The Design of A High-Bandwidth Sigma-Delta Modulator

Brian Limketkai and Bret Victor

***Abstract*—The design of a high-bandwidth $\Sigma\Delta$ modulator which achieves 10 bits of resolution with a conversion rate of 20 MS/s is presented. The oversampling ratio is 16, requiring a sampling frequency of 320 MHz. The modulator is implemented as a fourth-order 2-1-1 cascade using switched-capacitor integrators, which is amenable to simple implementations in a CMOS process. The component values are scaled to maximize the dynamic range at the outputs of each integrator. In a typical 0.25μm process, we estimate the overall modulator power consumption to be around 40 mW.**

I. Introduction

THE growing demands for high-data-rate wireless devices have led to an abundance of wireless standards requiring high-bandwidth analog-to-digital converters (ADCs). One widely accepted standard designed for wireless local area network (LAN) applications is the IEEE 802.11a standard. This paper describes the design of a fourth-order 2-1-1 cascade sigma-delta ($\Sigma\Delta$) modulator which can be used in devices implementing this standard. The $\Sigma\Delta$ modulator was chosen over other modulators because of its greater insensitivity to analog circuit imperfections, which is becoming more of a problem as device dimensions shrink.

The 802.11a standard [11] defines a set of channels, each with a bandwidth of 16 MHz. Each channel is further divided into 53 subchannels of 300 kHz each. This allows for the repartitioning of subchannels to different users as data rate requirements change. However, the standard also defines the middle subchannel as an empty subchannel where no information is transmitted. This allows for the use of direct downconversion when implementing the standard. Thus, the Nyquist rate required for this ADC is 20 MS/s, allowing for the excess bandwidth of the transmit filter. The required resolution has been determined to be 7–10 bits [10], where the use of higher resolutions can be traded off with analog RF design complexity.

Section II presents the system design of the $\Sigma\Delta$ modulator. Section III describes the system-level modeling aspects of the modulator and the simulation framework. The nonidealities due to the circuit implementation and an estimate of the modulator power consumption are described in Section IV, followed by simulation results in Section V, followed by our conclusions.

Although this paper describes the design of a high-bandwidth ADC tailored to the 802.11a standard, much of the results are also applicable to other high-data-rate wireless standards.

II. System

It can be shown [9] that the ideal SNR of a $\Sigma\Delta$ converter with a one-bit quantizer, assuming an input amplitude equal to the modulator output amplitude and a perfect differentiator noise transfer function of $(1 - z^{-1})^L$ is

$$10 \log_{10} \left(\frac{3}{2} \frac{2L+1}{\pi^{2L}} M^{2L+1} \right) \tag{1}$$

where L is the modulator order and M is the oversampling ratio. We wanted the oversampling ratio to be a power of two, as is typical, in order to ease design of the decimation filter. However, an OSR of 8, with any reasonable modulator order, does not provide enough dynamic range for our application, and an OSR of 32 presents circuit speed requirements that are too stringent for our technology. Hence, an OSR of 16 was selected. This provides an ideal maximum SNR of 64 dB for a third-order modulator and 80 dB for fourth order. The third-order case is too close to our target dynamic range of 60 dB, not leaving enough margin for considerations such as noise, reduced input range, and integrator output scaling. Thus, we chose a modulator order of four.

At this point, we had to decide between two different families of modulators. A so-called single-loop modulator consists of a single quantizer, and a loop filter with an order equal to that of the overall modulator. A cascade, or MASH, topology uses multiple low-order loops, with the error produced by each quantizer feeding into the next stage. A digital filter then uses the output of each loop to construct a signal that has noise shaping on the order of the overall modulator. The cascade structure has the advantage that if each low-order loop is stable, the entire modulator is as well. Because the stability of first- and second-order loops can be guaranteed with proper selection of feedback coefficients and input range, stability concerns are avoided in high-order cascaded modulators. A disadvantage is that digital filtering is used to cancel low-order shaped noise, and if the noise transfer functions implemented in the analog domain deviate significantly from the ideal, the cancellation will be imperfect and low-order shaped noise will "leak" into the modulator output. Single-loop modulators do not experience this effect, so they are much more robust against circuit non-idealities. However, stability of a useful high-order single-loop modulator is impossible to guarantee formally, and requires extensive simulations and heuristic design

