

A Nyquist-Rate Pipelined Oversampling A/D Converter

Susanne A. Paul, Hae-Seung Lee, *Fellow, IEEE*, John Goodrich, Titiimaea F. Alailima, and Daniel D. Santiago

Abstract—A pipelined Δ - Σ analog-to-digital-converter architecture is described that incorporates the high speed of pipelined converters and the high resolution of oversampling quantization. A prototype, containing both modulation and decimation circuits on a single chip, is implemented using a 1.2- μ m commercial CMOS process. It uses charge-coupled-device elements to perform pipelined analog operations. It exhibits a maximum data rate of 18 MHz, a signal-to-noise ratio of 74 dB, spurious-free dynamic range of 78 dB, differential nonlinearity of <0.15 LSB at 13 bits, and power dissipation of 324mW.

Index Terms—Analog-to-digital (A/D), charge-coupled device (CCD), delta-sigma, oversampling, pipeline, sigma-delta.

I. INTRODUCTION

OVERSAMPLING and noise-shaping techniques, such as Δ - Σ modulation, are widely used in analog-to-digital conversion to achieve accuracy that exceeds that of integrated-circuit components. Such converters have an inherent tradeoff between accuracy and speed, whereby resolution in amplitude is achieved at the expense of resolution in time. They have limited data rates because their internal circuits must operate over many clock cycles to produce a single result. Much attention has been focused on improving the speed of Δ - Σ analog-to-digital converters (ADC's) through use of higher order modulators [1], multibit feedback [1], and multibit architectures with single-bit feedback [2]. However, data rates remain limited to less than a few megahertz and are not easily extended.

A pipelined oversampling architecture is described here that circumvents the speed-resolution tradeoff of conventional oversampling ADC's by performing spatial, rather than temporal, oversampling. It combines the high-resolution quantization capability of Δ - Σ techniques with the high speed of pipelined architectures so that both of these attributes are achievable. In comparison to conventional oversampling converters, power is improved as a result of a charge-domain implementation, reduced sensitivity to thermal noise, simplified decimation, and a reduced circuit speed requirement, which permits voltage scaling and use of low-power technologies. A pipelined

architecture is also well suited for processing presampled signals because, like the parallel-channel architecture in [3], it performs Nyquist-rate sampling.

Section II presents pipelined oversampling quantization algorithms, their implementation in a converter, and their associated design considerations. Section III introduces circuit techniques and charge-domain building blocks, used to build such a converter using charge-coupled device (CCD)/CMOS technology. Last, Section IV presents details of the prototype implementation and measured test results.

II. PIPELINED OVERSAMPLING ARCHITECTURE

A conventional Δ - Σ ADC, shown in Fig. 1(a), includes a single modulator operating at a speed greater than the converter's output data rate by a factor equal to the oversampling ratio. A time sequence of many modulator outputs is generated. This sequence is filtered and downsampled in the decimator to produce each result. In contrast, the pipelined oversampling converter (POSC), shown in Fig. 1(b), performs oversampling in space. Its modulator and decimator loops are unraveled into a pipeline so that consecutive cycles of operation occur along consecutive pipeline stages, rather than within a single piece of hardware. Incoming signals are sampled by the first modulator stage, using Nyquist sampling, and are processed by an oversampling quantization algorithm along the pipeline. Each modulator stage produces a digital output, and these outputs are processed by the decimator, which is pipelined as well. The converter's digital result is produced by the final decimator stage. Although both pipelined and conventional devices have a long latency, the pipelined device computes a new result every cycle and achieves an output data rate that is many times faster than that from a time-oversampling device.

A POSC is not subject to a speed-resolution tradeoff. Its output data rate equals its internal clock rate, and no higher speed circuits are required. On the other hand, its resolution is determined by its pipeline length. Accuracy and speed for such a device can be independently adjusted within the constraints of a given process technology. The higher speed of a POSC is achieved at the cost of additional hardware. While time-interleaved ADC channels could also be used for such a speed improvement, these techniques differ in an important regard. All signals in a pipelined device exercise the same circuit path, eliminating the need for accurate channel-to-channel matching.

A POSC algorithm operates over two dimensions, time and space, as shown in Fig. 2. An example sinusoidal input is

Manuscript received April 22, 1999; revised July 8, 1999. This work was supported in part by the Department of Defense and in part by a National Science Foundation Graduate Fellowship.

S. A. Paul is with the MIT Lincoln Laboratory, Lexington, MA 02420 USA and the Massachusetts Institute of Technology, Cambridge, MA 02139 USA.

H.-S. Lee and J. Goodrich are with the Massachusetts Institute of Technology, Cambridge, MA 02139 USA.

T. F. Alailima and D. D. Santiago are with the MIT Lincoln Laboratory, Lexington, MA 02420 USA.

Publisher Item Identifier S 0018-9200(99)08969-6.




Fig. 1. Comparison between (a) conventional time oversampling and (b) pipelined oversampling.




Fig. 2. Pipelined oversampling operates over both time and space dimensions.




