A 4096 CELL SWITCHED CAPACITOR ANALOG WAVEFORM STORAGE INTEGRATED CIRCUIT

Stuart A. Kleinfelder

Physics Division
University of California
Lawrence Berkeley Laboratory
Berkeley, California, USA

Abstract

A full custom switched capacitor transient analog waveform storage and reconstruction integrated circuit containing 4096 sample and hold cells has been designed, fabricated and tested.[1] The Switched Capacitor Array (SCA) is organized as 16 parallel channels of 256 serially addressed samples per channel with multiplexed and buffered analog output. Signal sampling frequencies of up to 100 MHz have been accomplished, with a dynamic range, measured at 10 MHz, of at least 8000 to 1.

I. Introduction

Massively parallel data acquisition systems with as many as several hundred thousand individual waveform recording channels are required by researchers in particle physics and other fields. The general task to be accomplished is the decimation and temporary storage, followed by slower readout, of brief but high speed analog transients. For many applications, acquisition with sampling frequencies of up to 100 MHz or higher are needed. Usually, a trigger decision of a certain latency is made to judge the value of the stored signals. These signals are then either discarded or digitized, processed, and transferred to computer systems. Thus, these systems usually require both time delay via temporary storage and time stretching of the stored signal to match the bandwidth of the data processing system.

Traditional solutions to this problem have involved the use of flash analog to digital converters (FADC's) followed by digital memory, or use of charge coupled device (CCD) delay lines. The resolution of FADC's is limited to 6 to 8 bits and their dynamic range, even when using a non-linear transfer function, is often less than for CCD's. FADC systems have the substantial advantage of immediate digitization and therefore essentially zero conversion time. They may also allow very long record lengths (number of samples taken without pause). FADC systems, however, suffer from the worst density and highest cost and power consumption (> 1 watt per channel). The substantial size of monolithic FADC circuits has so far prevented incorporation of multiple channels or signal processing on a single integrated circuit chip. Charge coupled devices can have higher dynamic range ($10 - 12$ bits) and can be obtained with at least two channels per chip. They have so far been somewhat slower than FADC's, topping out at about 50 MHz. CCD's are often more difficult to use than FADC's due to the high clock drive they require. They also often suffer from temperature sensitivity and variable signal droop, with rate and duty cycle dependencies. It is presently more conceivable that CCD's can be integrated with other electronics appropriate to data acquisition applications than for FADC's. Both FADC's and CCD's are mature technologies, and only slow, incremental improvements can be expected.

The new switched capacitor approach may offer lower cost, lower power and higher density while providing fast sampling speeds, wider dynamic range and a high degree of flexibility. High channel count per chip permits low cost per channel and high density installations. The high performance, low cost CMOS process used, available from prototype foundries such as MOSIS, encourages experimentation and specialization. Finer feature size processes that include MOS and bipolar transistors are now becoming available, permitting even higher performance. These processes are suitable for inclusion of signal processing power which may ultimately include on-chip preamplification, signal shaping, analog to digital conversion and sophisticated addressing schemes including sparse or selective data readout.

[1]This work was supported in part by the U.S. Department of Energy, under contract No. DE-AC03-76SF00098

