Architectures and Issues for Gigasample/second ADCs

Ken Poulton, Robert Neff, Brian Setterberg,
Bernd Wuppermann, Tom Kopley
Agilent Labs, Santa Clara, California

Abstract

Architectures for ADCs at 1 Gigasample/second (1 GSa/s) and beyond now include flash, folding and interpolating as well as the time interleaving of slower unit converters such as pipeline and even successive approximation ADCs. In addition, CMOS is taking over in this former bastion of bipolar technology. We describe the issues common to all architectures: bandwidth, power, I/O, data storage, and cost. We examine these issues in detail for the time-interleaved approach as exemplified by two 8-bit ADCs operating at 4 GSa/s and 20 GSa/s, implemented in CMOS.

1. Introduction

Before 2001, analog-to-digital converters (ADCs) at rates over 1 gigasample/second (1 GSa/s) were mainly used in digital oscilloscopes. The architectures used were flash and folding-and-interpolating, some of them time-interleaved up to 4 ways. Sample rates went up to 4 or 5 GSa/s; the technologies were bipolar and power and cost were high.

The last five years have seen a huge change in ADCs used at 1 GSa/s and beyond. On the application side, increasing disk drive data rates has created the initial need for cheap, relatively low-power 6-bit ADCs in the gigasample region. More recently, future high-bandwidth communications applications (both wireless and wireline) are driving interest in ADCs with sample rates near 1 GSa/s with power compatible with battery operation.

In technology, the last five years have seen a wave of mostly-CMOS converters. This is enabled by the increasing speed of CMOS processes, and required for integration of ADCs into larger CMOS chips for cost-reduction reasons. The surprising development is the use of CMOS and pipeline architectures even for

non-embedded applications, up to the fastest 8-bit ADCs. This has been enabled by the development of massively time-interleaved architectures.

The issues that are specific to, or more critical in, gigasample ADCs include

- Wide bandwidth with low distortion and noise. The analog performance of an ADC (e.g., BW, SNDR, SFDR) depends most strongly on the front-end sampling of the input signal. The dynamic range required depends on the application.
- High sample rate with low metastable error rate. Flash and folding ADCs can make their metastable rates arbitrarily low simply by adding more pipelining after their comparators (at the expense of power), but this is more difficult to resolve for multi-stage converters. Applications with a lot of memory or peak detection can drive the need for metastable error rates as low as 10^{-17}.
- Power dissipation. Gigasample ADCs are usually close to (or beyond) the knee in the power curve where power must be increased much more than linearly to increase the sample rate. Even converters with low absolute power may be near this point.
- Interleave overhead. While time interleaving allows one to increase sample rate, it brings along its own issues:
    - Generation of multiple clocks with sub-picosecond time alignment
    - Gain and offset alignment
    - Physical distribution of analog signals and clocks

  These add both power and complexity.
- I/O or sample memory. Driving the samples off-chip, or storing them on-chip can take as much power as the ADC itself.

These issues will be illustrated in the rest of the paper.

2. 20th-Century Technologies for Gigasample ADCs

Most gigasample ADCs before 2001 shared several characteristics:

- Bipolar IC technology. Briefly stated, the idea was that to go fast, you should use the fastest technology. In the late 90's, 25-GHz bipolar technology could hit 1.5 GSa/s [1] and 2 GSa/s [2] and 50-GHz GaAs HBT technology reached 4 GSa/s [3]. 8-bit CMOS ADCs were still in the 100-MSa/s range. 6-bit CMOS ADCs were technology-limited to 500 MSa/s.

  Secondly, bipolar transistor accuracy (e.g., offset in a diff pair) is about 10 times better than in CMOS technology. These two characteristics made bipolar the "obvious" choice for gigasample ADCs.
- Emphasis on low complexity. In most bipolar processes, the number of transistors that can be fabricated with high yield is very low compared to CMOS. In addition, a relatively large current (on the order of 1 mA) is

required to achieve the full f_T of the transistor, which leads to a practical limit of around 3000 full-speed transistors on a 10-W chip.

- Front-end track-and-hold (T/H). It is difficult to maintain low distortion at high input frequencies through the analog preprocessing in a folding ADC. The multiple parallel comparators inherent in both flash and folding architectures provide another avenue for high-frequency errors due to sample-time mismatches. The usual solution was to include a front-end T/H, but bipolar track and hold circuits such as diode bridge samplers or switched emitter followers are very power hungry.
- Parallel PECL outputs. Most designs limited data rates on their parallel output ports to 1-2 GHz by adding demultiplexing circuits and increasing the number of outputs by 2-4x. These are also power-hungry.
- High Power. The above-mentioned bipolar ADCs were all over 5 W well before microprocessors forced the development of many of today's heat-dissipation technologies.
- Custom packaging. While [1] is in a BGA package, earlier ADCs used custom ceramic packages to handle the bandwidth, I/O and power-dissipation requirements. Many used multi-chip assemblies [2] [4].

A package photo of such an ADC [2] is shown is Figure 1. The core was a 7-bit bipolar folding and interpolating ADC, operating at 2 GSa/s. Two cores were on one chip, time-interleaved to get 4 GSa/s and voltage-interleaved to reach 8 bits. Two chips were time-interleaved to reach 8 GSa/s in an oscilloscope product. The chip included a preamp and a bipolar diode-bridge T/H for each ADC core, which dissipated nearly as much power as the ADC cores. The total power was 13 W.

To reduce the power required to communicate parallel data to the memory, the custom CMOS memory chip was placed next to the ADC and connected with chip-to-chip wirebonding.

