

Approaches and analysis for on-focal-plane analog-to-digital conversion

Bedabrata Pain and Eric R. Fossum
Center for Space Microelectronics Technology
Jet Propulsion Laboratory, California Institute of Technology
Pasadena, California 91109 USA

ABSTRACT

This paper presents approaches for on-focal-plane analog-to-digital conversion (ADC). Common approaches and architectures for ADC and their utility for on-focal-plane integration are discussed. Candidate approaches are analyzed with respect to required amplifier gain, bandwidth, capacitance matching, noise and offsets as a function of ADC resolution. A column-parallel ADC architecture appears to be an effective compromise of chip area, power, circuit speed and ADC resolution. The discussion is valid for both infrared focal-plane arrays (FPAs) and visible image sensors.

1. INTRODUCTION

In the next generation of scientific and defense focal-plane arrays (FPAs) there will be an increasing emphasis on improving overall system performance and on increased integration to miniaturize camera size. One way to meet both these objectives is to integrate analog-to-digital converters (ADCs) on the focal-plane.




Fig. 1 System signal chain with off-chip ADC (above) and on-chip ADC (below)

The system signal chains for two cases, one with an off-chip ADC and the other with on-focal-plane ADC are shown in fig. 1. The incorporation of an on-focal-plane ADC will lead to a reduction of the total number of integrated circuits in the system, and in turn will lead to a number of system improvements. First, increased electronics integration has proven itself over the past 30 years to lead to an improvement in system performance, an increase in system reliability, and a decrease in system cost. Increased integration in spacecraft instruments also leads to reduction of total mass, volume, and system power. As a result the quantity, mass, and volume of power supplies are greatly reduced, and power system reliability is improved. Another high leverage advantage is that as the total electronics volume is reduced, radiation shielding mass can be greatly reduced. Increased integration also leads to a reduction in system complexity, since the reduction in the IC count reduces the number of interconnects required. This speeds system design time and reduces

the design error rate. Further, with a digital output, interfacing the FPA is also simplified, with a likely concomitant reduction in brassboard and flight hardware development time since considerable time is expended in "cleaning out" the signal chain.

An FPA system with on-focal-plane ADC is also expected to exhibit superior noise performance. Typically, the noise performance of state-of-the-art focal-plane arrays is limited by the performance of their signal chains. This is due to the inevitable introduction of unwanted and possibly unavoidable noise through cross-talk, clock pickup, power supply noise, electromagnetic interference (EMI), and other mechanisms. Since the serial data rate in the signal chain is the highest rate in the imaging system, white noise is introduced with a maximum bandwidth. Clock noise and other capacitively coupled pickup is also known to increase with increasing data rates. On-chip ADC-per-column technology operates at a significantly

lower bandwidth ameliorating these effects. Since no off-chip analog cabling is required, pick up and vibration sensitivity is eliminated. More fundamentally, multiple sampling or oversampling of the detector signal can be much more effectively performed on the focal-plane compared to off-chip. Thus, on-chip ADC both eliminates mechanisms for the introduction of noise, as well as permits increased SNR through oversampling techniques.

Despite intuition to the contrary, it can be shown that on-focal-plane ADC can lead to a reduction in FPA power dissipation. While the ADC will dissipate extra power on the focal-plane, the need for a high bandwidth amplifier to drive high fidelity analog signal off chip (through high capacitance loads) is eliminated, resulting in a net power savings.

Once the data has been converted to the digital domain, digital signal processing (DSP) can be performed on-chip as a further stage in increased integration. The on-chip DSP can be used for autonomous sensor control such as exposure control or for control of windowed region-of-interest readout. Image compression can also, in principle, be achieved on-chip to reduce off-chip drive requirements. However, it is the view of the authors that once one is "outside of the feedback loop" to the detector array, increased integration on-chip is less justifiable and a separate DSP application specific integrated circuit (ASIC) may be a more attractive approach.

In this paper ADC architectures and approaches suitable for focal-plane implementation are presented. The achievable conversion resolution as a function of circuit parameters, speed, area, and power for a given CMOS implementation topology are derived. Approaches to column-parallel implementation are discussed.