0018-9499/90/0600-1230$01.00 © 1990 IEEE

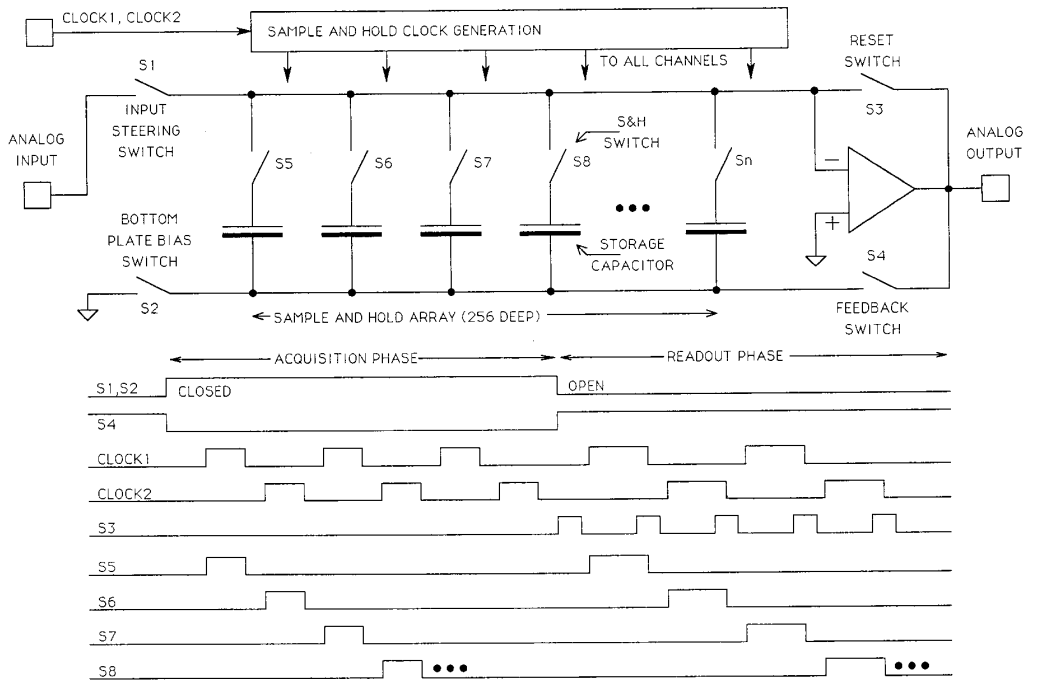

Figure 1: Simplified circuit and timing diagrams.

II. Circuit Description and Operation

The SCA (Switched Capacitor Array) integrated circuit uses a two micron double metal CMOS process with two layers of polysilicon forming the high density capacitors. Designed at L.B.L. using Berkeley software, this device was prototyped through the MOSIS service. It has 16 parallel analog signal inputs, 16 parallel analog outputs and 1 multiplexed and buffered analog output. Two digital clock inputs, four steering switch and reset control inputs, partially redundant, and four channel address inputs are used to operate the circuit. Figure 1 shows a simplified circuit diagram and timing information, Figure 2 a complete circuit diagram, and Figure 3 contains a die photograph of the fabricated integrated circuit.

The circuit contains 4096 sample and hold cells subdivided into 16 parallel channels of 256 cells per channel. Each of the 16 channels has a dedicated analog input which is steered onto a bus distributing the input signal to 256 sample and hold cells. Steering is accomplished by transmitting the input signal through a low on-resistance complementary switch (S1). Each sample and hold cell consists of a complementary CMOS transmission gate (complementary switch) and a 1.5 pF double polysilicon capacitor. The capacitors use a high quality silicon-oxide dielectric of 700 angstroms thickness. Complementary switches consist of an n-channel and a p-channel transistor whose sources and drains are paralleled. When complementary, rail to rail, clocks are applied to the gates of both transistors, transmission of analog signals from rail to rail can be accomplished. An externally supplied reference voltage is applied to the bottom plate of all capacitors through another set of steering switches (S2). The voltage stored on each sample and hold capacitor then corresponds to the difference between the input signal from a given channel at the time its sample and hold switch is opened and the applied reference voltage. The reference voltage can be adjusted to shift the baseline of the acquired signal for better level matching.

A shift register based clocking mechanism sequentially turns on and off sample and hold switches using a controllable break before make action. Break before make insures that no charge sharing between subsequently engaged capacitors can occur. Break before make can be controlled by varying the amount of time between non-overlapping clock phases. Normally, only a few nanoseconds are required. The 256 pairs of sample and hold clock signals are generated once and shared by all channels for lower power consumption and higher density. The sample and hold clock generation is designed to capture one sample per clock phase so two samples are captured per clock cycle, reducing external clock generation demands (100 MHz sample rates are attained with a 50 MHz clock speed). All clock and control input receivers are buffered internally for minimal external loading.

