Analog Integrated Circuits and Signal Processing, 30, 69–81, 2002

An 8.25-MHz 7th-Order Bessel Filter Built with Single-Amplifier Biquadratic MOSFET–C Filters

HANSPETER SCHMID

Bernafon AG, Morgenstrasse 131, 3018 Bern, Switzerland
E-mail: h.p.schmid@ieee.org

Abstract. This paper is a practical guide to building higher-order filters with single-amplifier biquadratic MOSFET–C sections. Theory, design guidelines, and measurement electronics are discussed by example of a 7th-order current-mode filter built to the specifications of a 1× DVD read channel filter. The 7th-order filter was fabricated with the double-poly 0.6-micron CMOS process by AMS. It is continuously tunable from 4.5 MHz up to 10 MHz, covers a chip area of only 0.24 mm^2, and consumes 49 mW from a 3.3-V supply. The SNR at −40 dB of harmonic distortion is between 48 dB and 50 dB over the whole tuning range. The comparatively low power consumption and chip area could be achieved by using single-amplifier biquadratic building blocks implemented as MOSFET–C filters and generating the control voltage of the MOSFET resistors with an on-chip charge pump. The technique is, with a small loss of SNR, also applicable on fabrication processes where only gate-oxide capacitors are available.

Key Words: MOSFET-C filter, single-amplifier biquad, current-mode filter, current conveyor, charge pump, integrated filter, continuous-time filter

1. Introduction

Discrete-component single-amplifier biquadratic filters (SABs) have long been used in the industry: they are cheaper and more power-efficient than integrator-connected filters because they require only one amplifier to generate a pair of complex poles. This advantage has to be paid with a higher variance of the pole Q, which is acceptable in most low-Q and medium-Q filters [16,27].

In spite of their advantages in terms of power consumption and cost, SABs have seldom been used on integrated circuits. The reason is that their pole frequency depends on the values of passive components only, which means that in order to continuously tune the pole frequency, adjustable passive components have to be used. The simplest way to achieve this is to build the SAB as a MOSFET–C filter by replacing the resistors with transistors operating in the triode region.

This paper is a tutorial that demonstrates how MOSFET–C SABs can successfully be used to build higher-order video-frequency filters. The filter described here was built to verify the theory developed in [25], where a comprehensive discussion of most theoretical and practical aspects of MOSFET–C SABs including detailed descriptions of the measurements can be found.

In the following sections, we describe the ideal and non-ideal transfer functions of a Sallen-and-Key low-pass filter (Section 2), the structure of a second-order MOSFET–C feedback network (Section 3), a suitable video-frequency current amplifier (Section 4), a self-oscillating charge pump used to increase the dynamic range by driving the MOSFET resistors with a gate voltage above V_{DD} (Section 5), a 7th-order Bessel filter (DVD read channel filter) built with MOSFET–C sections (Section 6), measurement circuits on and off chip (Section 7), the main measurement results and a comparison with other filters (Sections 8 and 9), and a discussion of design trade-offs (Section 10).

2. Single-Amplifier Biquadratic Filters (SABs)

Fig. 1 shows one way to build a low-pass SAB, a so-called Sallen-and-Key filter. The component values are given in terms of the geometric means R and C and the component spread factors n and m. This simplifies

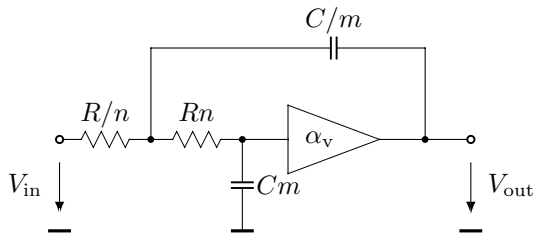

Fig. 1. A voltage-mode low-pass SAB.

the analysis and is also closer to the reality on the IC, because variations of R and C come from process variations, whereas variations of m and n come from mismatch.

Ideally, the filter in Fig. 1 has the following transfer function:

$$T(s) = \frac{\alpha_V \omega_p^2}{s^2 + \frac{\omega_p}{q_p} s + \omega_p^2}, \tag{1a}$$

$$\omega_p = \frac{1}{RC}, \qquad \frac{1}{q_p} = mn + \frac{m}{n} + \frac{1-\alpha_V}{mn}, \tag{1b}$$

where ω_p is the pole frequency in rad/s and q_p is the pole Q. Only R and C appear in the formula for ω_p, so ω_p varies because of process variations (and temperature ...) and must be tuned, either by adjusting C, or by adjusting R. Only m, n, and α_V appear in the expression for q_p, therefore the variance of q_p depends only on matching.[1] The q_p normally does not need to be adjusted if it is low or moderate ($q_p < 5$), but if necessary it can be done by adjusting α_V, the gain of the amplifier.