2. ON-FOCAL-PLANE ADC ARCHITECTURES

In conventional focal-plane readout systems shown schematically in fig. 1, a single off-focal-plane ADC serves the entire array. For on-focal-plane ADC implementation the most obvious architecture will be to integrate a single ADC monolithically along with the readout electronics. Although conceptually simple, a serial on-focal-plane ADC architecture




Fig. 2 Two architectures for infrared readout with on-focal-plane ADC: a serial architecture (left), and a parallel architecture (right).

is not an optimal one. First, incorporation of high resolution ADCs on the focal-plane itself poses a challenge, since compared to a stand-alone ADC, there is much less silicon area available. Secondly, analog circuits in a serial ADC are required to operate with the highest bandwidth of all focal-plane components, since the conversion rate is the same as the pixel data rate. Typical pixel data rate in scientific applications is around 100 kHz, while for defense applications and even in certain scientific applications, data rates in excess of 1 MHz are required. Thus, reliable operation of CMOS circuits pose a concern, especially for applications requiring high data rate. The problem is compounded by the fact that scientific applications routinely require resolutions greater than 16-bits. This can only be achieved with oversampling techniques, so

that the ADC clock rate is even higher. Finally, increased focal-plane power dissipation due to high speed operation of several analog circuits (as opposed to a single driver amplifier in conventional focal-plane readouts) is another concern.

If it is possible to incorporate an ADC per readout pixel, as shown in fig. 2, a massively parallel ADC architecture results. However, since only a relatively small area is available for most applications (typical pixel size being about $30 \mu\text{m}^2$), conventional ADC approaches may not be feasible with this architecture. A few unconventional ADC approaches such as voltage-to-frequency converters [1] have been investigated. A voltage-to-frequency converter operates by generating a pulse train whose repetition frequency depends on the analog input voltage. The feasibility of this approach is limited by the ability to read the pulse train of varying repetition frequency from each pixel. For extremely low flux levels, a photon counting approach to in-pixel ADC is being investigated [2]. In this approach, A/D conversion can be achieved by counting individual photoelectrons with high-gain amplifiers operating with sub-electron input-referred noise.

Speed limitations in a serial architecture and area limitations in parallel architecture abbreviate their utility for on-focal-plane applications. On the other hand, use of a semi-parallel architecture can be expected to preserve the advantages and mitigate the adverse consequences of both the architectures. A semi-parallel architecture featuring an ADC for every column of the readout, shown in fig. 3, affords virtually unlimited chip area in one dimension, and tight but feasible design space in the other dimension. Such tall, skinny ADCs would operate in parallel on a row of image data at a time. The




Fig. 3 A column-parallel architecture for on-focal-plane ADC

conversion rate is the row read out rate, and is 100-1000x slower than the serial pixel data rate. For example, for a serial pixel data rate of 50,000 pixels per second typically employed in a scientific CCD, the corresponding ADC data rate for an ADC per column architecture would be just 50 pixels per second for a 1024x1024 array size. In the case of a line array, an ADC per column converges to the limiting case of an ADC per pixel. It is also possible to envision an ADC per n columns, where n is a small number, perhaps 32. This derivative architecture requires sub-multiplexing of the n -columns to the ADC. It affords more design area in the transversal direction at the expense of increased data rate.

A semi-parallel architecture also affords lower power operation. This is due to the fact that unlike digital circuits, power dissipation in analog circuits depends super-linearly on the operating frequency. As a result the total power dissipation for a semi-parallel architecture can be lower than for a serial architecture for the same pixel data rate. Additionally, the semi-parallel architecture allows the use of a "shared circuits", i.e. a single circuit block (e.g. reference generator) can serve the entire ADC bank. This can lead to considerable power and area saving. Since the conversion rate is much lower than that in a serial architecture, the

analog circuits may be biased in sub-threshold, affording further reduction in power. The highest performance on-focal-plane ADC has been reported using a column-parallel architecture [3, 4]. One concern, however, is the matching of the ADC characteristics between the columns. A lack of matching will generate a fixed-pattern noise in the output data.

3. ALGORITHMS FOR ON-FOCAL-PLANE ADC

A number of ADC algorithms are available for use in focal-plane applications. These conversion methods differ from each other in terms of operating speed, power consumption, achievable accuracy, and chip area. An important difference between on-focal-plane ADC and a single-chip monolithic ADC is that, unlike the latter, an on-focal-plane ADC must occupy a relatively small chip area. The real estate becomes an even more serious concern for column-parallel approaches. Due to unavailability of a large chip area, focal-plane ADCs cannot usually take advantage of elaborate trimming techniques for resolution enhancement. Thus, in choosing an ADC architecture, the immunity of the ADC performance to circuit

parameter mismatch is an important issue.