Before readout, the analog input and reference potential steering switches (S1 and S2) are turned off, disconnecting the signal bus and bottom capacitor plate from the input side of the circuit. The reconstruction operational amplifiers, one per channel, can then be connected to the

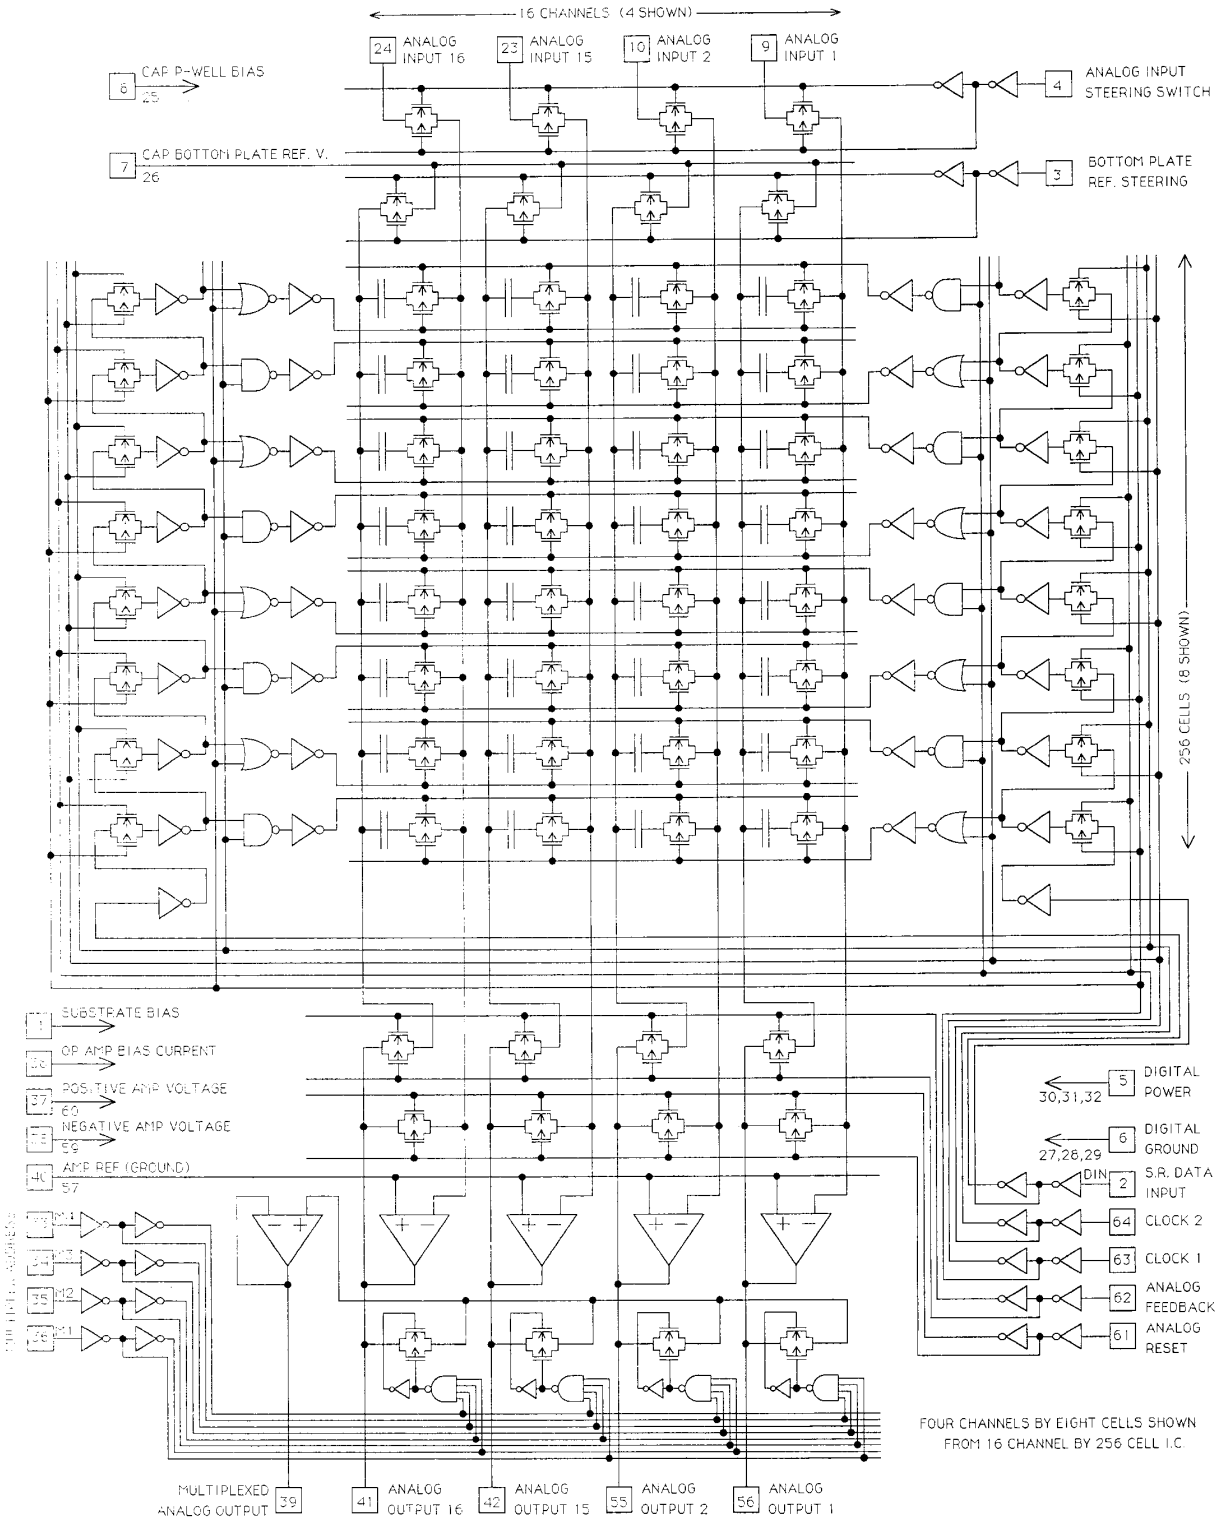

Figure 2: SCA analog transient recorder circuit diagram.

sample and hold arrays (via S4). These amplifiers are low power single stage differential folded cascode devices, selected for their low power, low noise, and large output range. The readout sequence then proceeds in much the same way as the acquisition phase; under control of the shift registers, each sample capacitor is placed, sequentially, in the feedback path of the reconstruction op-amp. The present chip implements an inverting output, but a non-inverting configuration is also possible. Once the amplifier outputs have settled, multiplexing of all 16 channel amplifier outputs through a single buffered analog output can be accomplished, selected via the four digital binary encoded channel multiplex inputs. This allows all data from one or many SCA chips to efficiently share a single high resolution analog to digital converter. Alternatively, each channel also has individual analog outputs available for use. After reconstruction and readout of each sample, stray charge on the sample capacitors and signal distribution busses are cleared by shorting each amplifiers inverting input nodes and output nodes (via S3). This prevents charge remaining on parasitic nodes from the previous samples from influencing the next read out sample. Reset of the sample capacitors also prevents any memory effect from influencing the next acquisition - this is only strictly necessary at the highest acquisition speeds.

III. Circuit Performance