Fig. 3. Computational view of pipelined oversampling.

sampled, using Nyquist sampling, at the beginning of the pipeline along the time dimension. Each sample is then passed unchanged along the pipeline and used as the input to a Δ - Σ quantization algorithm, which occurs along the space dimension. Each slice in time, corresponding to a single input sample, is processed independently from its neighbors. Consequently, a POSC is indistinguishable from a Nyquist-sampling converter and supports input bandwidths up to half the clock rate. As with any Nyquist sampling ADC, the advantages of Nyquist sampling in a POSC occur for the price that its input must be bandlimited with an anti-alias filter.

Computationally, a POSC can be viewed as a time-oversampling converter that is configured as shown in Fig. 3. A sampling clock is used to reset integrators in the modulator, reset memory in the decimator, and perform sampling of an analog input. The sampled input is held constant while the modulator and decimator operate on it over many cycles using a higher frequency internal clock. The converter's result is transferred to its output on the next sampling clock edge, at the same time that the next phase of resetting and sampling occurs.




Fig. 4. Second-order analog-integration quantization algorithm, shown in cyclic form. (a) Modulator. (b) Decimator.

A. Analog-Integration Quantization Algorithm

Second-order modulation is best suited for a POSC. A first-order approach requires significantly more hardware, power, and area and is susceptible to pattern noise because signals are constant throughout a POSC's quantization process. Third- or higher order modulators have little advantage in a pipelined device, where speed and resolution are decoupled, because they do not improve speed, they bring the danger of instability, and they provide only a small reduction in hardware.

A second-order analog-integration algorithm for the quantization of a single input sample is shown in cyclic form in Fig. 4. The discrete-time index n in the analysis below represents time in a cyclic configuration. When the algorithm is unraveled along a pipeline, each successive time step occurs in a later pipeline stage. In this case, n represents both time and space indexes during the quantization of a single input sample but does not represent the index of successive input samples.

The modulator input at pipeline stage n , which is zero before the first pipeline stage and constant over all pipeline stages, is described by

$$s_i[n] = S_g u[n]. \quad (1)$$

It is a step function whose amplitude equals the sampled value S_g that was captured before the first pipeline stage. Both integrators are reset to zero before the pipeline. A coarse ADC with a full-scale value of R_f generates an r -bit digital representation of the second integrator value i_2 . As in the truncated feedback approach of [2], the k most significant bits are included in the feedback path, and the remaining lower order bits form a digital truncation signal t_k . When quantization noise and circuit inaccuracies in the coarse ADC are modeled as an additive error e_r , the modulator output w_k is given by

$$w_k[n] = \frac{S_g}{R_f} u[n-2] + (e_r[n] - 2e_r[n-1] + e_r[n-2]) + (t_k[n] - 2t_k[n-1] + t_k[n-2]). \quad (2)$$

This signal contains a delayed version of the input sample, plus the second-order difference of errors e_r and t_k .

The decimator shown in Fig. 4(b) is a third-order accumulator. Since w_k tracks the analog input on average, d_1 , d_2 , and d_3 increase linearly, quadratically, and cubically with n at a rate proportional to S_g . The first two stages in the decimator reverse the differentiation in (2) and amplify the signal S_g . Their output is

$$d_2[n] = \frac{S_g}{2R_f}(n-3)(n-2)u[n-3] + e_r[n-2] - t_k[n-2]. \quad (3)$$

Truncation error t_k is digitally canceled before the final accumulation. The output of the third accumulator is

$$d_3[n] = \frac{S_g}{6R_f}(n-4)(n-3)(n-2)u[n-4] + \sum_{j=2}^{n-3} e_r[j]. \quad (4)$$

This decimator is referred to as error averaging because its final output contains a sum of e_r . After P stages, the downsampled result is reconstructed using the relation

$$S_g \approx d_3[P+4] \left(\frac{6R_f}{P^3 + 3P^2 + 2P} \right). \quad (5)$$

When quantization noise e_r has a constant mean square value, the signal-to-noise ratio (SNR) for a sinusoidal input with peak-to-peak value of $R_f(1-2^{-k})$ is

$$\text{SNR}[P] = \frac{2^r(1-2^{-k})(P^{5/2} + 3P^{3/2} + 2P^{1/2})}{2\sqrt{6}}. \quad (6)$$

Each doubling of the pipeline length provides an additional 2.5 bits of resolution. Each quantizer bit improves resolution by 1 bit and, therefore, allows a shorter pipeline length to be used. But this reduction in pipeline length comes at the cost of increased hardware per stage and a reduced degree of noise shaping, which increases the accuracy needed from elements in the feedforward path.

B. Matched Filter Decimation

An alternative decimation technique, not shown in the figure, is matched filtering. A POSC has two attributes, not present in time-oversampling converters, that make matched filtering possible. First, the quantization algorithm input has a precisely known form over n , which can be matched in the decimator. It is a step function whose amplitude varies with input sample. A second attribute that makes matched filtering possible is that a uniform decimator passband response is not required because the input spectrum, that of a step function, is independent of frequencies present at the converter input.

