Continuous-Time Sigma-Delta ADC in 1.2-V 90-nm CMOS with 61-dB Peak SNDR and 74-dB Dynamic Range in 10-MHz Bandwidth

● David Enright ● Ian Dedic ● Gavin Allen

(Manuscript received January 8, 2008)

This paper describes a continuous-time sigma-delta (CTSD) analogue-to-digital converter (ADC) with 14-bit resolution, and a full-scale input of 1.0 V p-p differential. The circuit is implemented in an 8-layer 90-nm 1.2-V CMOS process. Integrated into the macro is a low-jitter clock generator (LC-PLL). The macro is configurable to allow various oversampling ratios and signal bandwidths, which makes it suitable for multi-band applications such as WiMAX. The CTSD ADC provides the advantage of low power, requires no anti-aliasing filter, and is immune to noise outside the ADC bandwidth. The CTSD architecture uses a third-order single-loop sigma-delta with a 4-bit quantiser. Continuous background calibration is used to tune the quantiser thresholds and digital-to-analogue converter (DAC) currents. A prototype of the CTSD macro has been taped out and measured. It has a peak signal-to-noise-plus-distortion ratio (SNDR) (ENOB = 9.8) of 61 dB and dynamic range of 74 dB in a 10-MHz bandwidth and requires only 31 mW per channel.

1. Introduction

There is a large demand in the wireless communication market for high-performance analogue-to-digital converters (ADCs) that have a large signal bandwidth in the region of several tens of megahertz and a resolution of more than 10 bits. However, the reduction in supply voltages that comes with the shrinking of transistors is making it difficult to realise standard analogue circuits. With the reduction in supply voltage, the dynamic range of the input signal is reduced. To keep the same performance, either the architecture needs to be changed or the thermal noise of the analogue components must be reduced, which in turn will normally increase the power dissipation.

The conventional architecture for a wireless communication converter is a pipelined ADC.[1),2)] This circuit requires an analogue anti-aliasing filter (AAF) and, due to the reduction in supply voltage, it is increasingly difficult to meet the required accuracy of >10 bits. A discrete time sigma delta (DTSD) normally operates with very large oversampling ratios (OSRs).[3)] The circuit can operate at very low supply voltages, but still requires a simple analogue AAF because of the sampled nature of the circuit and is normally limited to bandwidths of less than 2 MHz. Both these types of ADC have degraded signal-to-noise ratios (SNRs) because the sampled-data analogue circuits alias wideband noise into the signal bandwidth.

The continuous time sigma delta (CTSD) converter does not require a power-hungry anti-aliasing filter or suffer from noise aliasing because of the continuous-time (non-sampled) loop filter, and low OSRs can be used by using multi-bit quantisers. The non-sampled CTSD circuits also make the circuit less susceptible to high-frequency noise pickup, for example substrate noise generated by digital circuits.

This paper describes a CTSD ADC that is

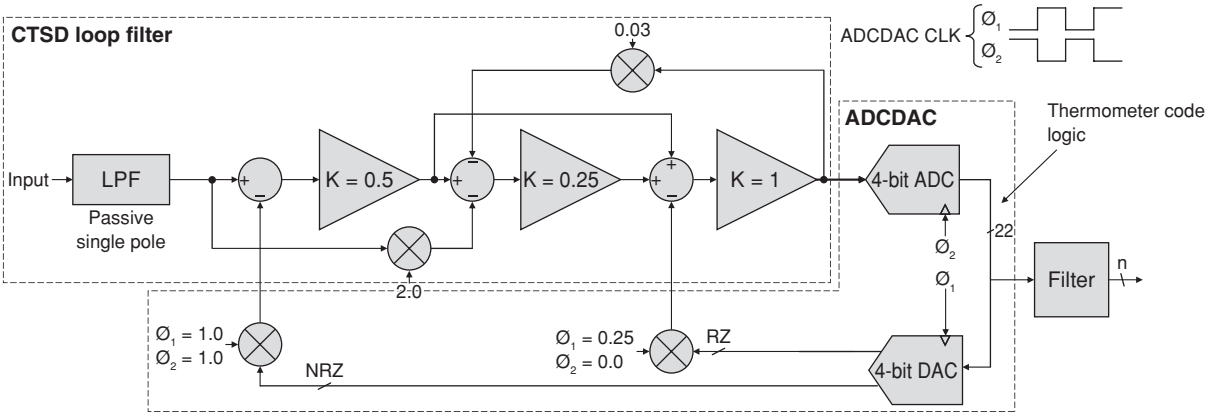

Figure 1
Sigma-delta macro.