The performance of the circuit is summarized in Table 1. With a 50 MHz sample rate, a 100 KHz readout rate, and using a $+/-3.5$ (or 0 to 7) volt supply, power consumption is a maximum of 10 mW per channel. Proper operation to 100 MHz sampling frequencies has been verified. The time constant for charging the sample and hold capacitors is 16 ns, yielding an analog bandwidth of about 10 MHz. The dc system gain with a 50 MHz acquisition rate is approximately -0.75, and at 10 MHz it is -0.95. Readout of all 4096 samples using the on-chip output multiplexing can be accomplished at an average rate of 1 MHz. The analog input range spans the entire power supply voltage used (maximum of 10 volts). The output operational amplifiers saturate within about 0.7 volts of each rail, and linearity degrades seriously within 1 volt of each rail. On a 7 volt supply, and within a $+/-2$ volt analog input range, reconstruction non-linearity (maximum deviation from straight line fit) remains within 3%, and within $+/-1$ volt, non-linearity is less than 1%. Using a quadratic fit, these numbers are several times better. Storage droop rate due to leakage on the storage capacitors is about 0.1 mV per ms at room temperature. Rising temperature to at least 50 degrees C causes no measurable baseline shift, but a 3 mV per degree C drop in gain at the extremes of the output voltage limits, worsening linearity somewhat (this is presumed, but not proven, to be due to worsening op-amp open loop gain at higher temperatures). At a 50 MHz acquisition rate, the cell to cell non-uniformity from all sources is less than 4 mV rms. At 10 MHz, the single cell output noise from all factors is less than 0.6 mV, and the dynamic range exceeds 8000:1, or 13 bits. Output offset is nominally 0 volts differential from the op-amp non-inverting input bias. Level shifting can be accomplished on both acquisition and readout phases, allowing the flexibility to match to existing external input and output side circuitry (ie, amplifiers and analog to digital converters).