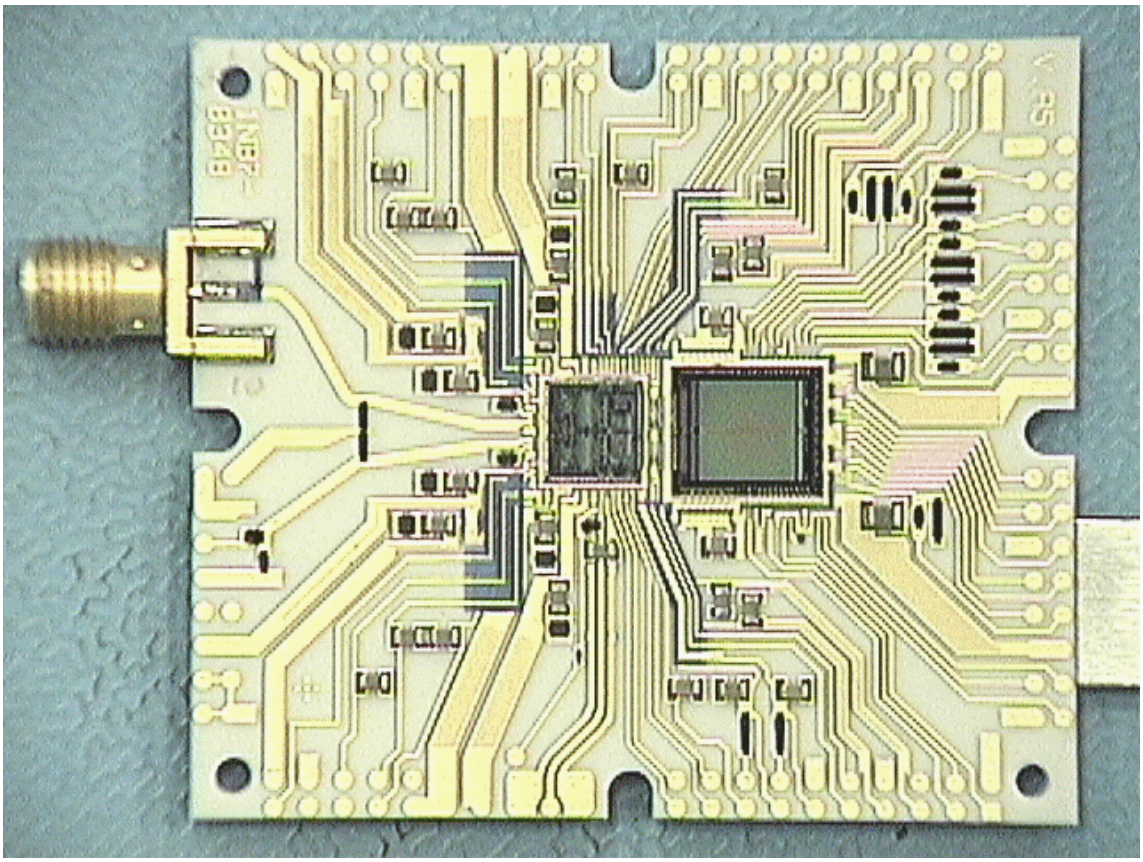

Figure 1: Custom thick-film hybrid package with custom bipolar ADC and custom CMOS memory chip

.

3. Trends and Tools

3.1. Trends

Several trends have driven a different approach to gigasample ADCs in the last 5 years. The first is in IC technology. For over two decades, CMOS technology has been following Moore's Law and the ITRS roadmap with amazing success and predictability. Progress in CMOS has been steady because the demand and rewards for digital circuits with smaller features have been so clear and the market has been growing at a huge rate. CMOS gate lengths have shrunk 33x in the last 25 years: from 3 um in 1980 to 90 nm, and gate delays have decreased by about the same amount.

Bipolar technology, by contrast, has moved in bigger steps, but much less frequently, because its shrinking share of the semiconductor market makes it harder to finance process development. Bipolar and BiCMOS technology has also been driven more by niche markets, such as CPUs (a mirage of the 90's) and RF circuits in the last 5 years. Silicon and SiGe bipolar f_Ts have changed by

about 40x, from 5 GHz in 1980, to 200 GHz now, but each period of stagnation has caused designers to consider other technologies. One of those periods was in the mid-90s.

A second aspect of technology scaling is the increasing level of integration required for lowest-cost solutions. This leads to the desire to integrate ADCs with other (mostly digital CMOS) circuits to reduce chip count and cost. Since CMOS wafers are much less expensive than BiCMOS, there is a strong motivation to move the ADCs into CMOS.

Even for chips that are standalone ADCs, the much lower cost of CMOS wafers and the dramatically lower power of CMOS logic makes CMOS ADCs attractive if the fundamental design problems around speed and accuracy can be overcome. The accuracy problems of CMOS were overcome very nicely in the late 80s by the use of switched-capacitor circuits and sigma-delta ADC architectures, but neither of these lent themselves to gigasample speeds. But they at least made it clear that CMOS could be used for some ADCs and other analog circuits.

Another trend is in ADC architectures: adding more calibration and digital corrections to augment the analog circuit performance of known ADC architectures. This both drives the need for more logic on the ADC chip and alleviates the need for the accuracy in the transistors.

3.2. Power Efficiency Figure of Merit

In order to compare different ADC approaches, the ADC power figure of merit (FOM) was developed (and first reported at AACD in 1992 in [5]):

$$FOM = P / (2^{ENOB} \ F_{S\text{-}NYQ})$$

where $F_{S\text{-}NYQ}$ is the minimum of F_S and 2*RBW, the BW where the SNDR has dropped by 0.5 bit from the low-frequency value. Smaller values indicate more power-efficient and accurate ADCs. The values for Nyquist (non-oversampled) ADCs reported at ISSCC through 1998 were 5-1000 pJ/conversion-level; reported values have dropped steadily since then to a best value of 0.16 pJ/conv-level in 2006.