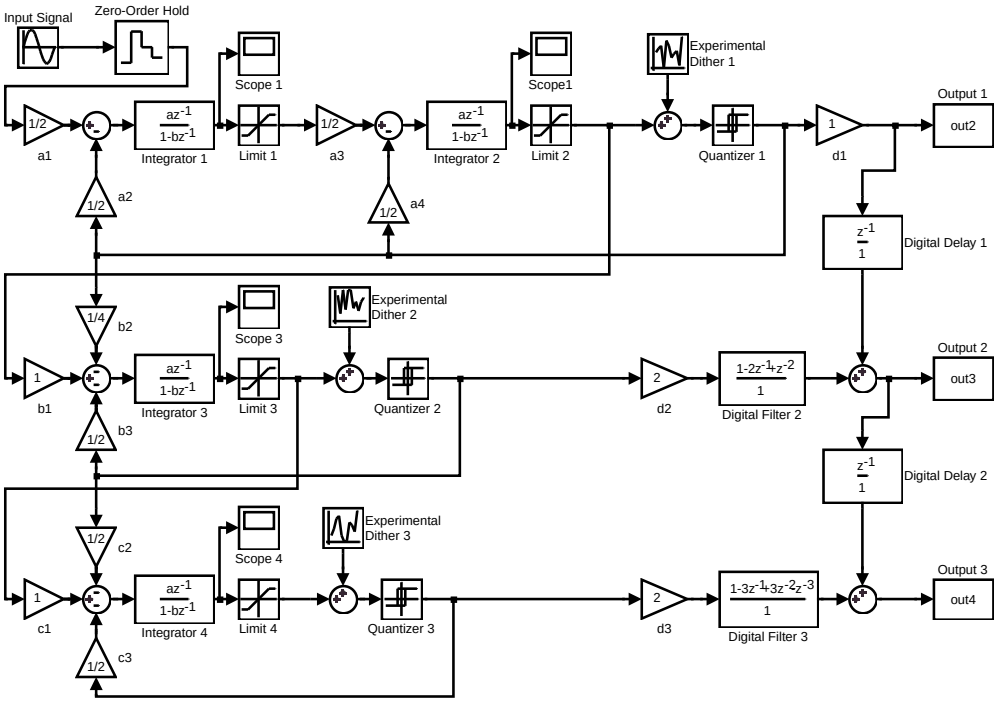

Fig. 1. Simulink model of modulator.

techniques. In addition, for stability concerns, high-order noise transfer functions must include poles to roll off the gain at high frequencies; thus, a high-order pure differentiator cannot be implemented, and the SNR will be degraded from what (1) predicts. The degradation required is much more significant at lower OSRs, and is almost prohibitive for our desired OSR of 16. A study in [8] demonstrated that even an eighth-order single-loop modulator can only achieve about 65 dB maximum SNR.

Because of the stability headaches involved with single-loop modulator design, as well as the higher order required for our low OSR, a cascade topology seemed natural for our application. The final topological decision was the selection of low-order modulators to cascade. It is typical to use a second-order modulator as the first stage, for mismatch concerns. With a second-order initial stage, integrator gain and pole error can cause leakage of second and first-order shaped noise, respectively. On the other hand, a first-order initial stage can cause leakage of first-order and non-shaped noise, which is worse. Two possible topologies were 2-1-1 and 2-2. We chose 2-1-1 because first-order sections are easier to design and scale, and it gave us an opportunity to directly view the second, third, and fourth-order noise shaping in our circuit.