designed to provide high performance whilst operating with a low supply voltage and having low power consumption. This makes the circuit ideally suited to mobile applications where supply voltages and power requirements are critical. Section 2 explains the CTSD architecture and operation of the circuit in detail. Section 3 describes the macro used to test the CTSD architecture. Experimental results for a test chip are presented in section 4.

2. CTSD

2.1 System level architecture

To meet the power and performance requirements, a multi-bit CTSD ADC was selected. This also greatly reduces the requirements for an AAF because the continuous-time integrators in the ADC already provide considerable aliasing protection. The CTSD architecture was optimised in Matlab to have a 10-MHz bandwidth and low noise whilst using the lowest possible power consumption. The filter order, OSR, and number of quantisation steps can all be varied to keep the same signal-to-noise-plus-distortion ratio (SNDR). The third-order loop allows a lower OSR than a second-order one. A fourth-order loop gives little improvement in SNDR at a similar OSR and is more difficult to stabilise (integrator gains are lower). If the OSR is too high, there can be fewer quantisation steps, but the clock path and switch power will be too large. If the OSR is too low, the switch power is low, but the complexity of the combined analogue-to-digital converter and digital-to-analogue converter (ADCDAC) is too high because of the large number of quantisation steps. A third-order filter with a 4-bit ADCDAC provides a good compromise between power and performance.

The architecture block diagram of the sigma-delta macro is shown in **Figure 1**. The third-order integrating loop filter uses a fully differential architecture. The loop filter consists of a passive low-pass filter at the input and three integrators, which have integrator gains (k) of 0.5, 0.25, and 1, respectively. The output of the loop filter connects to a 22-level flash ADC that outputs thermometer code logic. The 4-bit DAC has two current outputs: one output is a non-return-to-zero (NRTZ) pulse (low clock jitter sensitivity) that connects to the input of the integrating loop filter and the other is a return-to-zero (RTZ) pulse (reduced delay improves loop stability) connected to the input of the third integrator. The thermometer coded logic is decoded to binary logic and filtered with a decimating finite impulse response (FIR) filter to reduce the sampling rate and improve the signal-to-noise ratio (SNR). The loop filter