It's important to keep in mind the inconsistencies in the numbers used to compute FOM. The largest is whether digital power, especially I/O, is included. This may not affect low-speed ADCs very much, but can change the FOM by a factor of 2 for gigasample ADCs. Other circuits such as input buffering and T/H circuits may be key to performance but not included in FOM calculations.

But beyond those errors, there is typically a 10:1 range in FOM results reported in a given year, based on the goals of the IC. Some papers may be focused on demonstrating a new architectural twist or calibration method rather than power and accuracy. Many designs are aimed at hitting a specific sample rate and accuracy with the technology available, which may lead to "inefficient" implementations as measured by the FOM.

When designing a new converter, FOM results can be useful in trying to compare different architectures and extrapolating from published accuracy levels to the desired level.

One key pitfall in this process is that for ADCs limited by thermal noise, scaling for higher or lower SNR does not result in a constant FOM. In this case, power scales as the square of resolution 2^{ENOB}, not linearly as implied by the FOM. The reason to keep using this FOM is that when we look over the whole range of ADCs (from 3 bits to 20) this FOM provides a more nearly resolution-independent value than one that scales with a constant FOM. This is consistent with the observation that different architectures work better at different sample rate points.

4. Two 21st-Century Gigasample ADCs

Based on the 20th century ADC technologies, it seemed silly to consider using CMOS to build scope ADCs at 4 GSa/s and beyond. But the trends and cost differences were compelling enough to stimulate us to investigate.

Our first target for a CMOS ADC was for 8 bits at 4 GSa/s, aimed at replacing the above bipolar ADC with a lower-power and lower-cost CMOS ADC. We took the time-interleaved approach to an extreme: we would use many independent ADC "slices" operating in round-robin fashion to reach the desired sample rate. We settled on a number of new design principles:

- Choose the core ADC architecture primarily for power efficiency of the unit converter rather than raw speed
- Time-interleave as many slices as needed
- Let currents and device sizes be optimized for SNR rather than matching
- Use calibration to recover the necessary accuracy

These principles were chosen to take advantage to the strength of CMOS: lots of transistors, low costs for digital logic.

After the success of the 4-GSa/s ADC [6], we extended this approach to 20 GSa/s, creating what is still the fastest 8-bit ADC in any technology [7].

4.1. ADC Slice Architecture

The selection criteria for the unit ADC slice included:

- Power efficient (low FOM)
- Area efficient (low area/sample_rate)
- High sample rate (to minimize the overall complexity)
- Very low metastable error rate (conflicts with high sample rate)
- Scalability to future processes (especially for VDD scaling)

It quickly became apparent that pipeline ADCs were attractive at the 8-bit level. However, the typical switched-capacitor approach was both slow and the necessary capacitors were area-intensive.

Instead, we selected a current-mode pipeline. The simplified schematic is shown in Figure 2. The basic circuit is a current mirror with a sampling switch between diode and mirror devices. The signal is carried in the current values rather than the voltages. This has several advantages: open-loop operation for high speed; current mirrors are more linear than other MOS open-loop amplifiers; no linear capacitors required; simplicity for small area and low power. Although the previous literature showed current-mode pipelines in the 10 MSa/s range, we determined that a simple design could operate at 125 MSa/s in 0.35-um CMOS.

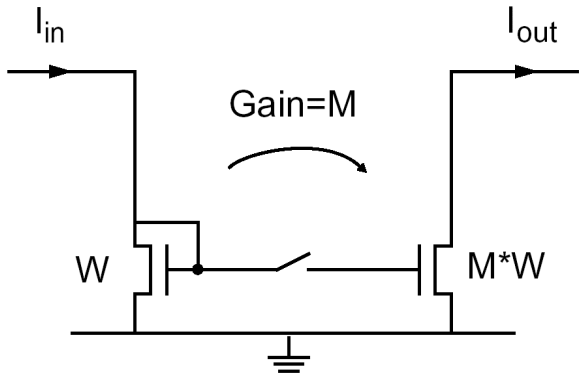

Figure 2: Simplified Schematic of the Current-Mode Amplifier in the Pipeline

When designed to meet the SNR needs of an 8-bit ADC, we had a small, low-power stage with very poor gain accuracy, far, far short of being able to achieve the required 8-bit gain accuracy in the first pipeline stage. We introduced redundancy with a reduced-radix architecture, resolving 1 bit per stage. (Note that many pipeline designs introduce redundancy with a 1.5-bit-per-stage design with stage gains of 2.00. This reduces the accuracy required of the comparators, but does not relax the gain accuracy required of the amplifier.) The nominal radix we chose was 1.6 to guarantee high yields with a large number of pipelines

on a die. This in turn called for 12 pipeline stages to reach a resolution of 256 levels.

The resulting pipeline requires three additional circuits to make a complete slice: 1) an input T/H, 2) an input transconductor stage to convert the input voltage samples into current and 3) a digital radix converter to convert 12-bits of radix-1.6 data to 8 bits of binary.

During the calibration process, the effective gain of each stage is extracted, which leads to the coefficients to be loaded into the radix converter. In this approach, we extract the fabricated stage gains with (at least) 8-bit accuracy instead of having to control them to that accuracy.

4.2. Input T/H

Time interleaving can get us to an arbitrarily large sample rate, but tends to reduce analog bandwidth and accuracy somewhat. The input T/H is the key to analog performance, limiting the bandwidth, linearity and SNR of the whole system.