III. System-Level Modeling

The Simulink model of the 2-1-1 modulator structure is shown in Fig. 1. Note that both of the integrators in the second-order loop have a delay. Although this is not a "textbook" second-order loop, it is easier to realize with switched-capacitor circuitry and results in a fully pipelined structure. Setting all of the coefficients to 1, except for a_4 which is set to 2, results in a valid modulator. That is, this set of coefficients ensures that the quantization error from the first and second stages is completely cancelled at the modulator output. Unfortunately, this configuration is not suitable for real-life implementation, because it requires the outputs of the some of the integrators to swing several times the quantizer output voltage. Thus, in order to equalize the dynamic range at the integrator outputs, coefficient scaling is required.

Because of the extreme non-linearity of the one-bit quantizer, scaling is not a trivial task. A one-bit quantizer has the property that the signal gain that comes before it is arbitrary, because the quantizer only detects the signal polarity. The $\frac{\Delta^2}{12}$ noise power approximation that is commonly used in data converter calculations only holds if the quantizer input is uniformly distributed between $\pm\frac{\Delta}{2}$. It can be derived (see Appendix) that if the input is distributed between $\pm V_{qi}$ and the quantizer output is ± 1, then the power of the quantization error is $1 - V_{qi} + \frac{1}{3}V_{qi}^2$, and the distribution is not uniform for $V_{qi} \neq 1$. This shows that the pre-quantizer signal gain affects the quantization error, which is important if that error is forwarded to the next cascade stage for subsequent cancellation. One approach to this problem is the *unity gain approximation* [2], which can be interpreted as inserting a "phantom gain" before the linearized model of the quantizer. This gain is adjusted such that the total gain around the outermost feedback loop of the modulator is unity. It should be noted that this technique has little rigorous justification, and simply produces analytical results that agree with simulation.

Solving the linear model of the second-order loop in terms of the coefficients leads to the following constraints:

$$a_2 = a_1 \tag{2}$$

$$a_4 = 2a_1a_3 \tag{3}$$

$$a_5 = \frac{1}{a_1a_3} \tag{4}$$

where a_5 is the phantom gain in front of the first quantizer, and a_1 and a_3 can be chosen to adjust the dynamic ranges at the output of the integrators. The initial choice for b_1 would then be a_5, so the second stage in the cascade sees the first stage's phantom gain correctly. Coefficients b_1 and b_2 can be scaled together, as long as d_2 and d_3 are scaled inversely; this effectively pushes a scale factor into the digital domain. Furthermore, b_1, b_2, and b_3 can all be scaled together, which affects the phantom gain of the second stage, and thus c_1. Scaling of the third stage can proceed similarly. Again, however, simulations are required in order to see the true effect of coefficient scaling. For example, it was found that in some configurations, b_2 could be scaled almost arbitrary, or even set to zero, with only minimal impact on the modulator SNR.

We selected $a_1 = a_2 = a_3 = a_4 = \frac{1}{2}$ because this selection of coefficients nicely bounds the integrator outputs to ± 1 for a reasonable input range, and the switched-capacitor circuit can be implemented with only one sampling capacitor per integrator [1]. This provides a phantom gain a_5 of 4. Half of this is pushed into the digital filter, setting d_2 and d_3 to 2 and leaving b_1 as 2 and b_2 as $\frac{1}{2}$. The inputs to the second stage are then scaled as well, reducing b_1 and b_2 further to 1 and $\frac{1}{4}$, respectively, and setting b_3 to $\frac{1}{2}$. The second stage's phantom gain, and thus c_1, is then 2. The inputs to the third stage are then scaled, leaving c_1, c_2, and c_3 as 1, $\frac{1}{2}$, and $\frac{1}{2}$, respectively. This affects the third stage's phantom gain, but this has no effect on the output because it is the last stage in the cascade.

This scaling procedure was guided at every step by Simulink simulations of the dynamic range at the output of each integrator, and was verified at every step by SNR calculations on the simulated data. Note that, because d_3 is 2, the modulator output is of the form

$$Y(z) = z^{-4}X(z) + 2(1 - z^{-1})^4 E_3(z) \tag{5}$$

That is, the uncancelled, fourth-order shaped quantization noise is amplified by a factor of two, and thus the SNR is 6 dB worse than that of the unscaled modulator.