A real amplifier will have a certain input capacitance, an output resistance, and a phase lag. All of them cause pole shifts, but the non-zero output resistance also introduces a pair of complex zeros that causes the transfer function to rise to a certain level for higher frequency and thus limits the achievable stopband attenuation [26,29]. This behaviour is shown using idealised transfer functions in Fig. 2 for different sets of input capacitances and output resistances, and for the basic band-pass and high-pass filters as well as for the low-pass filter from Fig. 1 (see [26] for more details).

It follows from the discussion in [26] that if the stopband attenuation and a voltage amplifier with a certain C_{in} and R_{out} are given, the maximum achievable pole frequency becomes

$$\omega_{p\,max} \leq \frac{1}{\max(m, 1/m)\, C_{in} \cdot \max(n, 1/n)\, R_{out} \cdot A_{stop}}, \tag{2}$$

which is largest for $m = n = 1$. Note that equality in (2) occurs only if C_4 consists of the amplifier input capacitance without an additional external capacitance. This may cause an unacceptably high q_p variance, because the precision of q_p then relies on the matching of a poly-poly capacitor and a parasitic capacitance. It

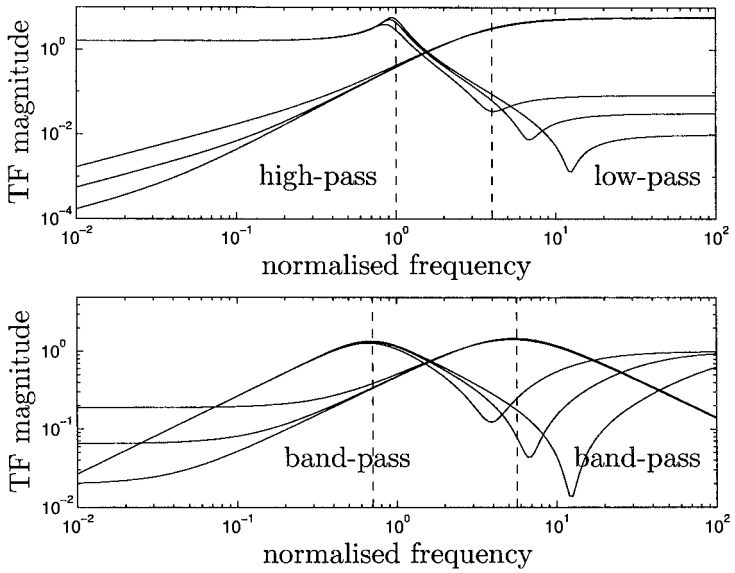

Fig. 2. Ideal transfer functions (TF) of the four basic Sallen-and-Key filters. The dashed lines indicate the different pole frequencies that were used to make all lines visible.

should, however, not be a problem for low-q_p sections such as the ones used in this paper. The output capacitance is of course non-linear, but simulations made in [25] show that this non-linearity scarcely affects the linearity of our video-frequency filter. However, it would become a problem in low-distortion low-frequency filters.

Note that in order to minimise the variance of q_p, the component spreads should generally be as large as possible, and not $m = n = 1$, as was shown in [27]. Therefore there is a trade-off between the maximum pole frequency that can be achieved with a given amplifier and the variance of q_p.

It was also shown in [27] that a Sallen-and-Key low-pass filter with minimum q_p variance always needs an amplifier with $\alpha_V < 2$. This is good because it means that the low gain α_V can be derived from components of similar size that can be made to match well.

3. Second-Order MOSFET–C Filter

The SAB from Section 2 can now be converted into a MOSFET–C filter by building it in a balanced form and replacing the resistors by transistors operating in the triode region, as shown in Fig. 3. The discussion in the previous section is also valid for a current-mode filter, with two changes. First, C_{in} and R_{out} become C_{out} and R_{in}; and second, if a current flowing into the output of the current amplifier is considered positive, then $\alpha_I = -\alpha_V$.

We have chosen to transpose the filter into a current-mode filter [17] because it then has a lower power consumption, at the expense of more harmonic distortion. We will now briefly explain why this is so. The amplifier will not require a gain above two, as we mentioned in Section 2. There are two possibilities to build an amplifier with such low a gain: either a high-gain amplifier (opamp) is used whose gain is then set to the desired low value by feedback resistors, or a low-gain amplifier without feedback is used. The former can be built easily with an opamp, where "opamp" can be a voltage opamp or any of the other eight types of opamps, e.g., a current opamp or a current-feedback opamp [24]. The most straightforward way to build an amplifier without gain-setting feedback is to use current mirrors, which results in a current-mode filter.