Fortunately, CMOS has one truly fast circuit: a series-FET sampling switch (Figure 3). There are several conditions required to get the highest bandwidth from this sampler: there can be only one NMOS device in the signal path; the signal swing is small; the input common mode is near ground; the clock has the fastest-possible falling edge; the circuit is differential; and C_{HOLD} is composed only of circuit parasitics.

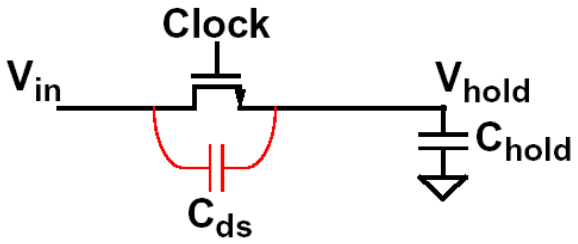

Figure 3 Simplified Schematic of the Input T/H

In 0.35 um technology, we got a 2 GHz bandwidth with –50 dBc HD3; in 0.18 um, we got a 7 GHz bandwidth with –50 dBc HD3.

4.3. T/H Architecture

There are several basic approaches to front-end T/Hs: single-T/H, tree structures, and time-interleaved T/Hs (one per slice, each directly connected to the input). The single-T/H and tree approaches have the key advantage that sample timing

is set by a single T/H clock. However, they tend to be limited by the loading after the first T/H to small interleave factors.

Multiple interleaved T/Hs makes the T/H-to-ADC-slice connection very simple and relatively easy to lay out, but creates very demanding input distribution and timing alignment problems. The alignment accuracy required can be calculated by considering the apparent voltage error created by a small timing error in sampling the fastest-slewing input signal. As a rule of thumb, for a 1 GHz input signal, 1 ps rms of jitter and clock misalignment will limit performance to 7 effective bits.

For the 4-way interleaving in [4] we chose a tree structure and we implemented the T/H tree in a faster technology (GaAs MESFET) than the ADC. For the high interleave factors we chose for the CMOS ADCs, interleaved samplers was the only feasible approach. In the 4-GSa/s ADC, we were able to simply drive all the interleaved samplers directly from the 100-ohm differential input to the chip. For the 20-GSa/s ADC, driving 4 pF at 5 GHz was not feasible from 100 ohms, so we added a SiGe input buffer chip to drive the T/H capacitance with a low impedance.

For the 4-GSa/s ADC, we used 32 slices, for the 20-GSa/s ADC, 80 slices. For the 20-GSa/s ADC, our goal was 5 effective bits at 5 GHz input signal which required better than 600 fs rms timing error including both clock jitter and residual misalignment. The uncalibrated alignment we achieved with careful design of parallel paths had around 10 ps rms misalignment, mostly due to random device fabrication mismatches in the clock paths. The calibration step adjusts digitally-controlled delay elements in each clock path.

4.4. T/H Clock Architecture

The chosen interleaved T/H architecture for the 4-GSa/s ADC requires the generation of 32 clocks at 125 MHz, each delayed by 250 ps from the previous. We targeted 1 ps rms for jitter and 1 ps rms for misalignment.

The simple way to generate these clocks would be a delay-locked-loop (DLL) with 16 stages of 250-ps delays (using both edges of the 16 stages to the 32 clocks). But the power required to reach 1 ps rms jitter with 16 stages would have been much larger than the total chip power budget. This is because power scales inversely with jitter squared as well as the accumulation of jitter down the delay chain.

The clocking architecture we use is shown in Figure 4. To generate the 250-ps delays between adjacent clocks, we use a 4-stage DLL at 500 MHz, which

reduces the power requirement by 16x. To get the 32 clocks at 125 MHz, we divide each 500 MHz phase by 4.

To get the required timing alignment, we introduce digitally-controlled delays in two places: 1) after the DLL and 2) after the divide-by-4 stages. The combination gives us a total delay-adjust range of around 100 ps and resolution of 250 fs.

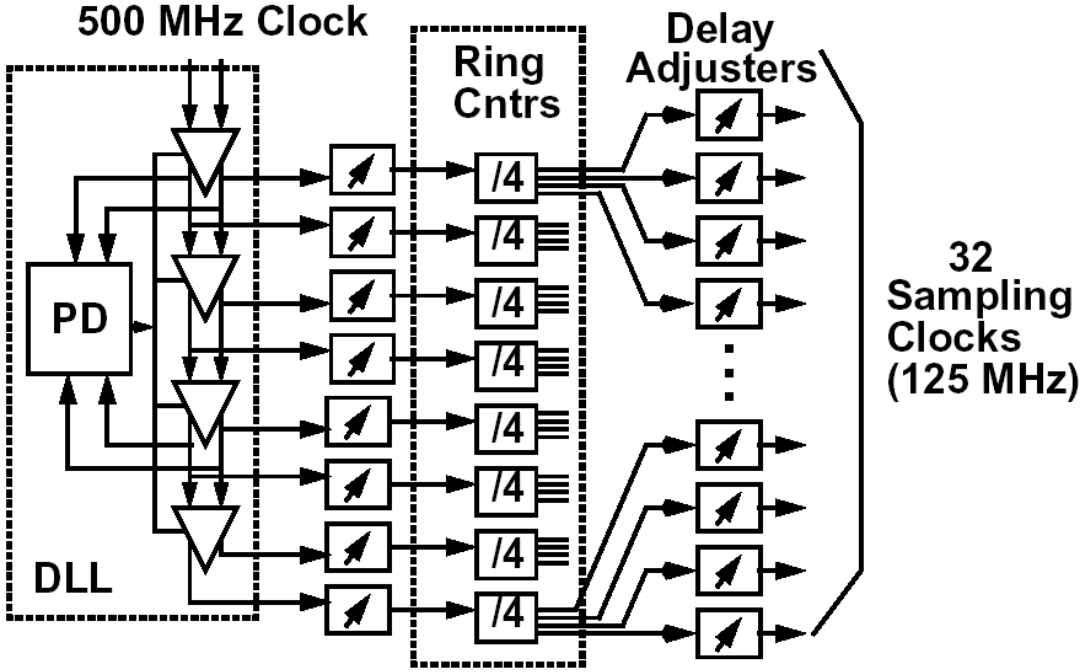

Figure 4 Clock Generator

4.5. Calibration

The overall block diagram of the 4-GSa/s ADC is shown in Figure 5 with the calibrations highlighted.