IV. Circuit-Level Modeling

Modeling the circuit aspects of the $\Sigma\Delta$ modulator design is crucial to figuring out whether the modulator is practical in a particular technology. Several different circuit-level implementation details are considered, such as thermal noise, slew rate limitations, and power consumption.

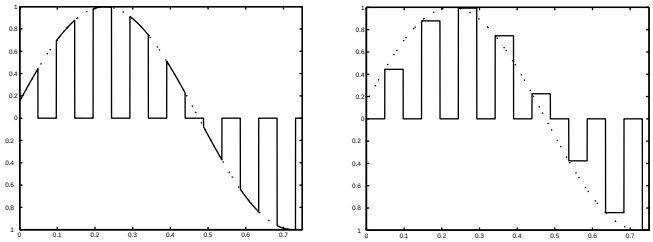

(a) (b)

Fig. 2. (a) Tracking phase. (b) Holding phase.

A. Thermal Noise

Thermal noise in sampling systems can be modeled as coming from a tracking phase and a holding phase [6]. The tracking phase can be modeled as a square wave multiplied with the signal, while the holding phase can be modeled as a sampling of the signal convolved with a sinc function (see Fig. 2). The output noise power due to multiplication with a periodic signal is simply the average power of the periodic signal multiplied with the input noise power. Thus, it can be shown that the noise component coming from the tracking phase is negligible when compared to the noise component coming from the holding phase due to the noise folding during sampling. For this reason, only the noise component due to the holding phase will be considered in our calculations.

The main noise sources in a $\Sigma\Delta$ modulator are the amplifiers and switches. Noise coming from blocks after the first integrator will be differentiated when referred back to the input, thus mitigating their effects on the overall SNR. Hence, the noise of the overall modulator will be dominated by the noise of the first integrator.

In this particular modulator design, the feedback path and signal path of the first integrator were chosen to have equal gains, as explained in Section III. As shown in [1], this allows for the use of a single-branch integrator which shares the sampling capacitor (see Fig. 3).

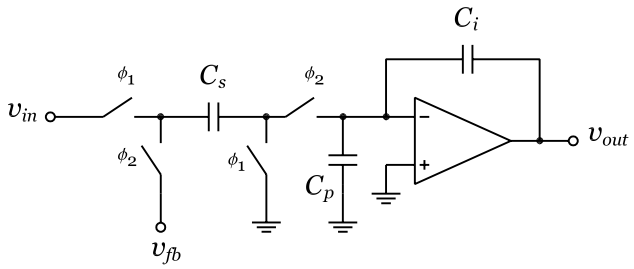

Fig. 3. Single-branch integrator.

This reduces the number of switches used, thus also reducing the thermal noise contribution due to the switches. The overall noise power can be calculated [6] to be

$$P_{th} = \frac{kT}{4MC_s} + \frac{kT}{2M(C_i + C_p)}\left(\frac{1}{3} + g_m R_{on}\right) \tag{6}$$

where R_{on} is the resistance of the switch when closed, M is the oversampling ratio, C_s,C_i,C_p are the sampling capacitance, integrating capacitance, and parasitic capacitance at the feedback node, respectively. It should be noted that this technique of sharing sampling capacitors cannot be used for the integrators of the cascaded stages because of the third input as shown in Fig. 1. This is not a problem for the noise calculations since their noise contribution is negligible when compared to the noise of the first integrator.

B. Charge Feedthrough

Using simple charge conservation, the amount of charge feedthrough in each integrator can be calculated. This effect manifests itself as a voltage spike in the opposite direction of the final step. The magnitude of this spike is

$$\frac{C_i C_o}{C_o\left(C_i + C_p\right) + C_o(C_i + C_p)} V_i \tag{7}$$

where V_i is the magnitude of the input step. Because of the opposite step, charge feedthrough lengthens the time required for the integrator to settle.