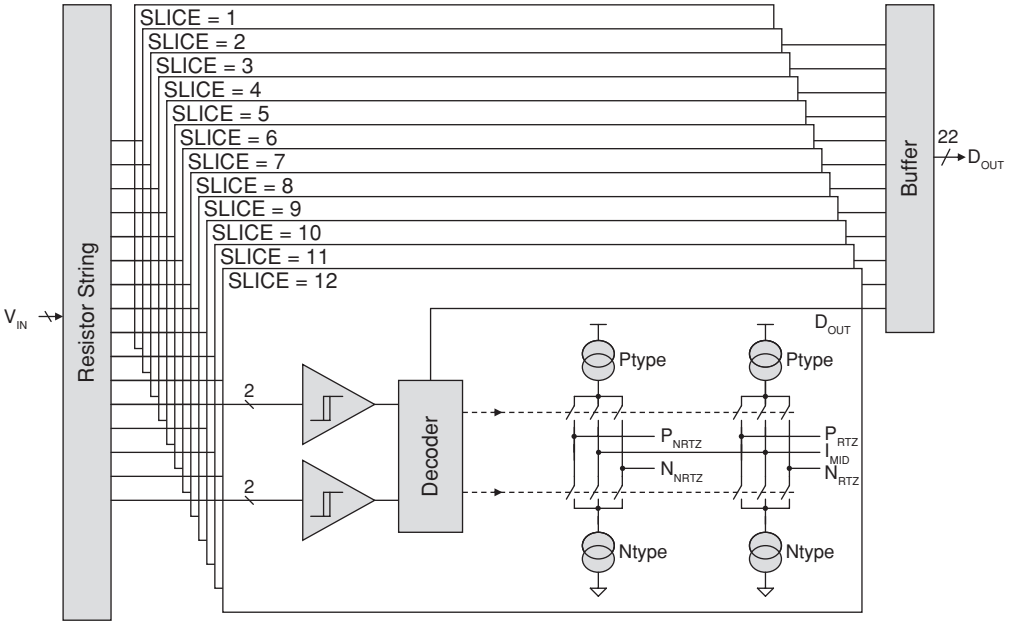

Figure 2
ADCDAC.

multipliers, which represent circuit gains, are constants (0.03 and 2.0), but the 4-bit DAC outputs have gains that change with the phase of the ADCDAC clock ($Ø_1$ and $Ø_2$) to describe the RZ and NRZ outputs.

There are four basic building blocks in the ADC: the ADCDAC, CTSD loop filter, digital filter, and reference clock. These are described in the following sections.

2.2 ADCDAC

The ADCDAC consists of a flash ADC and two current-steering DACs. The macro is designed using a slice approach, where two comparators of the ADC and two current steering sources/sinks are placed in a single slice (**Figure 2**). The resistor string at the input produces a different threshold voltage for each of the comparators. The slice structure allows the macro to be reconfigured easily for different applications, where the number of quantisation levels may be different. All the high-speed-data-dependant signals are self-contained inside each slice, which minimises data-dependant crosstalk between slices and channels.

The ADCDAC has been optimised for both area and power. Reducing the area of the current source array in the current-steering DAC increases the DAC output current distortion (due to transistor mismatch), which in turn reduces the SNDR of the CTSD. The largest power reduction in the ADCDAC was made by scaling down the transistors in the flash ADC and clock path. However, reducing the size of the transistors in the comparators increases their input offset voltages, which also affects the SNDR of the CTSD.

To counter the problems introduced by reducing the macro area and power, each current source and comparator is calibrated.

A slice can be switched out in real time without causing distortion in the input signal. When the slice is deactivated, all the elements are calibrated. The comparators are calibrated by connecting the inputs of the comparator together and then digitally programming the well voltage of the input pair until the output of the

comparator changes state.

This calibration process runs in the background continually whilst the CTSD ADC is running; this allows effects such as temperature and voltage supply variation, which can affect the circuit performance, to be minimised.

The current steering DAC outputs are calibrated by wiring the source to a sink current and integrating the error current. The result is then used to digitally tune the well voltage of the cascode in the current source cell.

The layout of the ADCDAC is telescopic to allow multiple channels to be abutted together and share clock routing, biasing, and supply rails.

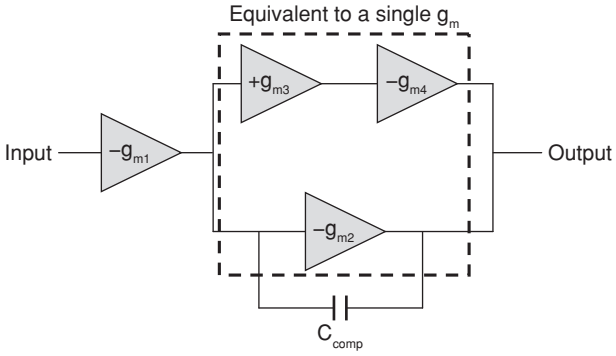

Figure 3
Loop filter amplifier.

2.3 CTSD loop filter

The loop filter (Figure 1) is a fully differential continuous-time third-order design implemented with constant internal resistors and variable capacitor units tuned to remove resistor and capacitor process variation. These capacitors are controlled to 5-bit accuracy over the full process variation.

2.3.1 Amplifier implementation

With a 10-MHz signal bandwidth and a high oversampling ratio, the bandwidth requirement for the loop filter amplifiers is at least 1.7 GHz. High-gain (90-dB DC gain) is also necessary for low in-band distortion. A four-stage fully differential architecture was used, which gives high bandwidth and DC gain. A single-ended representation of the amplifier in terms of transconductances is shown in **Figure 3**.