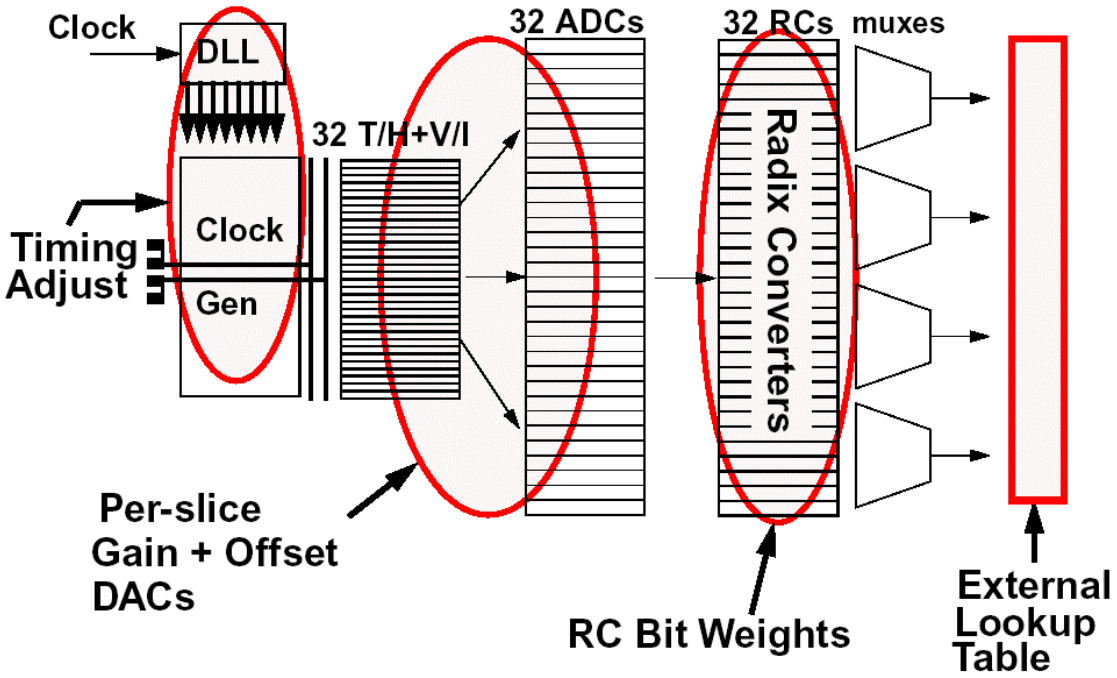

Figure 5 Block Diagram Showing Calibrations

All calibrations are done with special calibration signals switched into the input of the ADC. The calibrations are all digitally controlled, but we use a mix of analog and digital circuits to insert them into the signal path. The left-most calibration is the timing alignment, which uses digital values to adjust the timing of the clocks. The next one is the slice-to-slice alignment of gain and offset; this uses digital values to control per-slice gain and offset DACs. The third is the bit-weight values loaded into the Radix Converter. These may be thought of as the stage gains in each stage of each pipeline converter. The last is an external lookup table, which is used in some applications. This primarily corrects the minor third harmonic distortion created in the V/I stages in front of each current-mode ADC.

All of the calibration values must be calculated by software external to the ADC; we use the general-purpose CPU already present in an oscilloscope. However, once calibrated, the chip performs all the corrections except the optional lookup table in hardware in real time.

It has already been described how the use of calibration allows the use of less-accurate, and thus, smaller and lower-power circuits. In addition, the use of calibration relaxes the need to characterize and compensate for second- and third-order effects on gain accuracy, reducing design time. The use of iterative calibration allows the use of simple correction DACs with relatively poor gain control and linearity.

4.6. 20-GSa/s ADC

At 4 GSa/s, we sent the data out in LVDS in 4 parallel 8-bit words at 1 GByte/s for each word. For the next step to 20 GSa/s, we did not want to increase this to 20 words (320 pins), nor did we want to increase the word rate due to the difficulty of maintaining clock and data alignment. So we decided to put 1 MByte of sample memory on the ADC chip. The resulting chip is seen in Figure 6. It is evident that the memory dominates the chip. The SiGe chip to buffer the input capacitance is seen at the top; chip-to-chip wirebonding is used for the analog interface between the two chips.