A matched filter operates on d_2 in (3), where quantization noise is unshaped. Its impulse response is a time-inverted version of the signal term in (3). The filter output is

$$d_5[P] = \sum_{n=4}^{P+3} \left[\frac{1}{4}(n-3)^2(n-2)^2 \frac{S_g}{R_f} + \frac{1}{2}(n-3)(n-2)e_r[n-2] \right]. \quad (7)$$

The value of d_2 at each time step n is amplified in proportion to its SNR. After downsampling, SNR equals

$$\text{SNR}[P] = \frac{2^r(1-2^{-k})}{4} \sqrt{\frac{1}{6}(7P^5 + 38P^4 + 53P^3 + 46P^2)}. \quad (8)$$

The ratio between (8) and (6) shows that a matched filter provides about an additional 2.4-dB SNR over the error averaging decimator of Fig. 4.

Oversampling ratio, as this term is commonly used for time-oversampling converters, is not clearly defined in a pipelined device. With respect to quantization, the step function input to a pipelined algorithm is not bandlimited. With respect to sampling, input bandwidths occupy the full Nyquist range. The analogy is further complicated by the fact that in a POSC, integrators are reset and the modulator and decimator impulse response lengths are equal. The resolution of a time-oversampling converter depends on its oversampling ratio, whereas that in a POSC depends on its pipeline length P . Two comparisons are presented between these for the example of second-order modulation. First, a POSC with P stages and SNR given by (8) is compared to a time-oversampling converter with an oversampling ratio of P and a theoretical SNR of

$$\text{SNR}[P] = \frac{\sqrt{7.5}2^r(1-2^{-k})P^{5/2}}{\pi^2}. \quad (9)$$

The SNR of these approaches is nearly identical. Second, a POSC with P stages and SNR given by (8) is compared to a time-oversampling converter with a decimator impulse response length of P and SNR given by (9) for an oversampling ratio of $P/3$. In this metric, which compares SNR per modulator output, a POSC provides an improvement of about 24-dB SNR over time oversampling.

C. Digital-Integration Quantization Algorithm

For reasons described in Section III, a CCD implementation was chosen for the prototype. One difficulty with a CCD-based analog-integration approach is that signals in the first integrator must be replicated at the input to the second integrator. Charge replication circuits, described in Section III, provide limited linearity and are subject to thermal noise and coupling. Since this operation occurs after the first integrator, the impact of its linearity and noise is attenuated by a factor of $(P+2)/3$ due to first-order noise shaping and gain in the outer feedback loop. However, the need to suppress replicator nonideality still sets a lower limit on pipeline length P and, therefore, prevents full use of truncated feedback, which otherwise can be used to reduce P . This difficulty is eliminated in an alternative approach, referred to as digital integration.

The transition between the analog and digital-integration quantization algorithms is described with reference to Fig. 5. The feedback digital-to-analog converter (DAC) is moved from before to after i_1 so that the first stage of integration occurs digitally. In this configuration, signal and reference quantities enter the converter as analog references to the upper and lower DAC's and are multiplied by the digital DAC inputs. The digital input to the upper DAC's is u_1 , a signal-




Fig. 5. Transition between analog and digital-integration quantization algorithms.




Fig. 6. Second-order digital-integration quantization algorithm, shown in cyclic form.

independent value equal to the pipeline stage number. The digital input to the lower DAC is f_b , equal to $d_1 + 2w_k$, which is easily generated from signals in the decimator. The elements in Fig. 5 that compute signals u_1 and v_1 are, therefore, unnecessary. The digital-integration algorithm in Fig. 6 differs from that in Fig. 5 in that these computations are eliminated. The input-output transfer characteristics of the resulting digital-integration modulator are identical to those in (2) for analog integration, and identical decimation techniques can be applied.

The implementation of a digital-integration algorithm in a pipeline configuration is shown in Fig. 7. Two analog channels, s_i and i_1 , and three digital channels, d_1 , d_2 , and d_3 , flow through the pipeline. The converter's input sample, captured at the beginning of the pipeline, is passed unchanged along the s_i channel. At each stage, this sample is used as the reference input to the upper DAC. The digital feedback signal f_b , generated within the previous stage, is used as the digital input to the lower DAC. Values within the d_2 and d_3 channels are not needed until reconstruction at the end of the pipeline, and their computations can be distributed across multiple stages to reduce adder speed requirements.

Single-bit feedback is often used in time-oversampling converters because it has an inherent linearity advantage over multibit feedback. In such devices, digital integration is not desirable because it eliminates this advantage. However, single-bit feedback in a pipelined converter does not have an inherent linearity advantage. A similar characteristic occurs in time-oversampling converters with multibit feedback, MASH [4], or feedforward architectures. In a POSC, each of the




Fig. 7. Pipeline stage contents for digital integration.

converter's feedback operations occurs in a different stage and is performed using unique circuit elements. Mismatches cause nonlinearity regardless of the number of bits in each DAC. As a result, multibit feedback is desirable in a pipelined converter because it provides a larger input range, more predictable behavior, and less susceptibility to pattern noise.