Stages g_{m1} and g_{m2} form a usual Miller compensation amplifier, with compensation capacitor C_{comp}. This path within this amplifier provides all the high frequency response for the full amplifier. Stage g_{m3} is used within this amplifier for DC gain, and it has negligible gain at high frequencies. Stages g_{m2}, g_{m3}, and g_{m4} are essentially equivalent to a single g_m, so the full amplifier is roughly equivalent to a two-stage Miller amplifier. This architecture makes it possible to have the high bandwidth and high DC gain needed for the filter implementation. Note that g_{m3} needs to be implemented as a positive g_m for stability. The four-stage main amplifier is implemented as shown in **Figure 4**.

For minimum power, the biasing of the amplifier was accomplished as follows. The input stage (g_{m1}) is biased for a constant g_{m1}/C_{comp} ratio, which gives constant bandwidth and white-noise performance over process corners. All other stages are biased to follow internal resistance to track resistive loading of the amplifier within the filter. Common-mode feedback is accomplished by feeding back currents under the first-stage cascode devices.

2.3.2 Constant G_m biasing

For the amplifier to use the minimum amount of power for noise and speed requirements, the input stage was biased using a constant transconductance circuit. As the filter uses an RC tuning circuit to set the RC time constant to a constant value, this tuning word can be used to set a constant g_{m1}/C_{comp} for the filter amplifier. The bias loop is shown in **Figure 5**.

The bias loop produces a constant g_m related to the ratio of M1, M2, and the value of R1. This gives the g_m relation shown in equation (1).

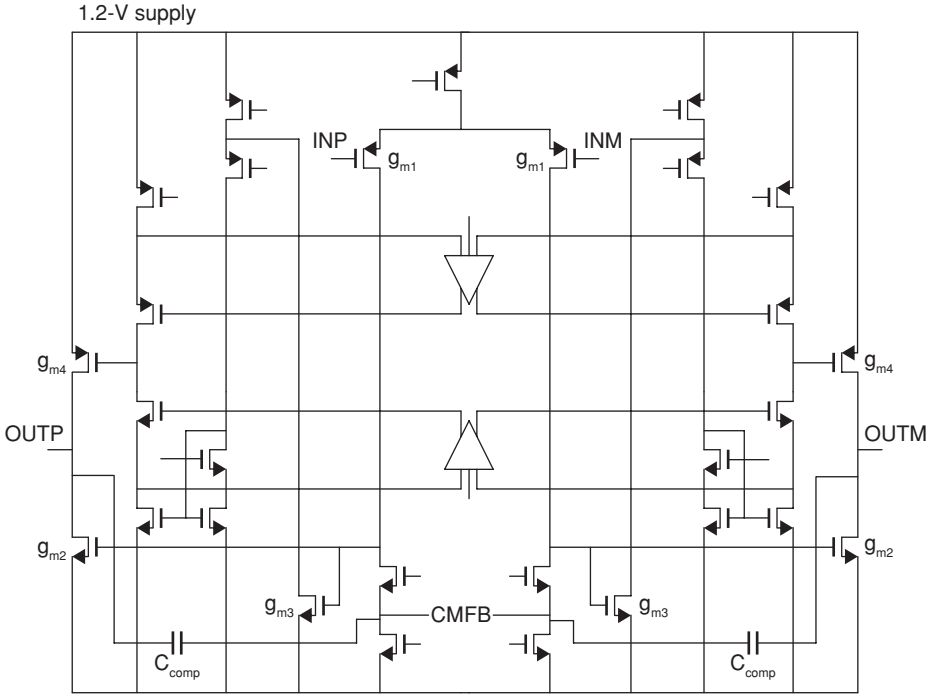

Figure 4
Amplifier implementation.

$$g_m \propto \frac{1}{R_{BIAS}} \quad (1)$$

The gain-bandwidth of the filter amplifier is given by

$$GBW = \frac{g_{m1}}{C_{comp}} \propto \frac{1}{R_1 C_{comp}}. \quad (2)$$

In the loop filter, the capacitance is digitally programmed to give the required RC time constant. The amplifier can use the same process word to vary g_m in the bias circuit to produced a constant g_{m1}/C_{comp} ratio in the amplifier.