Table 1: SCA Performance Summary

Number of channels per chip	16
Number of storage cells per channel	256
Total storage cells per chip	4096
Die size	$6.8 * 4.6mm$
Power consumption per channel at 7V	$10mW$
Input capacitance	$7.5pF$
Input voltage range	$> 5V$
Output voltage range	$> 5V$
Non-linearity within ± 1 volt range	1 %
Non-linearity within ± 2 volt range	3 %
Non-uniformity in baseline (50 MHz)	$< 4mV$ rms
Non-uniformity at $\pm 2V$ (50 MHz)	$< 4mV$ rms
Stored voltage droop rate	$0.1mV/ms$
Sample and hold time constant	$16ns$
DC Gain at 10 MHz sample frequency	-0.95
DC Gain at 50 MHz sample frequency	-0.75
Output noise, 10 MHz sample freq.	$< 0.6mV$
Dynamic range (signal to noise)	$> 8000:1$
Maximum sample rate	$\geq 100MHz$

IV. Limits on Circuit Performance

The SCA chip uses complementary switches exclusively in the analog signal path. This maximizes the usable input voltage swing to increase dynamic range. SCA can store signals from rail to rail (7 to 10 Volts). The on-chip readout operational amplifier then becomes the limit to usable range. It was chosen partly because it has a wide output range - within one threshold from each rail. Reduced linearity near these extremes tends to limit the usable range to within 1 volt from each rail. Thus, from 5 to 8 volts of total useful signal range can be achieved. The noise performance of the operational amplifiers limits the minimum resolvable signal, and therefore the also dynamic range. The measured noise value on any given single cell when operated with a 10 MHz acquisition rate is less than 0.6 mV RMS. Thus, the total dynamic range (single cell) exceeds 8000:1, or 13 bits.

Non-uniformities in the sample and hold elements also limit dynamic range unless corrections are applied. The double polysilicon capacitors, although of excellent quality, are not likely to be uniform to better than 99.9%, or about 10 bits, but the SCA readout scheme is not sensitive to capacitor values to first order. Non-uniformity of the sample and hold switches can be important, as the charge