In a time-oversampling converter with multibit feedback, a few elements are used repeatedly. Fortunately, the effects of mismatches are less severe in a POSC because it inherently achieves the benefits of dynamic element matching. Mismatches are reduced due to averaging among elements in the many stages. Tolerable mismatch is defined as the capacitance variation of each DAC element, with respect to its nominal value, at which the rms values of ideal quantization noise and DAC-related noise are equal. Neither multibit feedback nor digital integration changes DAC matching requirements appreciably. Tolerable DAC mismatch for an analog-integration architecture is given by

$$\sqrt{\frac{2^{-2r+k}P}{3\sum_{j=1}^P(j^2+3j)^2}}. \quad (10)$$

Mismatches in a digital-integration architecture experience first-order noise shaping, whereas those for analog integration do not. However, digital-integration DAC's have larger full-scale references. The result is that the mismatch tolerances of these two approaches, with respect to their DAC full-scale values, are similar.

III. CCD/CMOS CONVERTER IMPLEMENTATION

A. CCD/CMOS Technology

The POSC prototype is accomplished using a combination of CCD and CMOS circuits, fabricated in a generic CMOS process. Although CCD's are not essential to the concept, a combination of these circuit techniques enables performance that would be difficult from either one alone. CMOS plays a vital role in such devices by providing digital logic and




Fig. 8. Dual-gate CCD's in standard CMOS.

CCD support circuitry. CCD's provide fully depleted circuits, such as charge transfer, addition, integration, and conditional transfer, which are highly accurate, low in power, simple, and compact [5]. Charge transfer efficiencies as high as 10^{-7} have been demonstrated in imager applications [6]. Because CCD's are not subject to thermal noise, charge injection, or coupling from clocks or the substrate, high signal integrity is possible throughout hundreds of transfers, amidst noisy digital circuitry. Since their gain and linearity are determined by charge conservation, circuit transfer characteristics are insensitive to device parameters, and highly accurate circuit-to-circuit matching is possible. Finally, fully depleted circuits are strictly dynamic with only capacitive switching current and can, therefore, be performed with low power and high speed. These features make analog pipelines with hundreds of stages feasible in a CCD device.

Structurally, CCD's are similar to NMOS transistors. Their difference lies in their methods of interconnection. In a CCD circuit, adjacent gates are brought sufficiently close that their channel regions overlap and no diffusion is present between them. Although CCD devices are traditionally built using specialized fabrication, their most basic requirements are met by standard CMOS processes that include double-poly for capacitors. An example of a CCD structure in CMOS processing is shown in Fig. 8. Overlapping structures are formed by use of parasitic second-poly active gates. Such surface-channel CCD's have lower charge transfer efficiency than buried-channel devices because of surface state effects. However, they bring other advantages. First, because they use surface channels and are enhancement mode, they are compatible with CMOS voltage levels and a grounded substrate and do not require high voltage to drain their charge. Second, their threshold difference between poly1 and poly2 gates provides a built-in barrier that eliminates the need for offset barrier and storage clock potentials.

B. Dynamic Double Sampling Circuit

Functions such as charge generation, wire transfer, charge sensing, D/A conversion, and D/A subtraction are possible from nondepleted CCD circuits. They are accomplished in the POSC prototype using a technique referred to as dynamic double sampling (DDS). A core DDS circuit is shown schematically in Fig. 9(a). Its energy level diagrams in (b)–(e) illustrate conditions under each gate. Solid lines depict empty-well channel potential, and filled regions depict electron energy. Higher levels correspond to higher electron energy and lower potential. The objective of this circuit is to integrate incoming charge, introduced at the circuit's input, in




Fig. 9. (a) Dynamic double sampling circuit. (b)–(e) Four phases of operation.

a CCD receiving well. A precharge path contains gates G1 and G24. A sensing path contains primary and secondary circuits, the first consisting of G3 and G24, and the second formed from G8 and G9. Although G24 is shown as two gates in the figure, in most cases it is implemented as a single gate.

Operation occurs over four phases. During the fill phase, in Fig. 9(b), v_{fg} is pulled low through M1, the region underneath G24 is flooded with charge, and electrons consumed during the previous generation cycle are replenished. During the spill phase, in (c), v_{fg} is left floating and the precharge path is enabled by raising G1. Initially, the amplifier output is saturated at V_h and electrons flow from v_{fg} to the drain V_d . Current decreases as v_{fg} rises because the gate-to-source voltage of G24 is reduced. Once v_{fg} reaches a voltage of $V_r - V_h/A$, where A is the amplifier gain, v_{fg} begins to fall. During this transition

$$\frac{dv_{fg}(t)}{dt} = \frac{1}{C_{fg}} I(AV_r - (A+1)v_{fg}(t)) \quad (11)$$

where $I(v_{gs})$ is the current-voltage relation of G24 and




Fig. 11. Dynamic double sampling charge generator.

linearity of these circuits is limited by a dependence on CCD well capacitance or MOS device parameters. Their speed is limited by long subthreshold time constants and sampling effects. These circuits are also sensitive to low-frequency noise. A DDS charge generator is shown in Fig. 11. Its input signal is capacitive displacement charge introduced by a CMOS clamp and sample circuit. During the fill and precharge phases, v_d is clamped high and v_m tracks the analog input. During the collection phase, v_m is clamped to a reference voltage V_n , which is usually ground. The resulting transition couples through C_1 and C_2 onto the DDS input. The output from this circuit is

$$Q_o = (V_s - V_n) \left(\frac{C_1 C_2}{C_1 + C_2 + C_{p2}} \right) \quad (18)$$

where V_s is the value of the analog input voltage s_i at the end of the precharge phase and C_{p2} represents parasitic capacitance on node v_d . Improved linearity and matching is possible from this circuit because its result depends on poly-poly capacitors, rather than CCD well capacitors. Parasitic capacitor voltage dependence on v_{fg} does not cause nonlinearity because this node is a virtual ground. Nonlinearity due to C_{p2} voltage dependence can be reduced by bootstrapping the PMOS well.