2.3.3 RC tuning circuit

The process tuner produces a digital word that makes the RC product of the filter constant by digitally programming the capacitors. The implementation of the RC tuning circuit uses a single integrator and matches an internal resistor ([reference voltage]/R) current and a switched capacitor unit current. Using a digital successive approximation algorithm, the integration is repeated until the 5-bit word is resolved. This

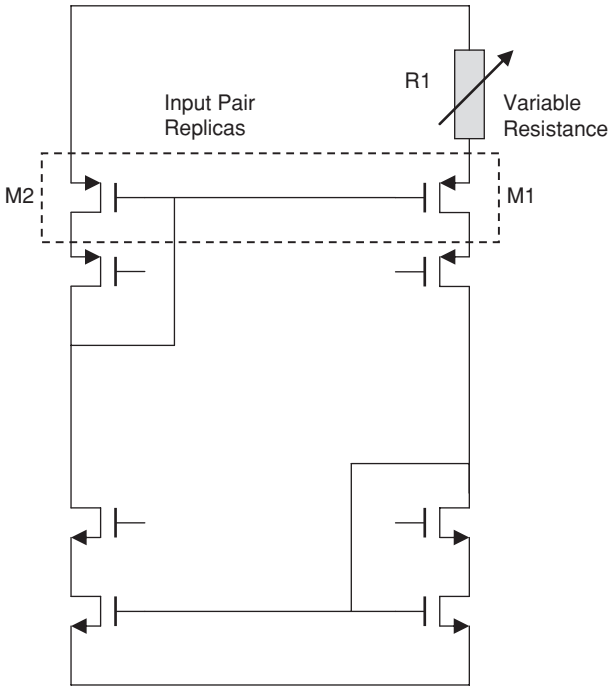

Figure 5
Constant Gm bias.

gives the correct RC time constant.

2.4 Digital filter

The main design priority for the digital decimation filter is low power consumption because the input data rate is up to 500 MHz. Another requirement is a variable decimation ratio to allow frequency planning to avoid interference with intermediate and radio frequencies and allow for different WiMAX channel bandwidths.

The conventional architecture would be a cascaded-integrator-comb (CIC) filter followed by an FIR filter, but this needs eight CIC sections, which would consume too much power because of the wide adders (up to 40 bits) clocked at the input data rate.

Instead of this, a single-stage multiplier-less FIR filter with variable decimation ratio was used. The multipliers are replaced by 2-bit canonic-signed-digit (CSD) stages, and, for a decimation ratio of M, the filter uses M parallel polyphase filters clocked at the output data rate. This gives the lowest possible power at the expense of an increased gate count, which is not an issue in 90-nm technology. The filter can realise decimation ratios of 9/10/11/14/15/18/20/22/28/30 with power consumption of 10 mW per channel.

2.5 Reference clock and biasing

A low-jitter high-frequency clock is implemented running at the OSR multiplied by the output data rate. This is implemented using a standard LC integer-N phase locked loop (PLL) that can be fully programmed via the serial interface. This allows a wide selection of external reference clocks to be used. The loop filter is external. The internal charge pump is driven by a type-4 phase-frequency detector and can be programmed between 65.2 and 437.5 µA to optimise the external loop filter. The tuning range has a 2-bit coarse tuning decoded off the CTSD OSR setting.

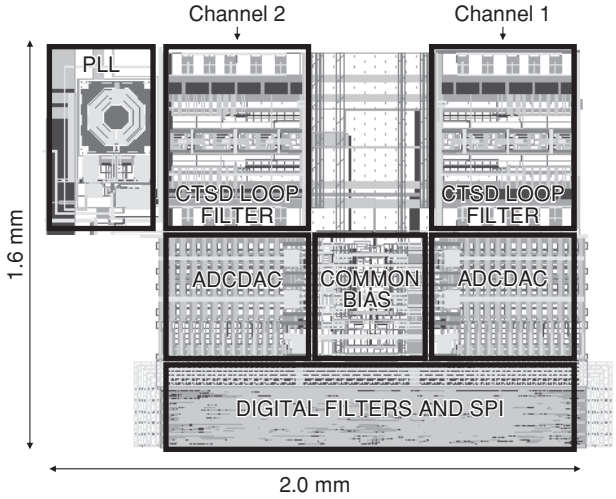

Figure 6
CTSD prototype layout.