The amplifier gain is comparably precise in both cases, since it depends on the matching of two polysilicon resistors for the voltage-mode filter and on the matching of current-mirror transistors in the current-mode filter. Since the feedback is necessary to set the loop gain, the voltage-mode amplifier can not be used to build filters with a pole frequency that is higher than the opamp unity-gain frequency divided by five times the closed-loop gain [25]. Such a restriction is not present in the current-mode filter, which will operate up to higher frequencies or, alternatively, will use less power for the same pole frequency. On the other hand, feedback also makes the amplifier more linear, so using a current-mode filter saves power, but increases the harmonic distortion.

The output of our current amplifier consists of composite transistors, as shown in Fig. 4. The cascode transistors are biased with a voltage $V_m \pm \Delta V$, where V_m is the mid-rail voltage. Note that in the process we used, both nMOS and pMOS have the same threshold voltage $|V_{T0}| = 0.85\,\mathrm{V}$. Increasing ΔV will of course increase the maximum signal swing, but it will also decrease the voltage margin available for driving the current mirrors, and therefore it will increase the noise produced by the current mirrors. It can be shown that the ΔV that optimises the SNR for a given level of

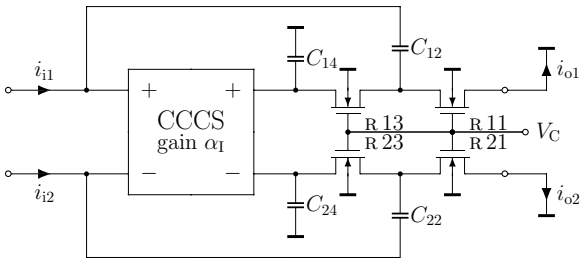

Fig. 3. Balanced-signal current-mode single-amplifier biquad containing a current amplifier with a negative gain α_I.

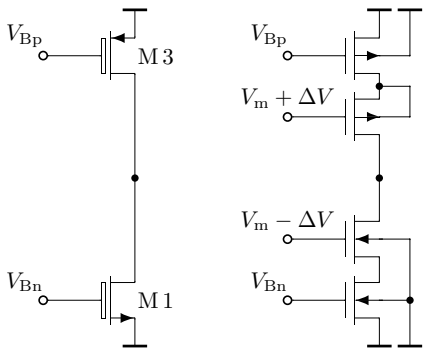

Fig. 4. Composite transistors (simple cascodes).

harmonic distortion (e.g., −40 dB) is [30]

$$\Delta V = \tfrac{2}{7} V_{dd} - V_{T0}, \tag{3}$$

which is only 0.09 V for $V_{T0} = 0.85$ V and $V_{dd} = 3.3$ V and can be set to zero in the implementation with a negligible loss of SNR.

Note that this discussion neglects the bulk effect of the *n*MOS cascode transistor. The increase of the *n*MOS threshold voltage, ΔV_T, can be taken into account by shifting V_m by $-\frac{1}{2}\Delta V_T$ and by replacing V_{T0} by $V_{T0} + \frac{1}{2}\Delta V_T$. However, since V_{SB} of the *n*MOS cascode transistor is between 200 mV and 300 mV, the necessary shift is very small, and not doing it in the implementation again results in a negligible loss of SNR.

The assumption leading to (3) is that the main source of distortion is the signal clipping that occurs when the cascode transistors at the output of the current amplifier leave the saturation region. This is indeed a good model of the reality if the gate control voltage of the MOSFET resistors is lifted above V_{DD} by a charge pump, as it is proposed in this paper. In MOSFET–C SABs without a charge pump, clipping introduced by saturating MOSFET resistors must be taken into account [25].

4. Video-Frequency Current Amplifier

MOSFET resistors are not very linear, but since the non-linearity is mainly of second order, the difference of the two balanced filter output signals is theoretically linear (c.f. Fig. 15 (a) for a typical spectrum with suppressed even-order harmonics). The current amplifier used in MOSFET–C SABs must therefore amplify the difference of its input currents, and its two paths should be identical. The amplifier whose block diagram is shown in Fig. 5 meets both criteria. Ideally, it

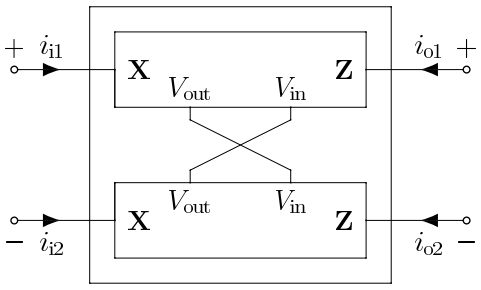

Fig. 5. Block diagram of a balanced current amplifier.