All focal-plane implementations require low power operation, maximum overall power dissipation in all the ADCs being limited to 1 - 20 mW. The required resolution and conversion rates vary widely depending upon applications. The conversion rate depends on the array size, the integration time, and the choice of ADC architecture, and is usually in the range of 1 kHz to 1 MHz. Scientific IR imagers usually demand high resolution (> 16 bits), but several other applications require only 8-10 bit accuracy. Thus there is a wide range of operating requirements, with the conversion rate requirement varying from 1 kHz to 1 MHz, and the bit resolution requirement varying from 8 to over 16 bits. There is no single ADC algorithm that optimally meets these widely varying requirements.

Approaches to the design of the ADC include flash ADCs, successive-approximation ADCs, single/dual slope ADCs, and oversampled ADCs. Flash ADCs achieve 8-10 bits of resolution at very high speed, but require a large chip area and dissipate unacceptably high power levels. Therefore, they are unusable for focal-plane integration. Successive-approximation ADCs achieve higher resolution at medium speeds. Power dissipation is medium. These types of ADCs are good candidates for an ADC per n columns since they require large chip area in order to achieve good component matching. Single and dual slope ADCs are architecturally simple circuits and can achieve high resolution at low speeds with minimal power dissipation. The robustness of these ADCs is judged to be too low to be suitable for a cosmic-ray-ridden space environment but may be suitable for many low resolution applications. Oversampled ADCs are an emerging approach to very high resolution (e.g. 18-20 bits) conversion at medium speeds and low power dissipation. Oversampled ADCs offer the possibility for SNR enhancement as discussed below. They also exhibit very high robustness and may therefore be extremely suitable for integration on the focal-plane. The characteristics of these algorithms and their feasibility for focal-plane use are summarized below in table 1.

TECHNIQUE	# BITS	PWR	AREA	ROBUSTNESS	SPEED	GOOD FOR FPA
Flash	8-10	High	High	Medium	High	No
Succ. App/Cyclic	10-12	Low	Medium	Low	Medium	Yes
Single/Dual-Slope	12	Low	Low	Low	Low	Maybe
Σ - Δ	18-20	Low	Medium	High	Medium	Yes

Table 1 Relevant characteristics of ADC algorithms for on-focal-plane applications

Cyclic ADC

For medium speed (data conversion rate < 100 kHz) operation, a cyclic, a successive-approximation (S-A) or a Σ - Δ ADC can be used. The cyclic algorithm is based on the multiply by 2 method of decimal-to-binary conversion and is given by [5]:

$$V(i+1) = 2 V(i) + (-1)^{b_i} V_{\text{ref}}$$

$$b_i = 1 \text{ for } V(i) \geq 0; \quad = 0 \text{ for } V(i) < 0; \quad i = 1, 2, \dots, n$$

$V(1) = V_{\text{in}}$; and V_{ref} is the reference voltage. b_i is the i -th bit in the binary representation of the analog signal. The range of the input voltage is given by:

$$- V_{\text{ref}} \leq V_{\text{in}} \leq V_{\text{ref}}$$

Successive-approximation ADC

The successive-approximation is an essentially "ranging" algorithm. During each conversion step, the algorithm tries to determine the upper and the lower bound within which the input voltage lies. By successively shrinking the bounds, the analog voltage is approximated to within a small error. The ranging can be done in several ways. A successive-approximation (S-A) algorithm [6] that is compatible with CMOS implementation is given by:

$$V(i+1) = V(i) - b_i \frac{V_{\text{ref}}}{2^i}$$

$$b_i = 1 \text{ for } V(i) \geq \frac{V_{\text{ref}}}{2^i}; \quad = 0 \text{ for } V(i) < \frac{V_{\text{ref}}}{2^i}; \quad i = 1, 2, \dots, n$$

$V(1) = V_{\text{in}}$; and V_{ref} is the reference voltage. The range of input voltage is given by:

$$0 \leq V_{\text{in}} \leq V_{\text{ref}}$$

In this algorithm, the residual voltage decimates to approximately $\frac{V_{\text{ref}}}{2^n}$ after n conversion steps. As a result, the algorithm can be susceptible to circuit noise, offset, and non-idealities more than a cyclic ADC.