3. Prototype implementation

A prototype CTSD was fabricated using 8-layer 90-nm CMOS technology. The layout is shown in **Figure 6**. A PLL reference of 22.4 MHz and external 0.5-V reference voltage are required. The CTSD occupies 3.2 mm^2, which includes the common biasing, reference PLL, two CTSD channels, digital filters, and serial interface. The CTSD macro was prototyped in a test chip to verify the CTSD architecture. The PLL is designed to output a fixed clock frequency of 716.8 MHz. The RC time constant for the loop filter components are also calibrated manually. The digital filter has a fixed 20-MHz bandwidth and eight times decimation. The ADCDAC calibration state machine is implemented in Matlab and communicates with the ADCDAC via a serial interface.

4. Measurement results

The performance of the CTSD is summarised in **Table 1**. The results for the prototype were measured with an ADCDAC clock of 716.8 MHz. The CTSD was measured with the digital filter deactivated and the output of the CTSD ADC routed directly to the external bus. Three succes-

Table 1
Performance summary of CTSD.

Technology	90-nm CMOS process
Resolution	14 bits
Sampling rate at digital filter	89.6 MHz
ADCDAC clock frequency	716.8 MHz
Supply voltage	1.2 V +/– 0.1 V
Signal bandwidth	10 MHz
SNDR (10-MHz BW)	61 dB
ENOB	9.8 bits
Dynamic range	74 dB
Single-channel power consumption	31 mW
Total power consumption including biasing PLL, and 2 channels	86.4 mW
Area	(2 channel) 3.2 mm^2
Peak input signal	1 V p-p differential
Phase noise of PLL (716.8 MHz @ 1 MHz offset)	–121.8 dBc/Hz

Table 2
Measured results.

Fundamental	–4.02 dBFS
2nd harmonic	–77.09 dBFS
3rd harmonic	–70.40 dBFS
4th harmonic	–86.14 dBFS
5th harmonic	–94.17 dBFS
SNDR (5-MHz BW)	64.10 dB
SNDR (10-MHz BW)	60.96 dB

sive thermometer coded words from the CTSD were translated into 2's complement words and clocked out of the macro at 238.933 MHz as a single 15-bit word. This word was captured by a logic analyser and the spectral performance was analysed using Matlab. All the ADC performance measurements were obtained using this test port.

The spectral performance of the prototype ADC is shown in **Table 2**. The input frequency was 1.99 MHz @ –4.02 dBFS. The prototype's spectral performance for a large signal is shown in **Figure 7** and that for a small signal is in **Figure 8**.

The results before and after calibration with two input signal levels of –4 dBFS and –40 dBFS are shown in **Table 3**. For the smaller input signal level, there was a 3.4-dB improvement in SNDR. For the –4 dBFS signal, the improvement was 4.2-dB.

The figure of merit (*FoM*, see equation 3) can be used to compare the prototype ADC with other converters having published data.

$$FoM = \frac{Power}{2^{ENOB} \times 2 \times F\max} \quad (3)$$

The two-channel prototype ADC dissipates 86.4 mW of power with a 1.2-V supply. This power figure includes the current from the common biasing, PLL, CTSD loop filter, ADCDAC, and digital filter. The power dissipation for a single channel without biasing and reference clock is 31 mW.

$$ENOB = \frac{SNDR - 1.76\ dB}{6.02}, \quad (4)$$

where *ENOB* is the effective number of bits (see equation 4) and *Fmax* is the bandwidth (using 61 dB from Table 3 and 10 MHz). The *FoM* for this ADC is 1.7 pJ per conversion. This figure is compared with some recently reported data[1),2),4)] in **Table 4**.