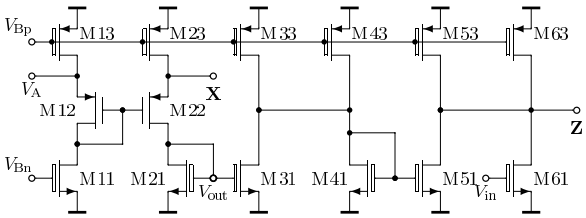

Fig. 6. Half circuit of the balanced current amplifier.

is described by

$$i_{o1} = \alpha_I (i_{i1} - i_{i2}) \qquad i_{o2} = -\alpha_I (i_{i1} - i_{i2}). \tag{4}$$

According to the discussion in Section 2, the output capacitance of the current amplifier should be low, its input resistance should be low, and it can also be shown that its phase lag should not exceed 10 . . . 20° at the pole frequency of the biquad to be built [25].

The half circuit of our current amplifier is shown in Fig. 6. It was developed from a circuit that is conventionally called class-A second-generation current conveyor [14, Chap. 11.5]. It consists of one voltage buffer and several current mirrors. M[1–6]3 and M11 are constant-current sources; M[2–6]1 form current mirrors. M22 is the input transistor. It provides, at its source, a current input with input resistance $R_{in} \approx 1/g_{m22}$. M12 is a voltage level shifter that sets the operating point voltage of node X to V_A ($V_A = V_m$ in our implementation). Any current flowing into X is mirrored from M21 to M31 and from M41 to M51 and flows out of Z; it is also mirrored from M21 to M61 of the other half-circuit, where it flows *into* Z. Thus the two input currents i_{i1} and i_{i2} are subtracted, and if all current mirrors have unity gain, the resulting gain is $\alpha_I = -2$. A different gain can easily be achieved by changing the width of all output transistors M[5–6][1,3].

Note that all transistors drawn with boxes as gates are composite transistors as explained in the previous section. All cascode transistors are biased by V_m, with one exception. The gate potential of the diode-connected transistor M12 is below $V_m - V_{T0}$, so the cascode transistor of M11 must be biased differently, namely by the voltage V_{bc} generated by transistor M81 of the bias circuit shown in Fig. 7. The main transistor of M11 does not necessarily need to be enlarged, since it only has to conduct the bias current, while all other *n*MOS composite transistors in the amplifier have to be able to conduct the bias current plus the maximum signal current. So in our design, M11 by coincidence has the same size as M[2–6]1.

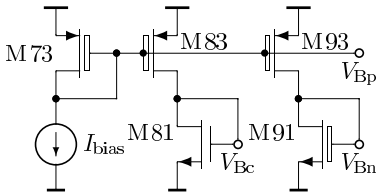

Fig. 7. Bias circuit.

In the composite transistors, the W/L ratio of the main transistor is about six times smaller than the W/L ratio of the cascode transistor. As was shown in [2], simple cascodes are fastest when the $V_{d\,sat}$ of the cascode transistor is about 40% of the $V_{d\,sat}$ of the main transistor. The factor of six results when the W/L ratios necessary to achieve the 40% are calculated from $V_{d\,sat}$ and I_d.

The actual transistor dimensions were found iteratively. First, we knew from experience that the bias current would have to be around 160 μA to achieve a maximum pole frequency around 16 MHz. The maximum signal current to be supported by the current amplifier was designed to be 60 μA, approximately the current at which the MOSFET resistors would saturate. This determined the sizes of all current source and current mirror transistors. The input transistors M[1–2]2 were sized such that they provide an X input resistance around 500 Ω, and then it was made sure that the cascode transistor in M11 would indeed remain in saturation by giving it a bias voltage 0.1 V below analogue ground, which determined the size of M81. This process was iterated several times until the simulated performance of the highest-Q biquad in the filter was satisfactory. The resulting transistor dimensions are given in Table 1.

Table 1. Transistor dimensions in the amplifier. The "×2" indicates common-centroid layout.

	Main Transistors	Cascode Transistors
M[1–6]1	45 × 1.8 μm	95 × 0.6 μm
M81	14 × 0.6 μm	—
M91	45 × 1.8 μm	95 × 0.6 μm
M[1–2]2	120 × 0.6 μm × 2	—
M13	87 × 1.8 μm	140 × 0.6 μm
M[2–8]3	70 × 1.8 μm	140 × 0.6 μm
M93	87 × 1.8 μm	140 × 0.6 μm

5. Self-Oscillating Charge Pump

It was shown in [25] that the signal swing of a MOSFET–C filter can be increased considerably[2] if the gate control voltage is lifted above the positive supply using a charge pump.

The charge pump shown in Fig. 8, whose element sizes are given in Table 2, combines features of a charge pump proposed in [6] with those of a five-inverter ring oscillator. It actually comprises two charge pumps. The main pump, consisting of M1, M4, M5, M6, C_1, and C_4, fills the reservoir capacitor C_0 with charge, where M5 and M6 alternatively conduct the charging current. A second pump driven by the same inverters, consisting of M2, M3, C_2, and C_3, sets the gate voltage of M5 and M6 to 2 V_{in} while they charge C_0. Thus the output voltage becomes $V_C \leq 2\,V_{in} - V_{T5}$, where M5's threshold voltage V_{T5} is comparatively large because of the bulk effect (we are using an n-well process). In our example, $V_C = 4.6$ V for $V_{in} = 3$ V. The charge pump operates properly for $V_{in} = 1.3 \ldots 3.3$ V, resulting in $V_C = 1.5 \ldots 5.0$ V.