C. Slew Rate Limitations

The finite current that the amplifiers can source or sink results in a circuit nonlinearity known as slewing. The slew rate is defined to be the largest change in voltage per time. Slewing affects the settling time by putting a lower bound to the time required for the output of the integrators to reach a certain level, which in an ADC should be less than half of an LSB. An expression for the settling time that takes slewing into account is (see Appendix for derivation)

$$t_{settle} = \frac{C_L^*\left(V_o - \frac{1}{F} v_{dsat}\right)}{I_{slew}} + \frac{C_L^* v_{dsat}}{F I_{slew}} \ln\left(\frac{v_{dsat}}{\epsilon V_o F}\right) \tag{8}$$

where C_L^* is the total load capacitance seen at the output of the integrator, V_o is the step size of the output, v_{dsat} is about 150 mV for a typical process, ϵ is the fraction that the integrator has to settle to (about 0.0005 for a 10-bit converter), and F is the feedback factor given by

$$F = \frac{C_i}{C_i + C_s + C_p} \tag{9}$$

D. Power Estimation

Given the maximum signal power of P_{sig}, for a dynamic range of DR, we require the total noise power P_{noise} to be

$$P_{noise} = 10^{-\frac{DR}{10}} P_{sig} \tag{10}$$

where

$$P_{sig} = \frac{\left(\beta V_{dd}\right)^2}{8} \tag{11}$$

and β is the ratio of the maximum signal swing to V_{dd}. As a first-order estimate to the power consumption of this modulator, we will overestimate the total noise power by approximating (6) by

$$P_{th} \approx \frac{kT}{MC_s} \tag{12}$$

since $g_m R_{on} \ll 1$ (switch time constant is usually much less than amplifier time constant—0.1 or less is a typical value)[6]. This also assumes that C_s, C_i, and C_p are of the same order, which they are since the gains in our modulator are either 1, $\frac{1}{2}$ or $\frac{1}{4}$. Equating (10) and (12), we get a minimum bound for the sampling capacitance,

$$C_{s,min} \geq \frac{kT 10^{\frac{DR}{10}}}{M P_{sig}} \tag{13}$$

For a 10-bit converter with an OSR of 16, we find the minimum capacitance, $C_{s,min}$, to be about 15 fF. The wiring capacitance of the circuit will easily surpass this value; hence, the actual power estimate will depend mostly on the wiring capacitance. Noting that the settling time is 1.4 ns for our required 20 MS/s, and that the accuracy is 0.0005, we can use (8) to calculate the estimated power of each integrator. We suppose that the total capacitance at the output, C_L^*, is about 500 fF, leading to an estimated power consumption of 10 mW per integrator, when assuming a folded cascode topology for the integrators. This also corresponds to a transconductance, g_m, of 15 mS, and I_{slew} of 2 mA, with an assumed 2.5-V supply. Based on simulations in Simulink, we used the signal swing factor, $\beta = 0.3$. Thus, our overall modulator power should consume about 40 mW. Any increase or decrease in the total capacitance, C_L^*, will increase or decrease the power consumption proportionally.

E. Thermal and Quantization Noise Comparison

Using (12) and $\frac{\delta^2}{12}$ where δ is the LSB of the overall converter, we can find the thermal to quantization noise ratio of the converter

$$\frac{P_{th}}{P_{quant}} = \frac{12kT}{MC_s\delta^2} \tag{14}$$

which indicates that the thermal noise of this converter is about 30 dB less than the quantization noise. This implies that this converter is quantization noise limited. Compared to the typical use of $\Sigma\Delta$ modulators in high-resolution low-bandwidth applications, low-resolution high-bandwidth modulators tend to have simpler integrator designs because of this relaxed circuit noise requirement. This is due to the latter's lower settling accuracy and the assumption that both converters have about the same sampling frequency.