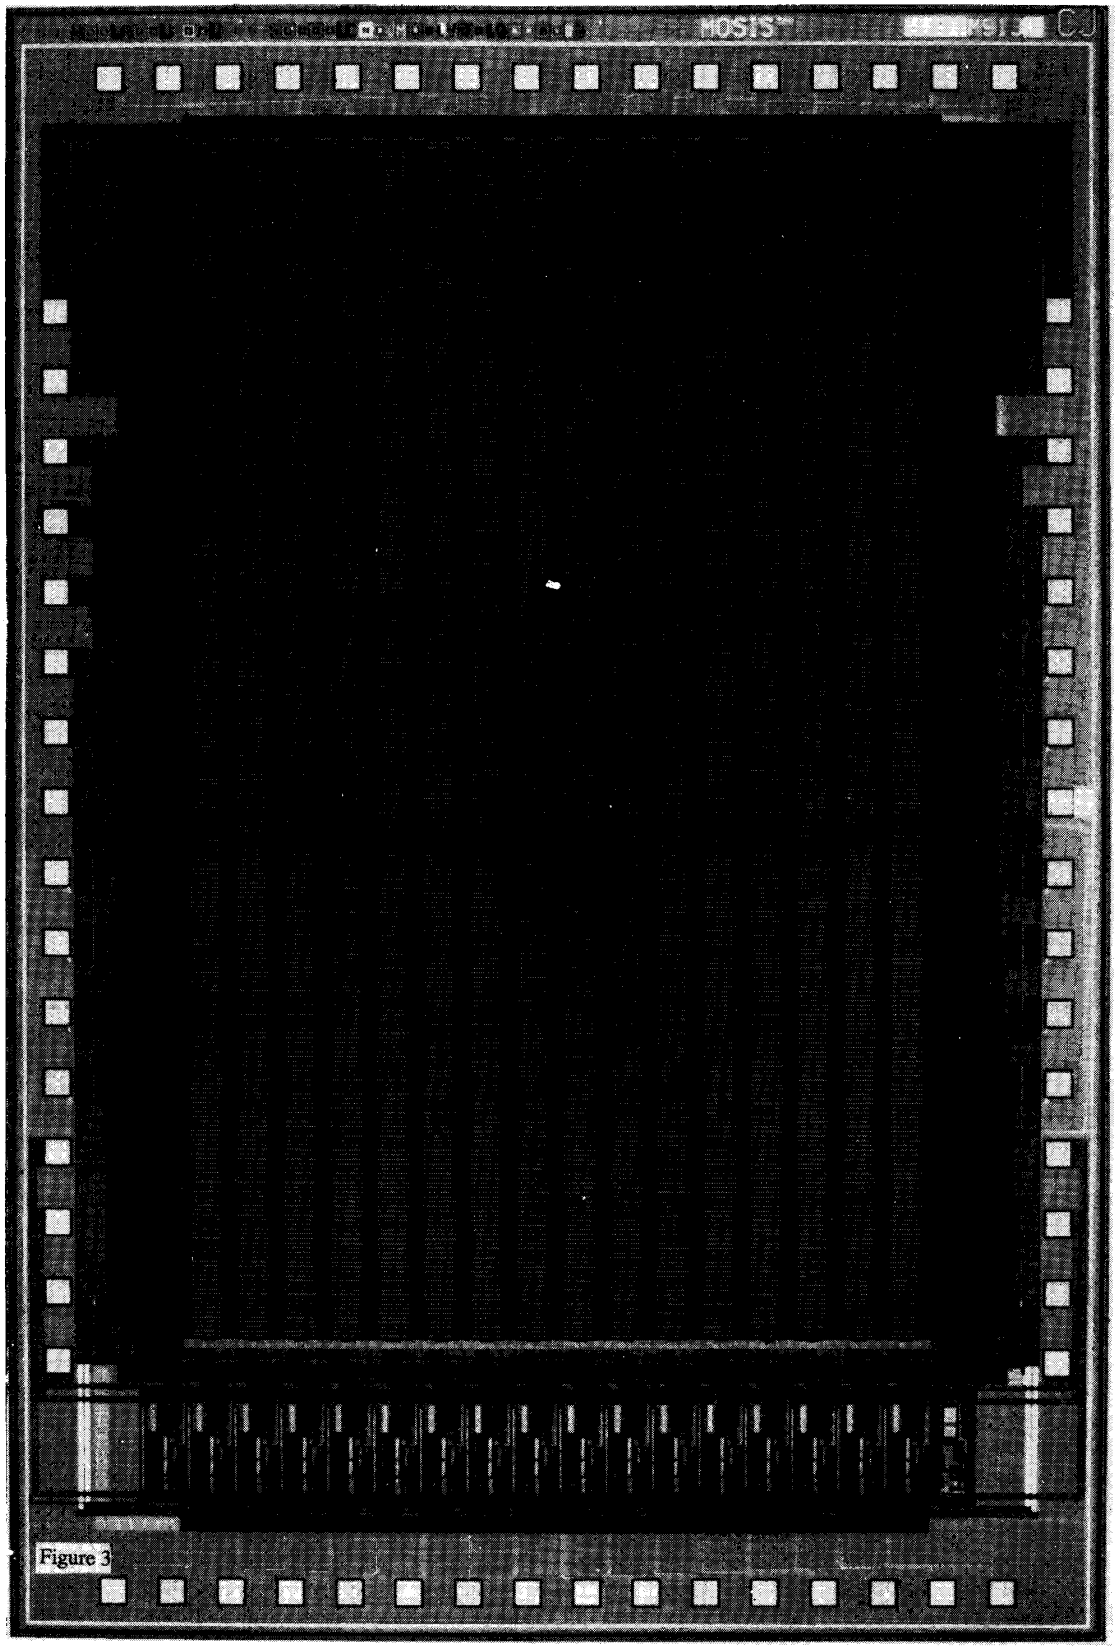

Figure 3

injection from these switches (due to gate to drain capacitance) will then vary. These switches are complimentary, and therefore the charge injection is canceled to first order. However, this cancellation is only approximate, and variations still exist. In addition, the rise and fall times of the clocks are very fast, and slight variations in clock skews (imperfect alignment of the complementary clocks) can also lead to imperfect and variable cancellation. Examination of the circuit diagram of the SCA chip reveals a potential asymmetry (in practice, if not in principle) in clock driving circuitry. The clock circuitry is split into two halves to increase circuit efficiency and and allow a finer cell pitch. These halves are on opposite sides of the chip, and facing in opposite directions, so mask misalignment and process variations accumulated during fabrication could induce extra random or systematic variations in clock edge delays and speed. This was indeed found to be the case, and a small odd/even asymmetry was measured. This means that odd samples had a small but systematically different offset value when compared to even samples. The total cell to cell variability from all sources, random and systematic, measured during operation with a 50 MHz acquisition rate was less than 4 mV RMS. Thus the uncorrected dynamic range, all cells, well exceeds 1000:1, or 10 bits. This situation can be improved several times by a redesign aimed at better clock matching. Since the present systematics dominate non-uniformity, elimination of any remaining sources of systematic mis-matches may increase the uncorrected dynamic range by about a factor of 4, to about 4000:1, or 12 bits.

During acquisition, several parasitic resistors and capacitors limit performance. The bus distributing the input signal to the sample and hold capacitors has a certain capacitance due to the metal distribution line and the sum of the switch transistor source capacitances. This capacitance must be driven by the signal input device (normally an amplifier or shaper). In the present chip this load is a modest 7.5 pF, much lower than most FADC inputs. Combined with the resistance of the input steering switches, this generates an RC charging time constant which may limit the analog bandwidth of the circuit. However, this time constant is presently under 2 ns., so this