The voltage ripple of this charge pump is smaller than that of a conventional charge pump by a factor

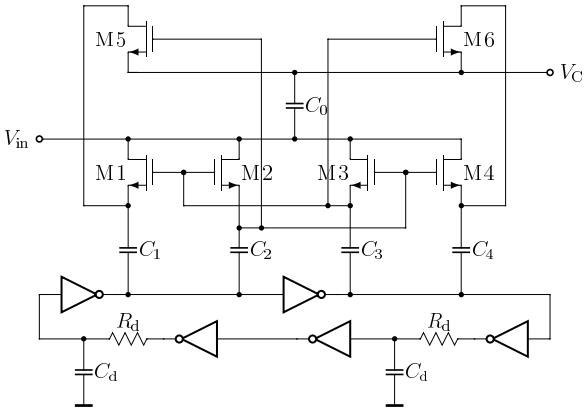

Fig. 8. Self-oscillating charge pump.

Table 2. Transistor and element dimensions in the charge pump.

All nMOSTs	10 × 0.6 μm
All pMOSTs	33 × 0.6 μm
R_d	4.8 kΩ
C_d, C_2, C_3	0.5 pF
C_1, C_4	1 pF
C_0	20.5 pF

of $g_{m5}/g_{ds5} \approx 30 \ldots 100$. As mentioned in [6], the voltage ripple becomes

$$V_{ripple} = \frac{1}{2} \cdot \frac{I_{out}}{C_0 f_{clk}} \cdot \frac{g_{ds5}}{g_{m5}}, \tag{5}$$

where I_{out} is the DC current flowing out of the reservoir capacitance C_0 and f_{clk} is the pump's clock frequency. This means that if V_{ripple}, I_{out} and f_{clk} are the same, the two-stage charge pump needs a reservoir capacitor which is 30 ... 100 times smaller than the one in a conventional charge pump.

The oscillator should have an oscillation frequency that is far above the edge frequency of the filter. Because the inverters need to deliver only small currents, they can be built with small transistors. Using only inverters would result in an oscillation frequency of almost 1 GHz, thus two passive one-pole low-pass filters, each consisting of one poly-silicon resistor and one poly-poly capacitor, were used to slow the oscillator down to simulated 90 MHz. This has the additional advantage that it reduces the temperature dependence of f_{clk}. The measured oscillation frequency was between 62 MHz and 71 MHz for a charge-pump supply voltage going from 2.7 V to 3.3 V, which is within the range given by the tolerances of poly-silicon resistors and double-poly capacitors on the IC process we used.

6. 7th-Order Bessel Filter

Using the biquads discussed above, the 7th-order Bessel filter shown in Fig. 9 was built. From left to right, the following building blocks can be seen: analog pads and power supply pads, poly resistors (four vertical gray lines) and the current amplifier of the on-chip V–I converter used for making the measurements (see Section 7 for details), a passive first-order low-pass filter with a normalised $f_p = 1.687$ that consists of a current amplifier and a first-order MOSFET–C low-pass filter, and three MOSFET–C SABs with $(f_p, q_p) = (2.053, 1.13)$, $(1.719, 0.53)$, and $(1.825, 0.66)$, in this order. This filter could, e.g., be used as a pulse equaliser in a 1× DVD read channel [12].

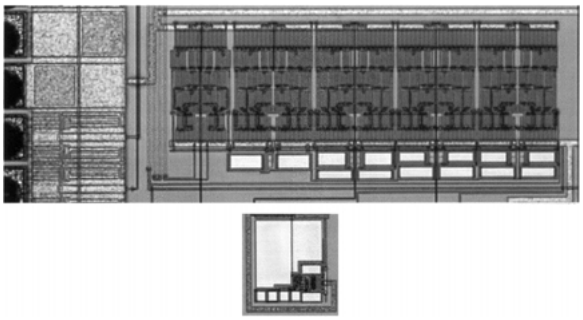

Fig. 9. Chip photos of the 7th-order filter and the charge pump (both photos have the same scale).

In order to maximise the SNR of a biquad cascade, the gains of the individual biquads should be set to unity to make the signal levels in all biquads equal. This will, however, increase the variance of the pole Q considerably compared to the pole-Q variance of the minimum-variance design. In our design, we decided to use a gain of two in all biquads, but measurements and simulations showed that it would have been sufficient to give the highest-q_p biquad a gain of two and the remaining stages a gain of one, which would have increased the maximum current before clipping occurs, and therefore the SNR of the filter, by a factor of three, or by 9.5 dB.