Σ - Δ ADC

A moderate to low speed ADC is the oversampled sigma-delta (Σ - Δ) ADC. It combines sampling at rates well above the pixel readout rate with negative feedback to exchange resolution in time with that in amplitude [7]. The algorithm for the most basic (first order) Σ - Δ ADC is given by:

$$V(i+1) = V(i) + \alpha(V_{\text{in}} - b_i V_{\text{ref}})$$

$$b_i = 1 \text{ for } [V(i) + V_{\text{in}}] \geq 0; \quad = 0 \text{ for } [V(i) + V_{\text{in}}] < 0; \quad i = 1, 2, \dots, n$$

$V(1) = V_{\text{in}}$; and V_{ref} is the reference voltage. The range of the input voltage is given by:

$$-V_{\text{ref}} \leq V_{\text{in}} \leq V_{\text{ref}}$$

The repeated integration and feedback force the average value of the output to track the average value of the input. As a result of the feedback and integration, any uncorrelated variations (i.e. circuit noise) are suppressed. By the same token, the algorithm is also tolerant to component mismatches and non-idealities as well as circuit offsets.

Incremental Σ - Δ ADC

A first order Σ - Δ ADC requires 2^n cycles to perform a n -bit analog-to-digital conversion. A higher order Σ - Δ , built by incorporating additional error integrator loops, can speed up the conversion. However, the addition of more than one integrators in the feedback tends to cause stability problems. This problem can be overcome by cascading two first order stages resulting in an incremental Σ - Δ ADC topology [8]. Since it consists only of cascaded first order stages, it is immune to loop instabilities. At the same time, compared to a first order topology, it vastly improves the conversion speed. The number of cycles (m) needed for a k -th order Σ - Δ ADC to achieve n -bit resolution is given by:

$$m = \left(k! 2^n \right)^{\frac{1}{k}}$$

The algorithm for a second order incremental ADC (for an integrator gain of unity) is given by:

$$V(i+1) = \frac{i(i+1)}{2} V_{\text{in}} - \sum_{k=1}^i a_k (i+1-k) V_{\text{ref}} - \sum_{k=2}^{i+1} b_k V_{\text{ref}}$$

$$2D = \sum_{k=1}^i a_k (i+1-k) + \sum_{k=2}^{i+1} b_k; \quad i = 1, 2, \dots, n$$

V is the output of the second integrator, and D is the digital word formed by appropriately weighing the digital outputs a_i and b_i , that are generated respectively by the first and the second Σ - Δ modulator stages.

Single and dual-slope ADC

Apart from a first Σ - Δ ADC, other slow speed ADCs include single or dual-slope conversion. A single-slope ADC consists of a ramp generator, a counter and a comparator. One input to the comparator consists of a sampled-and-held analog

voltage to be converted and the other input consists of the ramp generator output. The ramp generator is turned on synchronously with the counter. The output of the comparator is "0", until such time when the ramp output exceeds the sampled and held input voltage. The output of the comparator is used to latch the counter value, and is an equivalent digital representation of the input voltage. For a n -bit data conversion, 2^n counter cycles are required. A dual-slope ADC [9] can be designed to make the data conversion independent of the ramp inaccuracies and comparator offset voltage. However, it requires integration of the analog voltage and is therefore susceptible to the $1/f$ noise. This limits its use for focal-plane applications.

4. ANALYSIS OF MODERATE-SPEED ADC RESOLUTION

A building block common to all the moderate speed ADC schemes is an integrator. For cyclic ADC, the integrator is also required to have a gain of 2. In most cases, the accuracy of the integrator limits the resolution achieved by an ADC. The inaccuracy introduced by the integrator depends on the circuit implementation used, and on the clocking speed.




Fig. 4 Schematic diagram of a switched-capacitor integrator. v_n is the input-referred noise voltage generator of the amplifier.