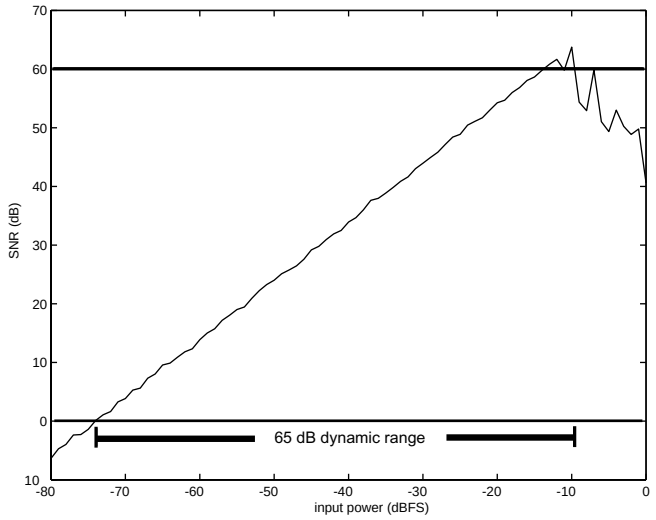

Fig. 4. Dynamic range simulation.

We can also view this as a tradeoff between the OSR and the overall converter resolution. From (1), we see that

$$DR \propto M^{2L+1} \tag{15}$$

and so

$$P_{quant} \propto \delta^2 \propto M^{-(2L+1)} \tag{16}$$

Thus, combining (14), (15), and (16), we see that the decreased oversampling ratio of the high-bandwidth $\Sigma\Delta$ modulators leads to a decrease in the thermal noise's ratioed contribution.

V. Results

The final modulator design was simulated in Simulink with limiters at the output of each integrator and no dithering. It was found that the modulator achieves the target SNR of 60 dB for an input signal amplitude of one-fifth full scale, or -17 dBFS signal power. This implies a noise power of $-17 - 60 = -77$ dB. Hand calculations on the ideal modulator predicted an SNR of 80 dB for a -3 dBFS input power. Recalling that our modulator has 6 dB more noise than the ideal due to coefficient scaling, hand calculations predict a noise power of $-3 - 80 + 6 = -77$ dB as well.

Fig. 4 shows the modulator SNR as a function of the input power. The SNR is positive for input powers of -74 dB and greater. The SNR is greater than 60 dB for input powers between -17 and -9 dB, after which it drops off as the integrator outputs overload. The achieved dynamic range is thus $-9 - (-74) = 65$ dB, which meets the application specifications. Fig. 5a shows the modulator output for a sample input after second, third, and fourth-order noise-shaping.

The modulator was then implemented in Spectre in order to model circuit-level effects. The switches were implemented as relays in Spectre to simulate ideal

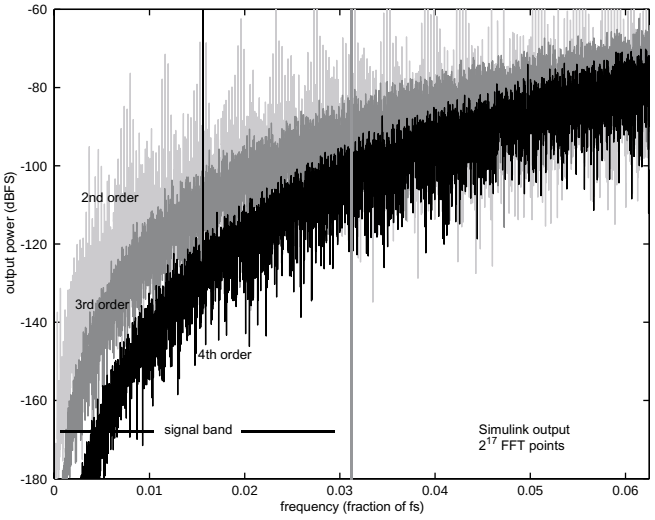

(a)

