 IEEE Communications Society's Best Tutorial Paper Award, the IEEE Electron Devices Society's 1997 Paul Rappaport Award for the Best Paper in an EDS publication during 1997, the 1999 Design Automation Conference Design Contest Award, and the 2004 DAC/ISSCC Student Design Contest Award. He has served as a technical program co-chair for the 1997 International Symposium on Low Power Electronics and Design (ISLPED), VLSI Design '98, and the 1998 IEEE Workshop on Signal Processing Systems. He was the Signal Processing Sub-Committee Chair for ISSCC 1999–2001, the Program Vice-Chair for ISSCC 2002, the Program Chair for ISSCC 2003, and the Technology Directions Sub-Committee Chair for ISSCC 2004–2006. He was an Associate Editor for the IEEE JOURNAL OF SOLID-STATE CIRCUITS from 1998 to 2001. He serves on the SSCS AdCom and is the meetings committee chair. He is the Technology Directions Chair for ISSCC 2007.

Authorized licensed use limited to: MIT Libraries. Downloaded on January 17, 2009 at 11:26 from IEEE Xplore. Restrictions apply.