Switched-capacitor integrator

In voltage mode CMOS implementation, the most commonly used integrator is a switched capacitor integrator, shown in fig. 4. In this circuit, during the clock phase ϕ_1 , the oncoming signal is sampled on the sampling capacitor C_s . During the ϕ_2 clock phase, the feedback loop is activated, and the charge on C_s is dumped on the integration capacitor C_i . The operation of the during phase ϕ_2 resembles that of a commonly used IR readout circuit called a charge transimpedance amplifier (CTIA). Assuming that the output is sampled during phase ϕ_1 , and that the gain of the op amp A_{vo} is extremely large, the recursion relation between the output and the input is given by:

$$V_{out}(n) = \frac{C_s}{C_i} V_{in}(n-1) + V_{out}(n-1)$$

The integration gain is given by $A_{go} = C_s/C_i$, and can be altered simply by changing the capacitance ratio. The switched-capacitor integrator is subject to several non-idealities that tend to affect the integrator gain [10]. These effects arise from mismatch of capacitors, finite gain and finite bandwidth of the op amp, noise in the op amp, and offsets associated with charge feedthrough during switching and the op amp. To a first order, the effect of these non-idealities is to modify the integrator gain, and has been widely studied. It can be shown that when the non-idealities are small enough, the integration gain can be expressed as [10]:

$$A_g = A_{go}(1 - \epsilon_g); \text{ where } \epsilon_g = \epsilon_c + \epsilon_A + \epsilon_w$$

ϵ_c is the error fraction of capacitance mismatch, ϵ_A is the error introduced due to the op amp gain A_{vo} , and ϵ_w is the error due to finite bandwidth effects. ϵ_A can be further expressed as:

$$\epsilon_A = \frac{1}{A_{vo}} (3 - \beta); \quad \beta = \frac{C_i}{C_i + C_s};$$

Assuming that the settling behavior of the op amp is governed by a single pole f_o , which is required by all op amps for reasons of stability, the error ϵ_w is given by [11]:

$$\epsilon_w = (1 - \beta) \exp\left(-\beta\pi \frac{f_u}{f_c}\right)$$

$f_u = A_{vo} f_o$ is the unity gain frequency and f_c is the clock frequency.

Modifying the analog-to-digital conversion algorithms presented in the previous section to include the effect of the integration gain variation and the offsets in the op amp and the comparator, and due to switch feedthrough, the error in a n-bit conversion can be determined for Σ - Δ ADC, cyclic ADC and successive-approximation ADC, for standard switched-capacitor implementation. Expressed in units of the number of least significant bits (LSB), the conversion error for a Σ - Δ ADC is given by [8]:

$$\epsilon_{\Delta} = 0.5 m \alpha_g; \quad \text{where } m \text{ is the number of conversion cycles needed to achieve } n\text{-bit resolution.}$$

The conversion errors for a cyclic ADC (ϵ_{cyc}) and a S-A ADC (ϵ_{sap}) can be derived using the approach detailed in reference 12, and are given by:

$$\epsilon_{\text{cyc}} = 0.5 \left[(1 - \alpha_g)^n \left(2^n \frac{\alpha_g}{1 - 2\alpha_g} + 1 \right) - \frac{(1 - \alpha_g)}{(1 - 2\alpha_g)} \right]$$

$$\epsilon_{\text{sap}} = \frac{\alpha_g}{1 - 2\alpha_g} \left[1 + 2^n (1 - \alpha_g)^n \right]$$

If the deterministic offsets (due to switch feedthrough, mismatches in the op amp and the comparator) are not to generate any missing code, the conversion error due to the offsets needs to be less than 0.5 LSB. This imposes a maximum limit on the offset voltage that can be tolerated by each ADC scheme. This limit can be expressed as follows for a Σ - Δ , a cyclic, and a S-A ADC respectively:

$$V_{\text{os}\Delta} \leq \frac{V_{\text{ref}}}{m(m+3)}$$

$$V_{\text{os}\text{cyc}} \leq \frac{V_{\text{ref}}}{2^n - 1}$$

$$V_{\text{os}\text{sap}} \leq \frac{V_{\text{ref}}}{n 2^n}$$

For a Σ - Δ or a cyclic ADC, the maximum allowable offset voltage for less than 1 LSB differential non-linearity is around $0.5V_{\text{LSB}}$. However, the successive-approximation scheme is most affected by the presence of offsets. The maximum allowable offset for this scheme can be as low as $0.1 V_{\text{LSB}}$ for a 13 bit or higher resolution. For the other two schemes, a 10 bit ADC, with $V_{\text{ref}} = 2\text{V}$, (typical voltage levels encountered for low power operation), requires the offset voltage V_{os} to be less than 1 mV. Thus, circuit offset is an important mechanism for limiting the accuracy of an ADC system. Accurate switch-feedthrough reduction and comparator and op amp auto-zeroing are required in order to reduce offsets in CMOS switched capacitor circuits.