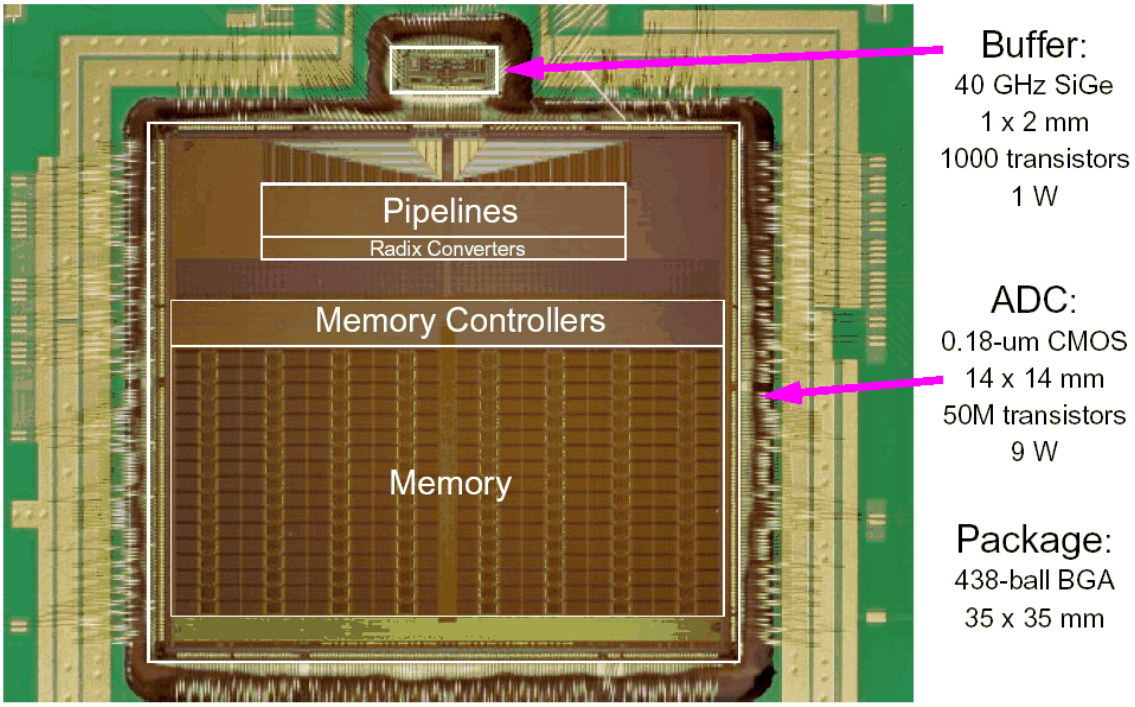

Figure 6: 20-GSa/s ADC Multi-chip Module Photo

4.7. Performance

The primary measure of performance for oscilloscope ADCs is signal to noise plus distortion ratio (SNDR), measured in dB or effective bits. In Figure 7 we plot the performance of the previous bipolar 4-GSa/s ADC along with the 4- and 20-GSa/s CMOS ADCs. Despite operating at 1/3 the power, the CMOS 4-GSa/s ADC is more accurate than the bipolar one by about 0.5 effective bits.

The gain of the 20-GSa/s ADC is seen in the lower graph; it is flat out to 5 GHz. The accuracy decreases inversely to input frequency at the high end; this is due mainly to the jitter in this ADC's sampling clocks, totaling 600 fs rms.

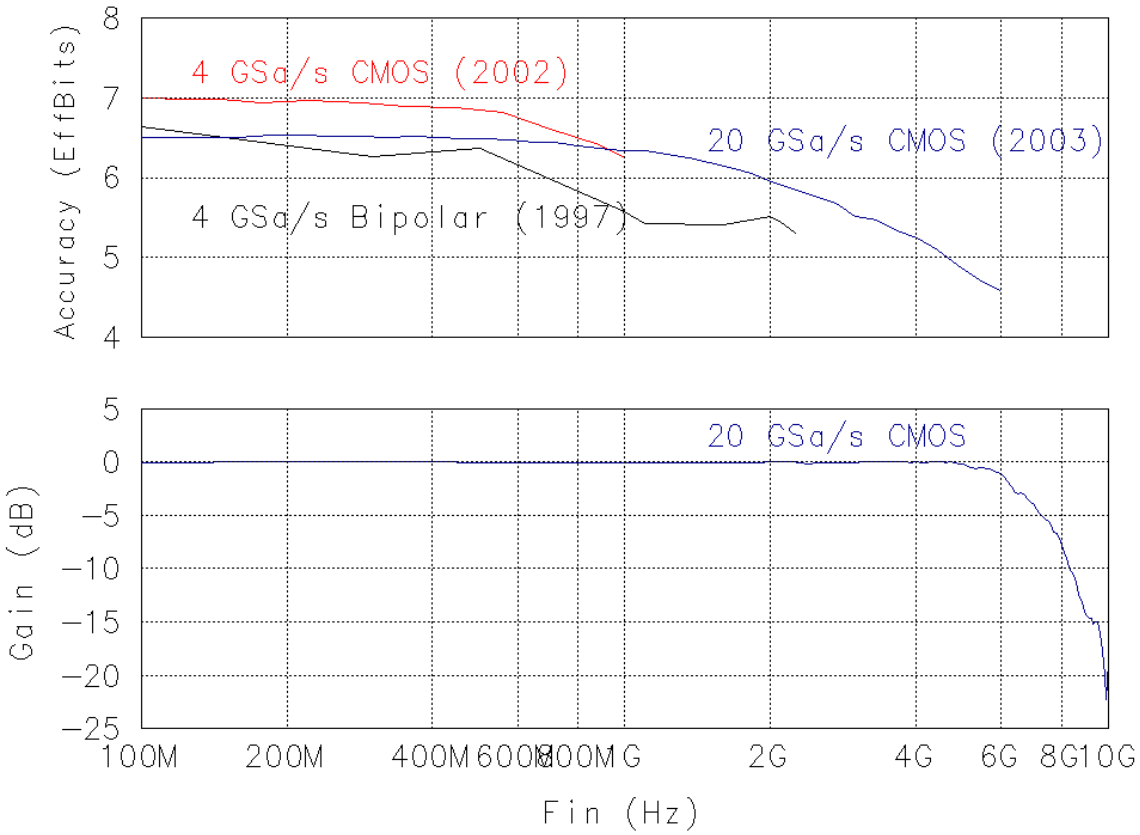

Figure 7: ADC Effective Bits and Gain vs. Input Frequncy

4.8. Power

The power breakdown of the two chips is shown in **Table 1**. The core ADC slices were chosen for power efficiency and provided good FOM values for their process generation.