The circuit in Fig. 11 can also be used to perform sampling. In this case, the precharge through M3 is turned off before input tracking through M1 and M2 is disabled. Sampling is performed by the turnoff transition of M3, which has fixed source and drain voltages, resulting in a signal-independent sampling time. Charge injection from the input tracking switch also does not affect the result because it does not change either V_s or V_n in (18).

Another application of the DDS technique is D/A conversion and subtraction. This circuit, shown in Fig. 12, combines negative charge Q_s from a CCD register and positive D/A charge on a wire, and the resulting packet is integrated in a CCD well by a DDS block. The source of negative charge is similar to that described above for wire transfer. The N -bit D/A charge is generated by an array of $(2^N - 1)$ identically sized capacitor circuits with a thermometer code digital input. The D/A capacitor inputs are clamped to ground during precharge and switched to the reference V_c during the




Fig. 12. Dynamic double sampling D/A conversion and subtraction.




Fig. 13. Dynamic double sampling charge replication and sensing.

collection and sensing phases. The circuit output is

$$Q_o = \left(Q_s + \sum_{i=0}^{2^N-1} w(i) C_1 V_c \right). \quad (19)$$

A final DDS circuit that is described is charge replication. Since charge-domain operations are destructive, they consume their signals. Charge replication is needed when packets are to be used multiple times. It is also used to convert charge to voltage, for purposes such as comparison, without altering the original packet. Previous charge replication techniques have been reported with 40-dB linearity at 20 kHz [9]. The linearity of this approach is limited by a dependence on parasitic capacitances, and the speed is limited by subthreshold time constants. A DDS charge replicator, shown in Fig. 13, contains a floating gate, within a CCD register, that is connected to the input of a DDS circuit. As incoming packets are shifted beneath the floating gate, they couple through the gate-oxide capacitance, and the resulting displacement charge

$$Q_o = Q_s \frac{C_{ox}}{C_{ox} + C_{cs}} \quad (20)$$

is integrated in the DDS receiving well. C_{ox} and C_{cs} represent the gate-to-channel and channel-to-substrate capacitances of the floating-gate well. A DDS charge replicator provides reduced linearity and matching, in comparison to other DDS circuits described above, because of voltage dependence of the capacitances in (20) and mismatches in their values.

In addition to replicating charge packets, the circuit in Fig. 13 can also be used for charge-to-voltage conversion and charge comparison. For charge-to-voltage conversion, the DDS receiving wells G10 and G5 in Fig. 9 are replaced by a poly-poly capacitor that is preset. The result is sensed as a




Fig. 14. Substages 1–3 of the subpipelined digital integration prototype.

voltage on an MOS gate. For charge comparison, the DDS receiving wells are replaced by the input capacitance of a CMOS comparator or charge-integrating amplifier.

IV. POSC PROTOTYPE AND MEASURED RESULTS

A. Subpipelined Prototype Implementation

A prototype device was built to demonstrate the pipelined oversampling concept. It utilizes a digital-integration architecture with 12 pipeline stages, a 5-bit feedforward quantizer, and 2-bit truncated feedback. The digital-integration configuration in Fig. 7 contains only a single delay, and its sequence of operations must be performed within one clock cycle. Since additional stages are easily added in a CCD pipeline, throughput is improved for the prototype by distributing its operations within each stage across eight substages. The first three of these substages are shown in Fig. 14. Substage boundaries are delineated by dashed lines and, for simplicity, delay elements are omitted from the figure. A differential structure is used because it allows cancellation of even-order harmonics and common-mode noise, reduces the need for common-mode charge rejection and accurate zero-reference biases, and allows complementary addition, which is a fully depleted and highly accurate operation, to be performed in place of subtraction, which is nondepleted and less accurate.

D/A conversion and subtraction occur during substage 1. Instead of replicating a single delayed input sample in each stage, an array of separate input samples, with the same nominal values, are captured before the pipeline and passed along the s_i channel. A total of n of them are used in

the n th stage to implement the signal DAC. Accurate level placement is possible from this DAC because it is determined by charge generator matching at the beginning of the pipeline. Nevertheless, a POSC is highly tolerant of misplaced levels, up to a few percent, within its signal DAC because every incoming signal exercises the same elements. Mismatches only alter converter gain, provided they are not large enough to overload the modulator.