The passive component values used to build the 7th-order filter are shown in Table 3. Note that the signal capacitors are large compared to the output capacitance of the current amplifier, which is around 0.6 pF. It will be shown in Section 10 that the linearity of the filter is scarcely affected if no explicit signal capacitors are used, such that C[1–2]4 consist of parasitic capacitance only [25]. This would make it possible to increase the edge frequency of the 7th-order filter without using more power, but it would reduce the SNR by 4 dB because the noise bandwidth would increase, and it would increase the variance of all q_p because the latter depends on the matching of C[1–2]4 and C[1–2]2, which would become worse.

Table 3. Element values in the four sections of the 7th-order filter.

(f_p, q_p)		Element Dimensions	Capacitance
(1.687, –)	C[1–2]4	68.9 × 28 μm	1.68 pF
	R[1–2]3	12 × 8 μm	
(2.053, 1.13)	C[1–2]2	60 × 26.75 μm	1.40 pF
	C[1–2]4	72.2 × 16.6 μm	1.05 pF
	R[1–2][1,3]	12 × 6 μm	
(1.719, 0.53)	C[1–2]2	55.5 × 19.1 μm	0.93 pF
	C[1–2]4	72.2 × 22.2 μm	1.40 pF
	R[1–2][1,3]	12 × 10.5 μm	
(1.825, 0.66)	C[1–2]2	56.5 × 21.25 μm	1.05 pF
	C[1–2]4	72.2 × 19.5 μm	1.23 pF
	R[1–2][1,3]	12 × 9.5 μm	

7. Measurement Setup

When voltage-mode filters are measured, one normally uses output drivers to drive the chip pads, but the input of the filter can often be connected directly to the pads. For current-mode filters, it is the output that can be connected directly to the pads and the input that needs a driver.

For all measurements, the filters were driven by the single-ended–to–balanced voltage converter shown in Fig. 10. It has a 50-Ω input to which the generator of the network analyser can be connected, and it provides precisely balanced output voltages. It bases on the current-feedback (CFB) opamp AD 8002, which is well capable of driving high loads, so all circuits on the chip could be connected to the same two input pads. Actually, the circuit in Fig. 10 is an adaption of a differential line driver proposed in the data sheet of the AD 8002. The conversion gain is one, and because the network analyser expects a resistance of 50 Ω, there is only the 6-dB loss of the power splitter at the analyser output to be taken into account.

The common mode of the input voltages is so constant that on chip, both balanced voltages can simply be converted to balanced currents by a poly-silicon resistor connecting the pad to the input of a single-ended current amplifier (i.e., only transistors M11–M33 of Fig. 6). The bandwidth of this setup is not much higher than the bandwidth of the filter to be measured, but the chip from which the measurements in this papers were made contained one path with nothing else but the on-chip V–I converter, such that its transfer function could be calibrated out of the measurements.

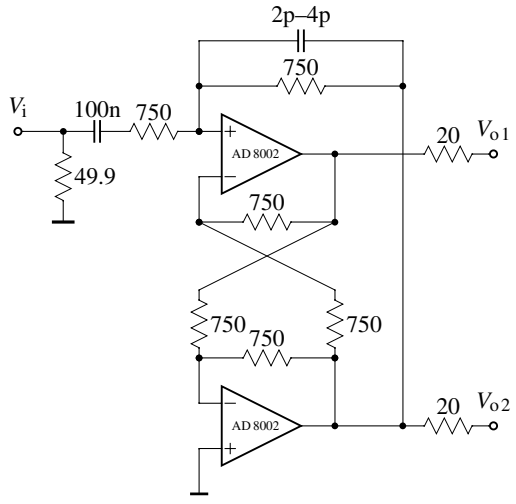

Fig. 10. Single-ended–to–balanced voltage converter using AD 8002 CFB opamps.

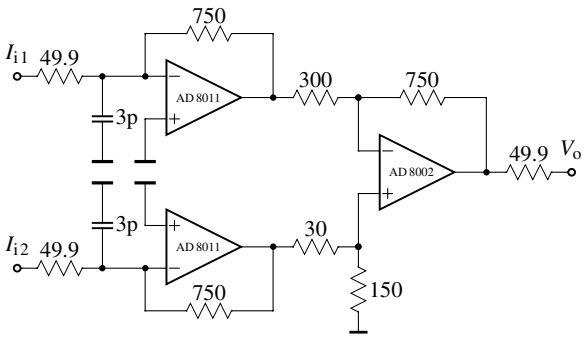

Fig. 11. Balanced-current–to–single-ended-voltage converter using AD 8002 and AD 8011 CFB opamps.