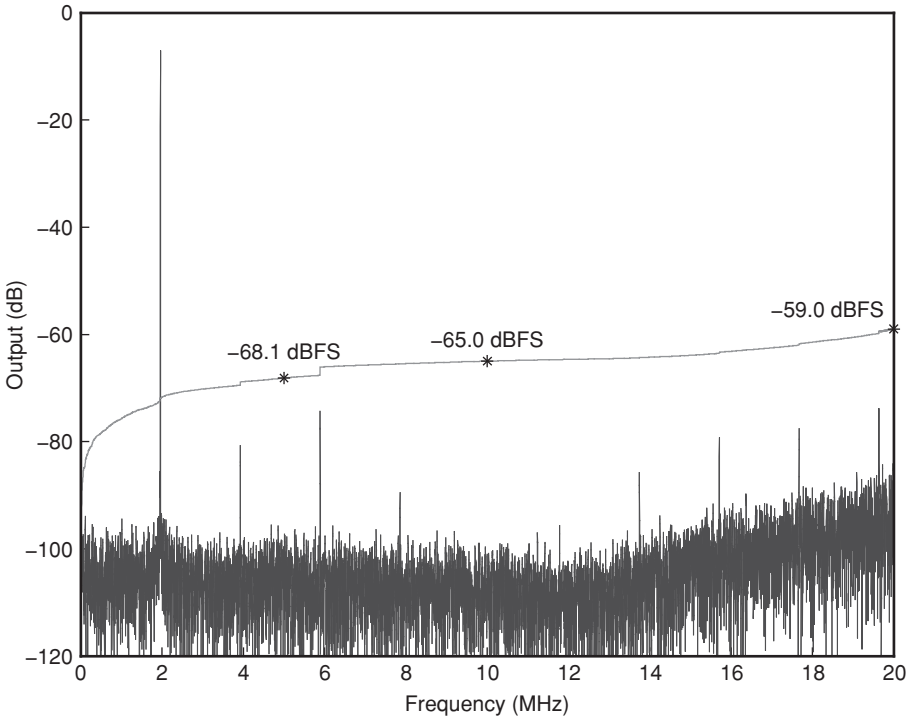

Figure 7
Measured spectrum (large signal –4 dBFS).

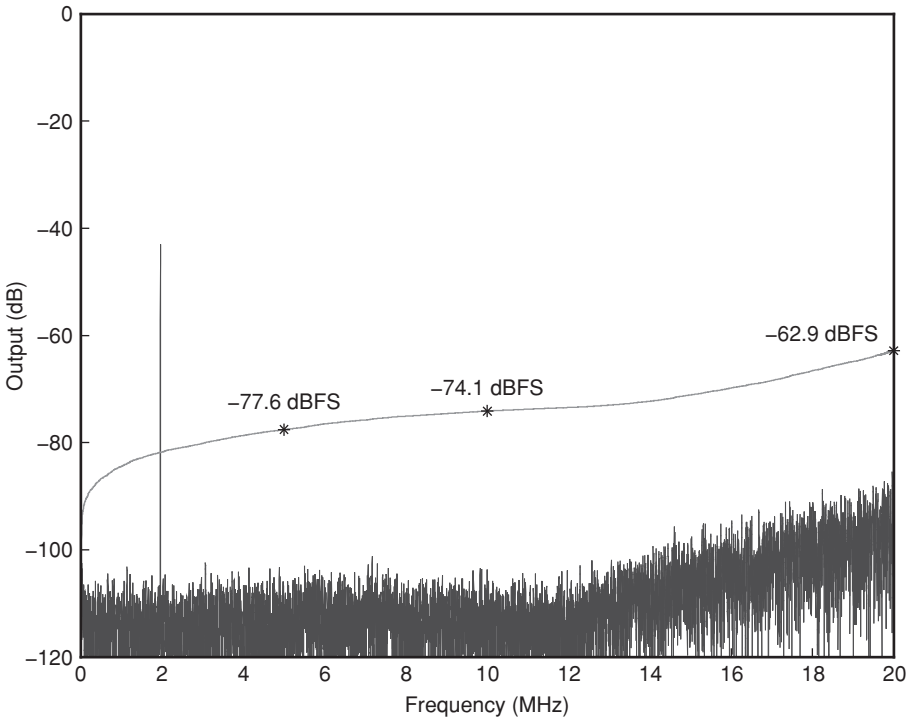

Figure 8
Measured spectrum (small signal –40 dBFS).