Noise

The major sources of noise in the switched capacitor integrator are the reset noise on the capacitors and the white noise from the op amp. The 1/f noise can usually be neglected since the op amp is usually clocked at the highest possible speed. During each integration cycle, the r.m.s. noise added is:

$$\langle v_n^2 \rangle = \frac{kT(C_i^2 + C_s^2)}{C_i^2 C_s} + \frac{kT(C_i + C_s)}{\left(C_L + \frac{C_i C_s}{C_i + C_s} \right) C_i} = \frac{kT}{C_{\text{eqv}}}$$

The first term represents the reset noise, and the second term is the white noise from the op amp. The 1/f noise contribution is ignored since in most practical cases, the op amp operates at a frequency above the MOSFET 1/f noise knee frequency. The op amp noise, to a first order, is dependent only on the capacitance of the network. This is because the transconductance of the input FET controls both the bandwidth as well as the noise power spectral density.

The noise in each cycle circulates through each conversion algorithm. Its effect is similar to that of the offset, except that noise appears as a random offset. If the ADC is to operate without missing a code due to the integrator noise, the maximum allowable r.m.s. noise is related to the reference voltage as:

$$\langle V_n \rangle_{\Sigma-\Delta} \leq \frac{V_{\text{ref}}}{\sqrt{m(m+3)}}$$

$$\langle V_{\text{os}} \rangle_{\text{S-A}} \leq \frac{V_{\text{ref}}}{\sqrt{n} 2^n}$$

$$\langle V_n \rangle_{\text{cyclic}} \leq \frac{V_{\text{ref}}}{\sqrt{(2^n - 1)}}$$




Fig. 5 Required capacitance as a function of desired bit resolution (for an 0.25 LSB error)

Fig. 5 shows the effect of noise on ADC resolution. Since the noise depends only on the circuit capacitance, the effect of noise can be equivalently represented by defining an equivalent capacitance C_{eqv} . Usually, a capacitance size of 1 pF or more is used in CMOS implementations in order to suppress the effects of the parasitic capacitance. It can be seen from fig. 5 that for a 1 pF capacitance, noise does not limit resolution for a Σ - Δ ADC or a cyclic ADC. However, for a S-A ADC, noise alone limits the resolution to about 14 bits for a 1 pF capacitance. The relatively poor noise performance of this ADC algorithm is due to the fact that in each cycle the residual voltage decimates monotonically, whereas in other two schemes, the residual voltage exhibits non-monotonic behavior due to repetitive addition or subtraction of the full reference voltage. Of the three schemes, a Σ - Δ ADC is the most immune to circuit noise due to the presence of negative feedback.

Op amp gain

Finite op amp gain (A_{VO}), finite unity gain frequency (f_{U}), and capacitance mismatch (ϵ_c) also limit the conversion resolution. Figs. 6, and 7 show the dependence of the bit resolution as a function of A_{VO} , and ϵ_c , for a given conversion error of 0.25 LSB. Amongst the three methods chosen, the second-order Σ - Δ ADC is the least sensitive to op amp gain - a two-fold increase in gain is required to achieve an extra bit of resolution. In contrast, for a S-A and cyclic method, a four-fold increase in gain is needed to achieve the same resolution enhancement. Further, for any given resolution, the required op amp gain for a Σ - Δ ADC is much lower than that of the other schemes. To achieve




Fig. 6 Required op amp gain as a function of desired bit resolution (for an 0.25 LSB error)

18-bit resolution with a Σ - Δ ADC, an op amp gain of only 50 dB is needed. This can be easily obtained using conventional CMOS op amps occupying a relatively small area. This indicates the utility of Σ - Δ algorithm for column-parallel ADC

implementations. In contrast, the S-A scheme and the cyclic scheme respectively 90 dB and 80 dB op amp gain to achieve a resolution of 13 bits. Although op amp gain requirement can be relaxed by using additional switched capacitor networks to compensate for the gain error [13, 14], it requires additional chip area, and therefore may not be useful for on-focal-plane applications.