The current output of every test circuit was converted to a single voltage by the circuit shown in Fig. 11. It consists of two independent I–V converters based on the AD 8011 (another CFB opamp). The converters have an $R_m = 750\,\Omega$. The following stage is a difference amplifier based on the AD 8002 with a voltage gain of 5. Together with the differencing, the overall R_m from a single current output to the converter output is 7500 Ω. The reason why two different CFB opamps were used is that the AD 8011 is basically slower; because of the stability problems that often occur with these high-speed amplifiers, we decided not to use amplifiers that are faster than necessary.

All transfer functions and characteristics were measured with the 500-MHz network analyser HP 8751 A; the noise and clock feed-through were measured with the 150-MHz spectrum analyser HP 3588 A. For the harmonic-distortion measurements, a 2-V_{pp} signal was generated with a Tektronix AFG 2020 function generator and then attenuated by a programmable RF attenuator, the Marconi MA 2186, in order to produce a harmonically clean signal for the measurements.

8. Measurement Results

The 7th-order Bessel filter was integrated in the 0.6-μm double-poly triple-metal CMOS process by Austria Mikro Systeme, of which some parameters are shown in Table 4. Fifteen chips were produced, and

Table 4. Process parameters of the 0.6-μm CMOS process by AMS.

	*n*MOS	*n*MOS	
V_{T0}	0.85	−0.85	(V)
$\mu \cdot C_{ox}$	120	40	($\mu A/V^2$)
γ	0.8	0.5	($\sqrt{V}$)
ϕ_0	0.94	0.91	(V)

all worked fine. The inter-chip matching measured for a single biquad with $f_p = 24\,\text{MHz}$ and $q_p = 3$ was sufficiently good, the standard deviations of f_p and q_p were 1.5% and 3%, respectively. On-chip matching can be expected to be even better, but was not measured.

The measured performance of the 7th-order Bessel filter is summarised in Table 5. The values for power consumption and chip area include the charge pump, the chip area also includes the wiring around the filter block and a margin of a few μm in every direction.

The input signal level at which 1% THD is reached was determined for every setting of V_C by first determining the edge frequency of the filter for that particular V_C and then using an input signal of one fifth of that frequency such that the first five harmonics lie in the passband.

The clock feed-through to the filter outputs was measured as well (actually, this was how we determined the oscillation frequency of the charge pump), but it was scarcely visible over the noise floor (see curve (c) in Fig. 15) and changes the SNR of the filter by much less than 1 dB.

CMRR and PSRR were simulated and measured as well and are exceptionally good because of the perfectly balanced structures and the good matching of the signal paths, but we decided to not publish the values because they are really irrelevant. There is another effect in MOSFET–C filters that is much stronger than the linear transfer of a power ripple or common-mode signal: any common-mode signal at the source/drain of a MOSFET resistor *will be modulated by the signal*.

Table 5. Measured performance of the 7th-order filter.

Charge pump supply	2.7...3.3 V
V_C	4.4...5.0 V
Edge frequency	4.5...10 MHz
SNR for 1% THD	48...50 dB
Supply voltage	3.3 V
Power consumption	53 mW
Chip area	0.28 mm^2

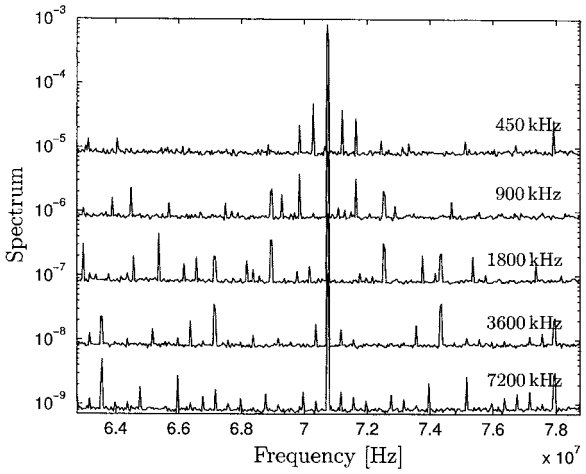

Fig. 12. Measurement of the intermodulation of signals with different frequencies with the power ripple caused by the charge pump (the spectrum units are chosen arbitrarily such that all peaks are visible).

To illustrate, Fig. 12 shows measurements of strong sine signals in the presence of a charge pump with a clock frequency slightly above 70 MHz. Clearly, the clock-feedthrough component dominates, but the mixing products of the signals with common-mode signals and power supply ripples caused by the charge pump clock can also clearly be seen. Thus the CMRR and PSRR will be non-linear and they will also depend on the nature of the differential-mode signals, the common-mode signals, and the power supply ripple signals. There is, to our knowledge, no standard way to quantify such non-linear CMRR and PSRR, all we can say is that our measurements showed that these effects were weaker than the linear clock feedthrough. Therefore, while we cannot say how small the CMRR and PSRR are, we still know that they are small enough.