is not a great limitation. More serious is the resistance of the sample and hold switch when combined with the sample capacitor. The tested circuit used minimum size transistors for the sample and hold switches, as the first application of this device requires only a 10 MHz sample rate. The measured RC time constant is 16 ns., yielding about a 10 MHz analog bandwidth. If operated at very high sample rates, this time constant causes the gain of the system to drop uniformly. To increase the bandwidth, either the capacitor may be made smaller, or the sample and hold switch made larger for lower on resistance, or both. At the time of writing, a non-inverting, higher analog bandwidth version is being manufactured for evaluation. The increased bandwidth is expected to come with the penalties of decreased cell to cell uniformity and higher input capacitance. The present minimum size switches are of benefit in lower speed applications because they minimize the absolute value of charge injected onto the sample and hold capacitors when they open, thus minimizing the offset errors (random and systematic) encountered. In addition, they keep the parasitic input capacitance minimized.

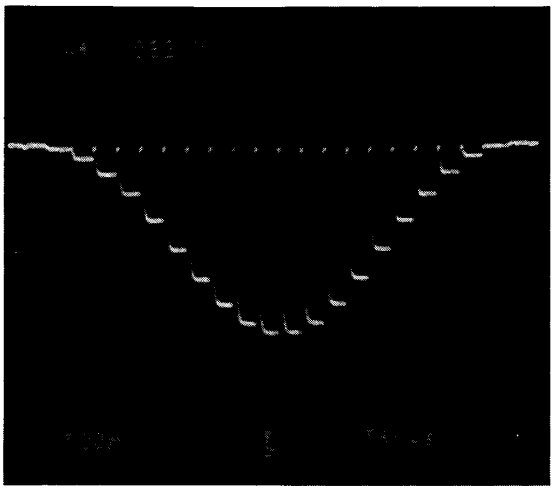

Figure 4: SCA output reconstruction of a sinusoidal pulse acquired at 50 MHz. Only the first 24 samples of a 256 sample long record is seen.

During readout, many of these same parasitics are important, as well as others. The parasitic capacitance between the input distribution bus and the bottom capacitor plate can cause a uniform signal attenuation during reconstruction. This effect is found to be small, however, as the value is only 0.03 pF per capacitor. The total attenuation due to all sources is less than 5% at a 10 MHz sample rate. Also important is the performance of the on-chip readout operational amplifiers. These amplifiers must drive the parasitic bottom plate capacitance of all sample capacitors - substantial at about 50 pF. This limits the speed of settling of the reconstructed output signal. Slower settling on this node is mitigated by the very much faster buffered output multiplexing. One sample in all channels settles in parallel internally, followed by fast switching between channels during readout. In this way readout performance is enhanced considerably while increasing readout convenience and lowering external digitization circuitry overhead.