Table 3
Pre- and post-calibration results.

	Uncalibrated output –4 dBFS	Calibrated output –4 dBFS	Uncalibrated output –40 dBFS	Calibrated output –40 dBFS	Dynamic range after calibration
SNDR (5-MHz BW)	59.88 dB	64.09 dB	33.52 dB	37.55 dB	77.55 dB
SNDR (10-MHz BW)	56.73 dB	60.96 dB	30.66 dB	34.10 dB	74.10 dB

Table 4
Performance comparison.

	Sampling rates (MS/s)	Power consumption per channel (mW)	FoM (pJ/conversion)
Reference 1)	220	135	1.6
Reference 2)	125	40	0.59
Reference 4)	40	38	0.23
This work	89.6	31	1.7

5. Conclusion

This paper described a low-voltage CTSD converter. The architecture is suitable for high-speed and high-resolution analogue-to-digital conversion. The results for the prototype represents the performance of an early version of the macro. The circuit has since undergone modifications to reduce the ADCDAC noise floor and power consumption, and a multi-mode digital filter has been introduced.

The structure of the macro allows for the addition of multiple channels that share the same reference clock and biasing circuit. Moreover, the building blocks can be used to produce different ADC variants and could produce a CTSD DAC.

The architecture allows background calibration of the ADCDAC, so if there are any supply or temperature changes, the calibration algorithm will track the effects. Through the use of calibration, the area and power consumption of the macro have been greatly reduced.

The additional advantage with the calibration scheme is that upon initial power up, the CTSD will function but the SFDR will be lower, though as the calibration continues in the background, this will improve.

References

1) B. Hernes, A. Briskemyr, T. N. Andersen, F. Telstø, T. E. Bonnerud, and Ø. Moldsvor: A 1.2 V 220 MS/s 10 b Pipeline ADC Implemented in 0.13 µm Digital CMOS. ISSCC Digest of Technical Papers, 2004, p.256-257.
2) M. Yoshioka and M. Kudo: 10-bit, 125 MS/s: 40 mW Pipelined ADC in 0.18 µm CMOS. *FUJITSU Sci. Tech. J.*, **42**, 2, p.248-257 (2006).
3) J. Yu and F. Maloberti: A Low-Power Multi-Bit ∑Δ Modulator in 90-nm Digital CMOS Without DEM. *IEEE Journal of Solid-state Circuits*, **40**, 12, p.2428-2436 (2005).
4) G. Mitteregger, C. Ebner, S. Mechnig, T. Blon, C.Holuique, E. Romani, A. Melodia, and V. Mellini: A 20-mW 640-MHz CMOS Continuous-Time ∑Δ ADC with 20-MHz Signal Bandwidth, 80-dB Dynamic Range and 12-bit ENOB. *IEEE Journal of Solid-state Circuits*, **41**, 12, p.2641-2649 (2006).

David Enright
Fujitsu Microelectronics Europe GmbH
Dr. Enright received the B.Eng. (Hons) and Ph.D. degrees from Essex University in 1995 and 2000, respectively. He joined Fujitsu Microelectronics Europe, Maidenhead, England in 2000, where he has been engaged in research and development of RF macros, HSIO front-ends, and high-performance analogue-to-digital converters.

E-mail: david.enright@fme.fujitsu.com

Ian Dedic
Fujitsu Microelectronics Europe GmbH
Mr. Dedic received the B.A. (Hons) degree from Cambridge University in 1980. He joined Fujitsu in 1990 in the Communications VLSI Research Laboratory in London. Since 1993, he has worked on all aspects of mixed-signal design for Fujitsu Microelectronics in Maidenhead.

E-mail: ian.dedic@fme.fujitsu.com

Gavin Allen
Fujitsu Microelectronics Europe GmbH
Mr. Allen received the B.Eng. (Hons) degree from the University of Western Australia in 1997. He joined Fujitsu Microelectronics Europe, Maidenhead, England in 2001, where he has been engaged in the research and development of high-performance analogue macros including fully differential operational amplifiers.

E-mail: gavin.allen@fme.fujitsu.com