Each of the substages 2–6 is used to compute one ADC bit. A 1-bit-per-stage pipelined CCD quantizer, described in [10], is used for this purpose. In substage 2, the integrator packets i_{2p} and i_{2m} are nondestructively sensed, using the floating-gate elements labeled S , and are compared to generate the most significant bit of w_k . Reference signals R_f are provided to a pair of scaling channels, and, in substage 3, these signals are divided in half using charge splitting circuits. The comparator result from substage 2 controls the conditional transfer elements labeled CT. Using these elements, a scaling packet is added to a modification channel on either the positive or negative side, whichever has a smaller value. No addition occurs on the side with a larger value. The comparator input in substage 3 is the sum of the integrator and modification signals. This is accomplished by covering both channels with a single floating gate. After the 5-bit quantization is complete, signals in both the scaling and modification channels are discarded. The binary scaled references in this configuration are used only as part of the 5-bit quantization and are not combined with signals in the integrator. Their inaccuracies are indistinguishable from comparator errors and are suppressed by second-order noise shaping.

The prototype includes a dual pipeline structure, illustrated in Fig. 15. The oversampling pipeline is divided into even and odd halves, and every signal is passed simultaneously through both of them. In each stage, the 2-bit feedback signal w_k is completed after substage 3 and is combined with d_1 in substage 4. The even pipeline is delayed by four cycles from the odd, so that stage n begins its operations just as the feedback signal from stage $n - 1$ is completed. Advantages of a dual pipeline structure include a reduction in pipeline latency by half, which alleviates accuracy requirements of signal-channel charge transfers, and a moderate decrease in signal channel hardware and power.

Constraints were imposed on the prototype's design by its 1.2- μm process geometries. For smaller process geometries, a straight oversampling pipeline, as described above, is believed to be the preferred approach. To accommodate a 1.2- μm process, the prototype includes the modified pipeline of Fig. 15, with front-end DAC's. In this configuration, a 6-bit estimate of the input signal is computed before the pipeline, and this estimate is used to digitally predict the value of f_b that will occur in each pipeline stage. The prototype makes use of this prediction and the fact that input signals are constant throughout the pipeline to move 4 bits from the reference DAC in each stage to DAC's at the beginning of the pipeline, where they are not subject to pipeline pitch constraints.

The resulting two-stage configuration does not change the converter's computations or change the minimum resolution required from the oversampling pipeline. No amplification




Fig. 15. Dual-pipeline configuration used for the prototype.

occurs, and signals after the DAC's are passed forward with unity gain, as determined by charge conservation. Only about 5-bit accuracy is required from the initial prediction because oversampling can reverse its decisions. The DAC is implemented using 256 identically sized elements, controlled by thermometer code inputs. Capacitors in the DAC are made large, with a smaller voltage swing, to improve matching, at the expense of an increase in charge-referred thermal noise. Since the initial DAC elements are functionally identical to those later in the pipeline, they have the same accuracy requirements. However, since fewer DAC elements are used throughout the converter in this approach, averaging is reduced. The pipelines are oriented so that gradual process variations across the chip are inverted in the even and odd channels, result in high frequency error, and are suppressed within the decimator.

B. Measured Prototype Results

One goal of the POSC prototype was to demonstrate that high-performance CCD devices are achievable using only standard double-poly CMOS processing. Two versions of the prototype were fabricated. The first, the POSC1, is from a $1.2\text{-}\mu\text{m}$, 5-V, double-poly, double-metal CMOS process from Orbit Semiconductor with $225\text{-}\text{\AA}$ gate oxide. The second, the POSC2, is from a $0.35\text{-}\mu\text{m}$, double-poly, double-metal MOSIS process from TSMC. The objective of the POSC2 was to determine the capability of more advanced processes for building dual-gate CCD circuits. The design is a direct $2\times$ shrink of the POSC1. It uses $0.6\text{-}\mu\text{m}$ design rules and thicker $150\text{-}\text{\AA}$, 5-V gate oxide that is available in this process. Both prototypes contain 6200 CCD and 8800 CMOS devices in their modulators and 22600 CMOS devices in their decimators. The modulator, decimator, and all necessary support logic are integrated onto a single chip. Major functional blocks are indicated on the die micrographs in Fig. 16. Despite the prototype's use of unsupported second-poly active gates, and




Fig. 16. Micrograph of the (a) POSC1 and (b) POSC2 prototypes.

the presence of overlapping four-layer structures, the fully functional yield of both devices was better than 90%.

Testing was done using an automated ADC testbed with synchronized clock and signal sources. Measured prototype performance is summarized in Table I. POSC1 measurements were done at an 18-MHz sampling rate with an input sinusoid of approximately 8 MHz. A 16 384-point spectral response plot is shown in Fig. 17. The spurious-free dynamic range (SFDR), given by the ratio of the fundamental to the largest harmonic, is determined by the third harmonic and is 78 dB. The relationship between harmonic distortion and input power level indicates that the second and third harmonics are dominated by nonlinearity in the input charge generation circuits and that third-order nonlinearity due to DAC mismatches is about 3 dB below this level. The impact of DAC mismatches is evident from the presence of higher order harmonics, near the others, with slowly decreasing magnitudes. These were anticipated for the design because of its front-end DAC's. The measured level is computed to correspond to an rms mismatch in DAC element values of about 0.25%. Nonlinearity due to the pipeline DAC's is reduced because some of it is translated into wide-band noise by inherent dynamic element matching. Among the devices tested, the harmonic distortion varies from 78 to 73 dB due to incomplete cancellation of even-order harmonics in the DDS circuits.