The shift register clocking mechanism can also affect performance. Coupling between the clock lines and the analog nodes may be a concern, but only one pair of differential sample and hold clocks (from the 256 pairs) are active at a time and their proximity to analog nodes has been minimized. The internal clocking is efficiently designed, and has been demonstrated to perform correctly at a minimum of 100 MHz. Simulations suggest the circuit may work at

more than twice as high a rate. The rise and fall times of the clocks to the sample and hold switches were measured using a micro-probe and found to be about 1 ns each. Low jitter in the sample and hold clocks is important at high sampling frequencies, as it directly limits sampling accuracy of high bandwidth signals. Timing jitter was found to be less than 0.1 ns rms.

V. Future Implementations

To increase intrinsic dynamic range further, it may be necessary to use fully differential techniques. This would further reduce systematic errors, in principle, to zero, though at a cost in channel density. With a similar cost in density, the familiar bi-range approach can be used. Linearity can be improved by increasing the gain of the readout amplifier and reducing parasitic capacitances. Analog bandwidth can be increased by the techniques discussed in the previous section. It is important to note that most analog performance parameters are sensitive to the number of samples per channel that the circuit includes due to increased parasitics associated with each sample cell.

Alternative clock generation schemes are being considered with the aim of increased uniformity and flexibility. Within the serially addressed paradigm, a system reset function which can quickly abort an acquisition or readout at any point would be useful. While serial addressing may always have its place, random addressing or a combination of serial and random access to capacitors would enhance functionality. The system could then skip to particular samples of interest during readout, or skip over saved samples during acquisition. The capacity for simultaneous acquisition and readout, in a randomly addressed fashion would allow extremely flexible operation. Such a system would necessarily be more complicated and, likely, of somewhat lower density. Planning would also be needed to prevent an explosion in the number of control and clock lines. That the switched capacitor approach is capable of such flexibility, however, is perhaps its greatest single advantage.

Acknowledgements

The help and advice received from Professor P. R. Gray of U.C. Berkeley and D. R. Nygren, F. Kirsten, and H. Wieman of Lawrence Berkeley Laboratory is appreciated. The efforts during tests of the SCA by K. Lee of L.B.L. and A. Rudge of U.C. Davis are particularly valued. Thanks go to the MOSIS crew for prototype fabrication.

References

Reference [1] reports on an earlier implementation of this circuit first fabricated in February of 1987. Reference [5] reports on an earlier NMOS device with similar intent but a different implementation. Reference [6] is on a small experimental device (8 cells) using a more related technique. Papers [2-4] and [7-9] are papers on transistor and capacitor matching issues, charge injection, and general circuit design and implementation.

1. S. A. Kleinfelder, "Development of a Switched Capacitor Based Multi-Channel Transient Waveform Recording Integrated Circuit," IEEE Trans. Nucl. Sci., NS-35, 151 (1988).

2. J. McCreary, "Matching Properties and Voltage and Temperature Dependence of MOS Capacitors," IEEE Journal of Solid-State Circuits, SC-16, no 6, (1981).

3. T. Choi, R. Kaneshiro, R. W. Broderson, and P. R. Gray, "High-Frequency CMOS Switched Capacitor Filters for Telecommunications Applications," IEEE Journal of Solid-State Circuits, SC-18, no 6, (1983).

4. K. Matsui, et al, "CMOS Video Filters Using Switched Capacitor 14 MHz Circuits," IEEE Journal of Solid-State Circuits, SC-20, no 6, (1985).

5. D. Freytag and J. Walker, "Performance Report for Stanford/SLAC Microstore Analog Memory Unit," IEEE Trans. Nucl. Sci., NS-32, 622 (1985).

6. W. Buttler, et al, "Design and Performance of a 10 MHz CMOS Analog Pipeline," Nuclear Instruments and Methods, A277 (1989) 217-221.

7. W. Wilson, et al, "Measurement and Modeling of Charge Feedthrough in n-Channel MOS Analog Switches," IEEE Journal of Solid-State Circuits, SC-20, no 6, (1985).

8. S. A. Kleinfelder, "Custom MOS VLSI Design: Tools and Methods," IEEE Trans. Nucl. Sci., NS-36, 484 (1989).

9. P. R. Gray and R. G. Meyer, Analysis and Design of Analog Integrated Circuits, New York: Wiley, (1984).