However, the overall power FOM of the chips is around 10 times worse. This is due to several factors. On the analog side, the interleaved clock generator increases the power by about 150 mW. The power required for the full-speed outputs or on-chip memory nearly doubles the chip power. The choice of sample rate at 3-4x the input bandwidth provides high-fidelity waveforms, but further reduces the FOM compared to the simple Nyquist requirement of 2x the bandwidth. Finally, extending the bandwidth into the jitter-limited region (as seen in Figure 7) increases the sample rate needed to 5-10 times the resolution bandwidth. These all contribute to better performance for oscilloscope products, but combine to make a very poor-looking whole-chip FOM.

Table 1: Power Breakdown and Figure of Merit of the CMOS ADCs

	4 GSa/s CMOS	20 GSa/s CMOS
Input Buffer Chip	-	1.0 W
Pipeline ADC core	1.3 W	3.4
T/H + V/I	0.6	0.6
Radix Converters	0.45	0.6
DLL	0.15	0.13
Outputs	1.7	-
Memory	-	4.2
Total	**4.2 W**	**10.0 W**
Slice FOM	**4.5 pJ/conv-level**	**2.3 pJ/conv-level**
Chip FOM	**41 pJ/conv-level**	**27 pJ/conv-level**

5. Other 21st-Century ADCs

5.1. Architectures

The 2006 International Solid State Circuits Conference included many papers demonstrating the arrival of CMOS in the gigasample range. A 4-bit flash converter at 1.25 GSa/s [8] achieved a record FOM value because its low resolution can tolerate the bad threshold matching of CMOS comparators with no calibration.

The conference also showed several examples of new architectures arriving in the gigasample arena. One paper [9] showed a 6-bit subranging ADC, probably the first of that architecture to operate directly at 1 GSa/s. Another [10] described a 1-GSa/s ADC using 4-way time interleaving of pipeline ADCs to reach 11 bits. A third paper of note [11] described the interleaving of two 6-bit successive approximation converters to reach 600 MSa/s. Although not yet a gigasample ADC, it seems likely that with 4-way interleaving, we could see one of the slowest of all Nyquist ADC architectures used to go faster than 1 GSa/s.

And of course, all of these ADCs were implemented in 130-nm or 90-nm CMOS. The only BiCMOS ADC [12] was a 5-bit flash at 22 GSa/s.

5.2. Power Figure of Merit

A plot of power figure of merit vs. ADC input BW is shown in Figure 8 for papers at ISSCC from 1998 through 2005 [13]. We observe a region in the lower right with no data points. This is not due to a fundamental limit, but rather many practical considerations that make it difficult to achieve a low FOM along with very high bandwidth: input buffering, clock generation, signal distribution, jitter and I/O power. These are the same factors that cause our massively-interleaved

ADCs to have reasonable FOMs for single ADC slices, but high FOM values for the whole chip.

The fact that this is not a fundamental limit is shown by the addition of points for ISSCC 2006 in Figure 9. Many of these points are pushing into this formerly-empty region.