Fig. 18 shows SNR, with a peak of 74 dB, and signal-to-noise plus distortion ratio (SNDR), with a peak of 71 dB, as a function of input power. Ratios are computed over a 9-MHz bandwidth. Since the converter produces a total of 16 bits, quantization error is not limited by arithmetic width at

TABLE I
SUMMARY OF MEASURED PROTOTYPE PERFORMANCE

	POSC 1	POSC 2
Output Data Rate	18 MSPS	30 MSPS
Peak SFDR	78 dB	70 dB
Peak SNR / SNDR	74 / 71 dB	66 / 63 dB
DNL / INL @ 13 bits	± 0.15 LSB / ± 1.0 LSB	± 0.25 LSB / ± 2.5 LSB
Analog / Digital Supply	5 V / 3.3 V	5 V / 3.3 V
Input Range	2 V p-p	1.5 V p-p
Power	324 mW @ 18 MHz	230 mW @ 30 MHz
Power Breakdown % (A-CMOS / CCD / D-CMOS)	65 / 20 / 15	50 / 35 / 15
Process	Commercial CMOS	Commercial CMOS
Design Rule	1.2 μm	0.6 μm
Yield	22 of 24	26 of 27




Fig. 17. Measured spectral response at an 18-MHz data rate with an input near 8 MHz.

the output. Sampling noise is reduced in the design because 32 separate samples are captured at the converter input. An examination of noise versus input amplitude shows that noise varies across the converter's input range. As a result, the dominant source of noise is thought to be wide-band dynamic element matching noise, caused by DAC mismatches, that is passed by the decimator. Other noise sources that may be significant include coupling in nondepleted circuits that is not common-mode, and surface state trapping in the CCD's. Theoretical analysis indicates that thermal noise in the device's nondepleted circuits is not a dominant component of noise. Integral and differential nonlinearity plots in Fig. 19 were generated using histogram techniques at an 18-MHz output data rate with a sinusoidal input near 8 MHz.

Device performance is unchanged at lower operating frequencies. As a result, incomplete charge transfer is believed to be an insignificant source of error. Performance degrades rapidly and exhibits digital failures at speeds above 18 MHz. This indicates that data rates are limited by the speed of CMOS supporting circuits and the need to generate and distribute four clock phases. A significant percentage of the clock cycle is lost because CMOS clocks must be nonoverlapping, CCD clocks must be overlapping, and each of these must be synchronized with the others.




Fig. 18. Measured SNR and SNDR over a 9-MHz bandwidth.




Fig. 19. Measured integral and differential nonlinearity at an 18-MHz data rate.

At a given operating voltage, POSC power scales linearly with sampling rate since its circuits are strictly dynamic. At full speed, the POSC1 operates from 5, 4, and 3.3 V for analog CMOS, CCD, and digital CMOS, respectively, and consumes 324 mW. At a reduced speed of 10 MHz, voltages can be reduced to 4, 3.3, and 3.3 V, and power is reduced to 122 mW. Of the total POSC1 power, 65% is in CMOS modulator circuits, such as DDS described above, 20% is due to CCD clock drivers, which are digital inverters, and 15% is due to the digital decimator.

POSC2 measurements, listed in Table I, were done at a 30-MHz sampling rate on an input of approximately 13.3 MHz. The speed and power improvements of the POSC2 are close to those expected from a $2\times$ geometry reduction and a gate oxide scaling from 225 to 150 Å. The reduced accuracy of this device is due primarily to the straight scaling that was applied to all circuitry except the pad frame and is not an inevitable result of CCD circuit implementation in reduced geometry processes. Ideally, when CCD circuits are scaled, CCD well sizes and polysilicon capacitor values are adjusted so as to preserve their corresponding charge packet sizes.

V. CONCLUSION

A new architecture for oversampling A/D conversion, referred to as pipelined oversampling, has been presented. It is capable of improved speed over conventional $\Delta\Sigma$ techniques because it performs oversampling quantization spatially along a pipeline, rather than sequentially in time. The architecture is also well suited for processing presampled signals because it performs Nyquist-rate sampling. Two pipelined oversampling quantization algorithms are described, and methods for unraveling these algorithms into a pipelined structure are presented. Differences between pipelined and conventional oversampling are discussed. Like conventional multibit or MASH architectures, a pipelined converter does not have inherent linearity because it uses multiple feedback elements.

Two prototype pipelined oversampling converters have been demonstrated. The devices are implemented using CCD/CMOS circuits in standard double-poly CMOS processes. They demonstrate that high performance is achievable from CCD circuits without custom processing. A set of new CCD circuit techniques, based on a technique referred to as dynamic doubling sampling, are presented, which provide improved linearity and speed over existing techniques. This technique is used to implement circuits for wire transfer, charge generation, charge subtraction, D/A conversion, and charge sensing in the prototypes. Measured performance is presented for the prototypes. The first uses a 1.2- μm process and achieves 74-dB SNR over a 9-MHz bandwidth and 78-dB SFDR from Nyquist sampling at an 18-MHz output data rate.

ACKNOWLEDGMENT

The authors would like to thank J. Holtham and E. Morales for their contributions and Mentor Graphics Corp. for providing their electronic design software to support this work.