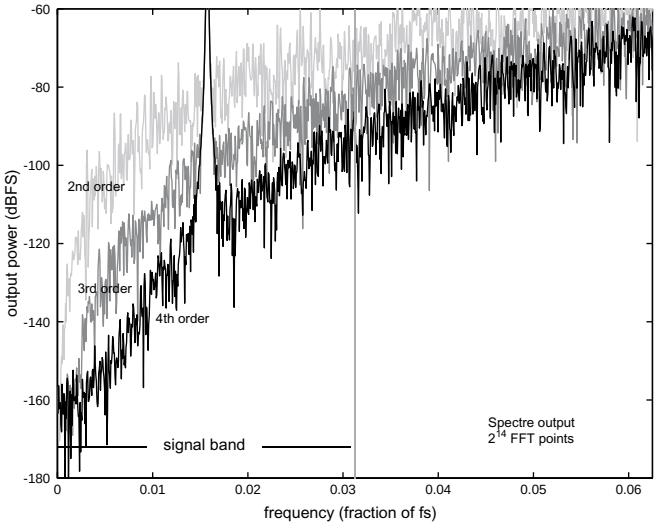

(b)

Fig. 5. (a) Modulator output from Simulink. (b) Modulator output from Spectre.

switches. Slewing and finite transconductance were also modeled in the OTAs by implementing the amplifiers using Verilog-A models. Even with slewing and finite transconductance, using values calculated in Section IV, the $\Sigma\Delta$ modulator still achieves the required dynamic range for 10 bits of resolution. Fig. 5b shows the modulator output for a sample input after second, third, and fourth-order noise-shaping.

Conclusions

A high-bandwidth $\Sigma\Delta$ modulator for high-data-rate wireless applications was designed, modeled, and simulated. Based on the results from Simulink and Spectre, we have demonstrated that the use of $\Sigma\Delta$ modulators in these high-bandwidth applications leads to relatively simple design specifications for the analog circuitry, specifi-

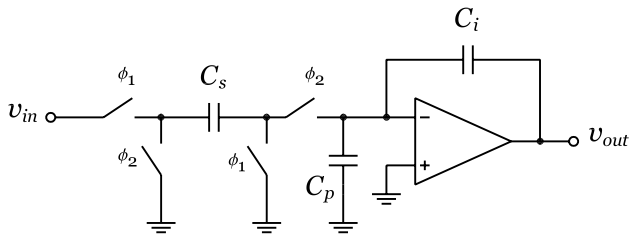

Fig. 6. Switched-capacitor integrator.

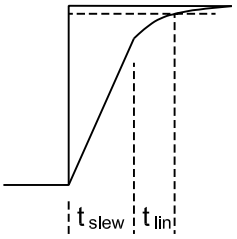

Fig. 7. Output step showing slewing.

cally the amplifier. In addition, we found that these lower-resolution high-bandwidth converters tend to be quantization noise limited. As a result, $\Sigma\Delta$ modulators can be a viable solution for integrated wireless applications where lower resolutions are acceptable.

Appendix

A. Derivation of quantization noise power

Assume the quantizer input is uniformly distributed between $\pm V_{qi}$ and the quantizer output is ± 1. The pdf distributions of the quantization error for V_{qi} less than, equal to, and greater than 1 respectively are:

$$\mathrm{pdf}_{V_{qi}<1}(x) = \begin{cases} \frac{1}{2V_{qi}}, & -1 \leq x \leq -1+V_{qi} \\ \frac{1}{2V_{qi}}, & 1-V_{qi} \leq x \leq 1 \\ 0, & \text{otherwise} \end{cases} \tag{17}$$

$$\mathrm{pdf}_{V_{qi}=1}(x) = \begin{cases} \frac{1}{2}, & -1 \leq x \leq 1 \\ 0, & \text{otherwise} \end{cases} \tag{18}$$

$$\mathrm{pdf}_{1<V_{qi}<2}(x) = \begin{cases} \frac{1}{2V_{qi}}, & -1 \leq x < -V_{qi}+1 \\ \frac{1}{V_{qi}}, & -V_{qi}+1 \leq x \leq V_{qi}-1 \\ \frac{1}{2V_{qi}}, & V_{qi}-1 < x \leq 1 \\ 0, & \text{otherwise} \end{cases} \tag{19}$$

The power of the noise in all three of the cases is:

$$\int_{\infty}^{-\infty} x^2 pdf(x)dx = 1 - V_{qi} + \frac{1}{3}V_{qi}^2 \tag{20}$$

In the $V_{qi} = 1$ case, the power is $\frac{1}{3}$, which is indeed $\frac{\Delta^2}{12}$ for $\Delta = 2$.