Capacitance mismatch

The error due to capacitance mismatch also shows a trend that is similar to that due to the op amp gain. With currently available CMOS technology, two capacitors can be matched no better than 0.01 %. Thus, unless additional circuitry is used, capacitance mismatch tends to limit the resolution of S-A and cyclic ADCs to 12-13 bits, and that of a Σ - Δ ADC to 22 bits as shown in fig. 7. The matching requirement can also be relaxed using additional switched capacitor circuit blocks, but at the expense of additional chip area [15].




Fig. 7 Required capacitance mismatch error as a function of desired bit resolution (for an 0.25 LSB error)

Speed

The finite unity gain frequency (f_u) of the op amp limits the maximum clock rate (f_c) at which the op amp can be operated because of slew and settling time requirements. The error due to finite op amp bandwidth affects all the three ADC schemes approximately in the same manner. However, it is possible to clock a Σ - Δ ADC almost twice as fast as the other two schemes, for a given resolution. On the other hand, if the clock frequency is at least 10 times slower than the unity gain frequency, the error due to settling time requirements affects resolution only when larger than 18-bit resolution is considered. Table 2 summarizes the resolution limitations of different ADC algorithms due to circuit imperfections.

TECHNIQUE	CAP. ERROR	AMP GAIN	UNITY GAIN FREQ.	OFFSET	NOISE
Assumed Mag.	0.01 %	70 dB	50 MHz	2 mV	60 μ V
Succ.-App. (S-A)	12	11.5	18	10	15
Cyclic	14	13	18	10	> 20
Σ - Δ	>18	> 20	> 20	10	> 20

Table 2 Effect of circuit non-idealities on the achievable resolution for three ADC algorithms

5. POWER DISSIPATION IN ON-FOCAL-PLANE ADCs

In CMOS digital circuits all power is dissipated during switching, and the power dissipated is directly proportional to the operating speed. This is not so for a CMOS analog circuit (e.g. an op amp). Such a circuit requires dc. bias current, which is the main source of power dissipation. The bias current also determines the circuit speed, since it controls the settling time (t_{set}) and the slew time (t_{slew}). For a well designed circuit, t_{slew} and t_{set} are chosen to minimize the bias current. Assuming a duty cycle of 50 %, the settling time t_{set} can be expressed as a function of the clock frequency as $f_c = \gamma / 2 t_{set}$; where γ is the fraction of the clock period spent in settling the signal. In that case, the required bias current (I_{bias}) as a

function of the clock frequency can be written as:

$$I_{\text{bias}} = \frac{2}{\mu C_{\text{ox}} \left(\frac{W}{L}\right)} \left[\frac{f_c C_L}{\gamma \beta \pi} \ln \left(\frac{1 - \beta}{\epsilon_{\omega}} \right) \right]^2$$

where C_L is the effective load capacitance, μ is the mobility of the input transistor, C_{ox} is the oxide capacitance, W and L are the gate and the length of the MOSFET, ϵ_{ω} is the settling error.




Fig. 8 Op amp power dissipation as a function of the clocking speed

power dissipation would have been the same in both the cases. However, since the power-frequency relationship of an analog circuit shows a square dependence, the total power dissipation in all the N op amps (used in column-parallel architecture) is less than that in the op amp (used in the serial ADC) operating N times faster. This allows a column parallel ADC architecture to be operated with lower power dissipation than a serial ADC.

6. CONCLUSIONS

This paper has presented several approaches for on-focal-plane ADC. While it is possible to consider a single serial ADC for the entire FPA, lower power, lower noise, and higher resolution can be achieved using a column-parallel architecture. For most applications, a column-parallel architecture seems to be most appropriate. In terms of ADC algorithms, for lower resolution ADC in a ideal environment (e.g. no cosmic rays, low data rate) a single-slope style ADC can be adequate. Successive-approximation or cyclic ADCs convert the data more rapidly but require more complex circuits. For high resolution, oversampled ADC such as Σ - Δ ADCs can be utilized. Oversampled ADCs require the most complicated circuits, but are most forgiving to circuit imperfections. The trade space for choosing an ADC for the FPA is rather flat, and several algorithms can be considered for a given imaging system. While there have not yet been many reported ADCs integrated on the focal-plane, the authors anticipate rapid progress in this area in the next few years. Future generations of FPAs will likely contain on-focal-plane ADCs.