REFERENCES

- [1] R. T. Baird and T. S. Fiez, "A low oversampling ratio 14-b 500-kHz $\Delta\Sigma$ ADC with a self-calibrated multibit DAC," *IEEE J. Solid-State Circuits*, vol. 31, pp. 312–319, Mar. 1996.
- [2] T. Brooks, D. Robertson, D. Kelly, A. Del Muro, and S. Harston, "A cascaded Sigma-Delta pipeline A/D converter with 1.25 MHz signal bandwidth and 89 dB SNR," *IEEE J. Solid-State Circuits*, vol. 32, pp. 1896–1905, Dec. 1997.
- [3] E. T. King, A. Eshraghi, I. Galton, and T. Fiez, "A Nyquist-rate Delta-Sigma A/D converter," *IEEE J. Solid-State Circuits*, vol. 33, pp. 45–52, Jan. 1998.
- [4] T. Hayashi, Y. Inabe, K. Uchimura, and A. Iwata, "A multistage $\Delta\Sigma$ modulator without double integration loop," in *ISSCC Dig. Tech. Papers*, Feb. 1986, pp. 182–183.
- [5] D. F. Barbe, "Imaging devices using the charge-coupled concept," *Proc. IEEE*, vol. 63, pp. 38–67, Jan. 1975.
- [6] B. E. Burke, J. A. Gregory, M. W. Bautz, G. Y. Prigozhin, S. E. Kissel, B. B. Kosiki, A. H. Loomis, and D. Y. Young, "Soft-X-ray CCD imagers for AXAF," *IEEE Trans. Electron Devices*, vol. 44, pp. 1633–1642, Oct. 1997.
- [7] E. R. Fossum, "Wire transfer of charge packets using a CCD-BBD structure for charge-domain signal processing," *IEEE Trans. Electron Devices*, vol. 38, pp. 291–297, Feb. 1991.
- [8] C. H. Sequin, "Linearity of electrical charge injection into charge-coupled devices," *IEEE J. Solid-State Circuits*, vol. SC-10, pp. 81–92, Apr. 1975.
- [9] E. R. Fossum, "A linear and compact charge-coupled charge packet differencer/replicator," *IEEE Trans. Electron Devices*, vol. ED-31, pp. 1784–1789, Dec. 1984.

- [10] S. A. Paul and H.-S. Lee, "A 9-b charge-to-digital converter for integrated image sensors," *IEEE J. Solid-State Circuits*, vol. 31, pp. 1931–1938, Dec. 1996.



Susanne A. Paul received the B.S., M.S., and Ph.D. degrees from the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, in 1988, 1995, and 1999, respectively.

In 1988, she was with Digital Equipment Corp., where she participated in the design of the Alpha microprocessor. Since 1990, she has been with the MIT Lincoln Laboratory, Lexington, MA. Her work there has focused on CCD and CMOS visible and infrared electronic imagers, CCD-based A/D converters, and low-power analog circuit design.

In 1993, she received a National Science Foundation Fellowship.



Hae-Seung Lee (M'85–SM'92–F'96) received the B.S. and M.S. degrees in electronic engineering from Seoul National University, Seoul, Korea, in 1978 and 1980, respectively, and the Ph.D. degree in electrical engineering from the University of California, Berkeley, in 1984.

At the University of California, Berkeley, he developed self-calibration techniques for A/D converters. In 1980, he was a Member of Technical Staff in the Department of Mechanical Engineering, Korean Institute of Science and Technology, Seoul, where he was involved in the development of alternative energy sources. Since 1984, he has been with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, where he is now a Professor. Since 1985, he has been a Consultant to Analog Devices, Inc., Wilmington, MA. His research interests are in the areas of analog integrated circuits with an emphasis on analog-to-digital converters, operational amplifiers, and microsensor interface circuits.


Prof. Lee received the 1988 Presidential Young Investigator's Award. He has been a member of a number of technical program committees for various IEEE conferences, including the International Electron Devices Meeting, the International Solid-State Circuits Conference, the Custom Integrated Circuits Conference, and the IEEE Symposium on VLSI Circuits. From 1992 to 1994, he was an Associate Editor of the *IEEE JOURNAL OF SOLID-STATE CIRCUITS*.

John Goodrich, photograph and biography not available at the time of publication.



Titima F. Alailima received the B.S.E. and M.S.E. degrees in electrical engineering from the University of Pennsylvania, Philadelphia, in 1994 and 1996, respectively.

His graduate work focused on digital microelectronics and computer architecture. He joined MIT Lincoln Laboratory, Lexington, MA, in 1996, working in the Analog Device Technology group. There, he has focused on mixed-signal silicon integrated circuit design, primarily in the use of charge-coupled devices for signal processing.



Daniel D. Santiago was born in Newport, RI, in 1970. He received the B.S. degree from Northeastern University, Boston, MA, in 1993, where he is currently pursuing the M.S. degree.

From 1993 through 1995, he was a Member of Technical Staff with Raytheon Co., Portsmouth, RI. Since 1995, he has been an Assistant Staff Member at the MIT Lincoln Laboratory, Lexington, MA, working on analog-to-digital converters utilizing CCD technology.