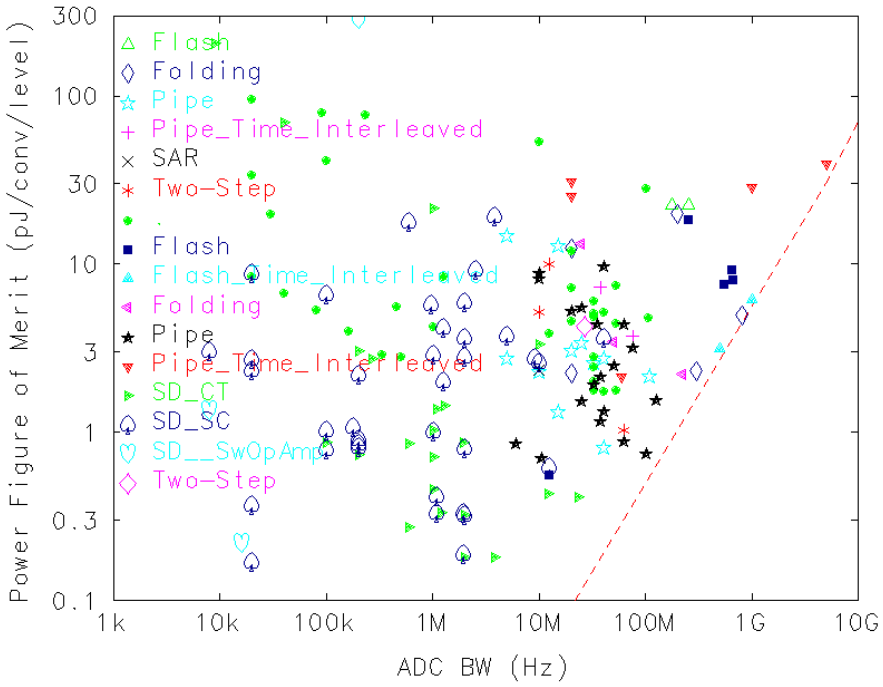

Figure 8: FOM vs. ADC Input BW for ISSCC papers 1998 though 2005

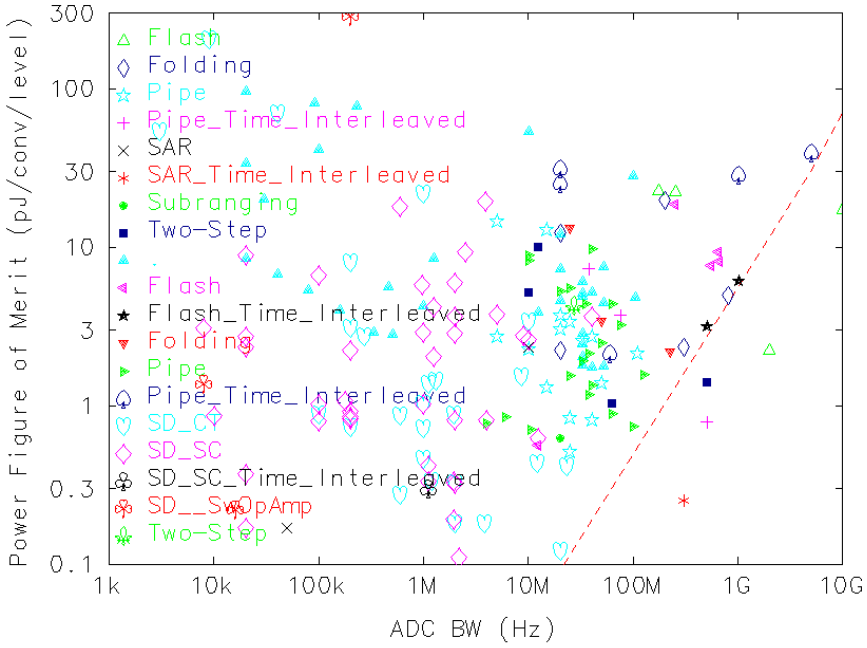

Figure 9: Power Figure of Merit including 2006

6. Summary

The turn of the millennium saw several major changes in gigasample ADCs. Most notable among these are: an increasing number of applications for such ADCs, an increased use of time interleaving enabling the use of slower ADC architectures as building blocks, and the march of CMOS into this formerly bipolar domain.

References

[1] MAX108 Datasheet, Maxim Integrated Products

[2] Ken Poulton, Knud L. Knudsen, John Kerley, James Kang, Jon Tani, Eldon Cornish and Michael VanGrouw,
"An 8-GSa/s 8-bit ADC System", IEEE VLSI Symposium on Circuits, 1997

[3] Ken Poulton, Knud L. Knudsen, John J. Corcoran, Keh-Chung Wang, Randy Nubling, Richard Pierson, Mau-Chung F. Chang, Peter Asbeck and R.T. Huang,
"A 6-bit, 4 GSa/s GaAs HBT ADC",
IEEE Journal of Solid State Circuits, Oct 1995

[4] Ken Poulton, John J. Corcoran, Thomas Hornak,

"A 1 GHz 6-bit ADC System", IEEE Journal of Solid State Circuits, Dec 1987

[5] J. Corcoran,
"High Speed Sample and Hold and Analog-to-Digital Converter Circuits", AACD 1992,
published in Johan H. Huijsing, Rudy J. van de Plassche, Willy M.C. Sansen "Analog Circuit Design", Kluwer Academic Publishers, December 1992, ISBN 0-7923-9288-4

[6] Ken Poulton, Robert Neff, Art Muto, Wei Liu, Andy Burstein, Mehrdad Heshami
"A 4-GSample/s 8b ADC in 0.35-um CMOS", IEEE International Solid State Circuits Conference Digest of Technical Papers, pp 166-167, Feb 2002

[7] Ken Poulton, Robert Neff, Brian Setterberg, Bernd Wuppermann, Tom Kopley, Robert Jewett, Jorge Pernillo, Charles Tan, Allen Montijo,
"A 20-GSample/s 8b ADC with a 1-MByte Memory in 0.18-um CMOS", IEEE International Solid State Circuits Conference Digest of Technical Papers, pp 318-319, Feb 2003

[8] Van der Plas, Geert; Decoutere, Stefaan; Donnay, Stéphane,
"A 0.16pJ/conversion-step 2.5mW 1.25GS/s 4b ADC in a 90nm Digital CMOS Process", IEEE International Solid State Circuits Conference Digest of Technical Papers, pp 566-567, Feb 2006

[9] Figueiredo, Pedro; Cardoso, Paulo; Lopes, Ana; Fachada, Carlos; Hamanishi, Naoyuki; Tanabe, Ken; Vital, João,
"A 90nm CMOS 1.2V 6b 1GS/s Two-Step Subranging ADC", IEEE International Solid State Circuits Conference Digest of Technical Papers, pp 568-569, Feb 2006

[10] Gupta, Sandeep; Choi, Michael; Inerfield, Michael; Wang, Jingbo,
"A 1GS/s 11b Time-Interleaved ADC in 0.13µm CMOS",
IEEE International Solid State Circuits Conference Digest of Technical Papers, pp 576-577, Feb 2006

[11] Chen, Shuo-Wei; Brodersen, Robert,
"A 6b 600MS/s 5.3mW Asynchronous ADC in 0.13um CMOS",
IEEE International Solid State Circuits Conference Digest of Technical Papers, pp 574-575, Feb 2006

[12] P. Schvan, D. Pollex, S-C. Wang, C. Falt, N. Ben-Hamida,
"A 22GS/s 5b ADC in 130nm SiGe BiCMOS",
IEEE International Solid State Circuits Conference Digest of Technical Papers, Feb 2006

[13] Data courtesy of Boris Murmann.