7. ACKNOWLEDGMENTS

The authors appreciate the support of M. Chahine and V. Sarohia of JPL and G. Johnston of NASA Headquarters. The research described in this paper was carried out by the Center for Space Microelectronics Technology, Jet Propulsion Laboratory, California Institute of Technology, and was sponsored by the JPL Director's Discretionary Fund and the National Aeronautics and Space Administration, Office of Advanced Concepts and Technology.

Fig. 8 shows the power dissipation of an op amp as a function of the clocking speed. It can be seen that the power dissipation even for a 1 MHz clocking rate is only about 20 mW. This means that a low-power serial on-focal-plane ADC can be operated even at relatively high pixel data rate (~ 1 MHz). Secondly, the power dissipation increases as the square of the clock frequency. The implication of this is as follows. Compared to a single serial ADC, a column parallel architecture requires N ADCs (and therefore N op amps) each operating at a frequency that is N times smaller than a similar serial ADC. If the ADCs were comprised of digital circuits only, the total

8. REFERENCES

1. A.I. Krymsky and E.I. Cherepov, "Spiking oscillator networks for on-focal-plane IR image preprocessing," to appear in Proc. SPIE, vol. 2226, paper 25, 1994.
2. S.K. Mendis, B. Pain, R.H. Nixon, and E.R. Fossum, "Design of a low-light-level image sensor with on-chip sigma-delta analog-to-digital conversion", Proc. SPIE, vol. 1900, pp. 15-26, 1993.
3. C. Jansson, A.P. Ingelhart, C. Svensson, and R. Forchheimer, "An addressable 256x256 photodiode image sensor array with an 8-bit digital output," Analog Intg. Circ. & Sig. Proc., vol. 4, pp. 37-49, 1993.
4. D.E. Ludwig, N.D. Woodall, and M.M. Spanish, "On-focal-plane analog-to-digital conversion with detector gain and offset compensation," Proc. SPIE, vol. 1097, pp. 73-84, 1989.
5. P.W. Li, M.J. Chin, P.R. Gray, and R. Castello, "A ratio independent algorithmic analog-to-digital conversion technique," IEEE J. of Solid State Circ., vol. SC-19, pp. 828-836, 1984.
6. S. Ogawa and K. Watanabe, "A switched-capacitor successive-approximation A/D converter," IEEE Trans. Instrum. & Meas., vol. 42, pp. 847-853, 1993.
7. B.E. Boser and B.A. Wooley, "The design of sigma-delta modulation analog-to-digital converters," IEEE J. of Solid State Circ., vol. SC-23, pp. 1298-1308, 1988.
8. J. Robert and P. Deval, "A second order high-resolution incremental A/D converter with offset and charge injection compensation," IEEE J. of Solid State Circ., vol. SC-23, pp. 736-741, 1988.
9. P.E. Allen and D.R. Holberg, "CMOS analog circuit design," Chapter. 10, Holt, Reinehart and Winston, New York, 1987.
10. R. Gregorian and G.C. Temes, "Analog CMOS integrated circuits for signal processing," Chapter 7, John Wiley, New York, 1986.
11. G.C. Temes, "Finite amplifier gain and bandwidth effects in switched-capacitor filters," IEEE J. of Solid State Circ., vol. SC-15, pp. 358-361, 1990.
12. S. Ogawa and K. Watanabe, "A high speed, switched-capacitor analog-to-digital converter using unity-gain buffers," IEICE Trans. Fundamentals, vol. E76-A, pp. 924-930, 1993.
13. P.J. Hurst, R.A. Levinson, and D.J. Block, "A switched-capacitor delta-sigma modulator with reduced sensitivity to op amp gain," IEEE J. of Solid State Circ., vol. SC-28, pp. 691-696, 1993.
14. L.E. Larson and G.C. Temes, "Switched capacitor building blocks with reduced sensitivity to amplifier gain," Proc. IEEE Int. Symp. Circuits Syst., pp. 334-338, 1987.
15. H.T. Yung and K.S. Chao, "An error-compensation A/D conversion technique," IEEE Trans. Circuits Syst., vol. CAS-38, pp. 187-195, 1991.