B. Derivation of slew rate

We assume that the amplifiers in the integrators are implemented by operational transconductance amplifiers (OTAs) since they will be driving mainly capacitive loads. A basic model of the switched-capacitor integrator can be seen in Fig. 6.

A simplified model of slewing can be seen in Fig. 7. The amplifier will be slewing as long as the required current draw for linear settling is larger than the maximum available current, or slew current, I_{slew}. For an output voltage step, V_o, the maximum rate of change of voltage for linear settling is V_o/τ, where $\tau = C_L^*/(Fg_m)$ and F is defined in (9). The OTA will slew until

$$\frac{V}{\tau} \leq \frac{I_{slew}}{C_L^*} \tag{21}$$

where V is the voltage difference between the current output of the integrator and the final output level. Equality in (21) determines V, or the remaining voltage at the junction where the OTA switches from slew-rate limited settling to linear settling. Therefore, we find

$$t_{slew} = \left(V_o - \frac{I_{slew}}{Fg_m}\right) / \left(\frac{I_{slew}}{C_L^*}\right) \tag{22}$$

Requiring an accuracy of ϵ, we use

$$V\left(1 - e^{-t_{lin}/\tau}\right) = V - \epsilon V_o \tag{23}$$

and get

$$t_{lin} = \tau \ln \frac{V}{\epsilon V_o} \tag{24}$$

Using $g_m = I_{slew}/v_{dsat}$ since half of the tail current is flowing through each transistor in a differential pair,

$$t_{settle} = \frac{V_o - \frac{1}{F}v_{dsat}}{I_{slew}/C_L^*} + \frac{C_L^* v_{dsat}}{FI_{slew}} \ln\left(\frac{v_{dsat}}{\epsilon V_o F}\right) \tag{25}$$

References

[1] B. Boser and B. Wooley, "The Design of Sigma-Delta Modulation Analog-to-Digital Converters," *IEEE Journal of Solid-State Circuits*, vol. 23, pp. 1298–1308, 1988.

[2] S. Rabii and B. Wooley, *The Design Of Low-Voltage, Low-Power Sigma-Delta Modulators*. Kluwer Academic Publishers, 1999.

[3] D. Sobel, "A 25 Ms/s Nyquist-Rate Sigma-Delta Modulator for a Wideband CDMA Receiver," *Master's Thesis*, 2000.

[4] S. Ardalan and J. Paulos, "An Analysis of Nonlinear Behavior in Delta-Sigma Modulators," *IEEE Transactions on Circuits and Systems*, vol. CAS-34, pp. 593–603, 1987.

[5] P. Aziz, H. Sorensen, and J. Van Der Spiegel, "An Overview of Sigma-Delta Converters," *IEEE Signal Processing Magazine*, January 1996.

[6] F. Medeiro, A. Pérez-Verdú, and A. Rodríguez-Vázquez, *Top-Down Design of High-Performance Sigma-Delta Modulators*. Kluwer Academic Publishers, 1999.

[7] *ed.* S. Norsworthy, R. Schreier, G. Temes, *Delta-Sigma Data Converters*. IEEE Press, 1997.

[8] R. Schreier, "An Empirical Study of High-Order Single-Bit Delta-Sigma Modulators," *IEEE Transactions on Circuits and Systems–II: Analog and Digital Signal Processing*, vol. 40, pp. 461–466, 1993.

[9] D. Johns and K. Martin, *Analog Integrated Circuit Design*. Wiley Press, 1997.

[10] D. Yee, *PhD Dissertation*. UC Berkeley, 2001.

[11] IEEE Standards Board, "Part 11: Wireless LAN Medium Access Control and Physical Layer specifications," *IEEE*, 1999.