9. Comparison with Other Filters

A figure of merit that is often used to compare filters is the power per pole and frequency as a function of the SNR at 1% THD. Fig. 13 shows this figure of merit for several filters published recently; Table 6 gives the references to where each filter can be found. The black circles denote nine CMOS filters working in the frequency range 5–50 MHz, three of them are MOSFET–C filters built by us, the other six are Gm–C filters. The gray circles are different filters. Note that Filters 15 and 25 are switched-capacitor filters, and

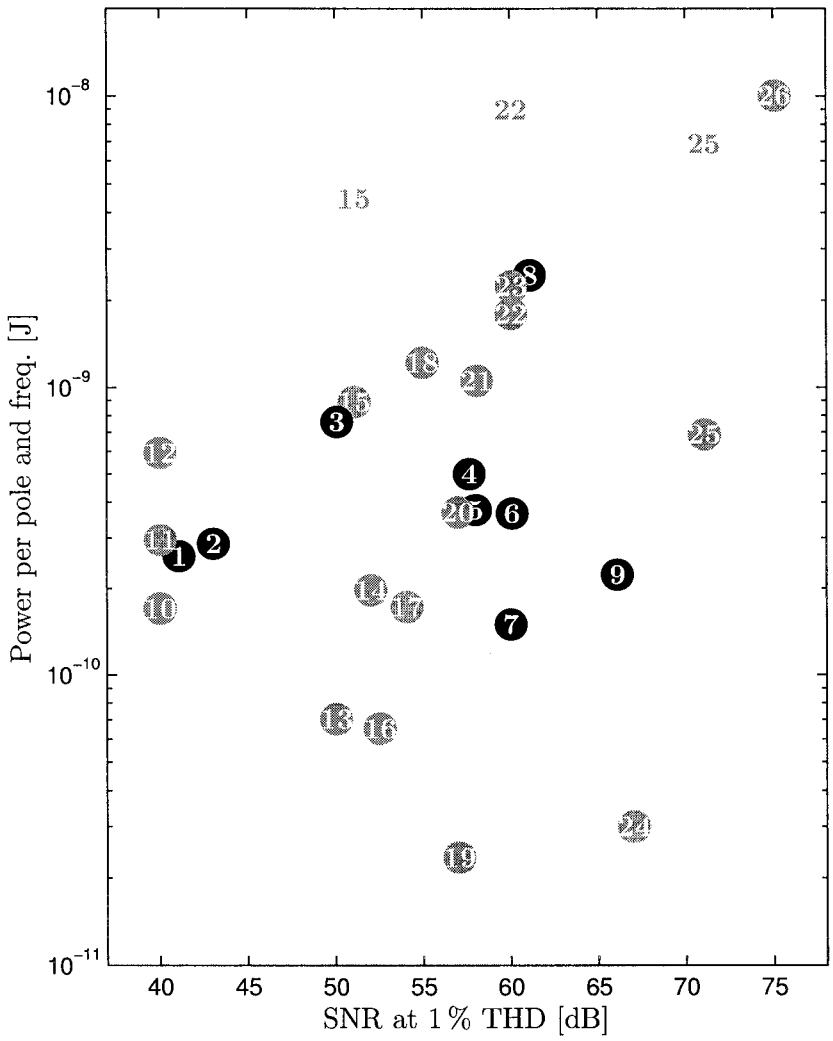

Fig. 13. Figure of Merit for different filters in the literature.

Table 6. References to the filters in Fig. 13.

1	**MOSFET–C**	8	[8]	15	[32]	22	[9]
2	[5]	9	**MOSFET–C**	16	[20]	23	[10]
3	**MOSFET–C**	10	[22]	17	[19]	24	[35]
4	[34]	11	[7]	18	[11]	25	[9]
5	[3]	12	[1]	19	[21]	26	[4]
6	[13]	13	[33]	20	[31]		
7	[15]	14	[23]	21	[18]		

Filter 22 is a switched-current filter. There are two entries for each of the switched filters, the gray number without a circle denotes the figure of merit for the pole frequency, and the gray circle stands for the sampling frequency. We included the three switched filters only to illustrate the well-known fact that switching costs power, and that the power per pole and sampling frequency of switched filters is comparable to the power per pole and frequency of continuous-time filters.

Several things can be seen in Fig. 13, e.g., that two filters, 19 and 24, lie far below the rest, but they use special technologies and techniques that have disadvantages which are not covered by the figure of merit: Filter 19 is a BiCMOS log-domain filter, and Filter 24 uses positive feedback. Both filters operate around 0.5 MHz.

We will now concentrate on discussing our filters and the filters that are comparable to them (i.e., the black circles). Filter 1 is a MOSFET–C filter without charge pump (discussed in [25]), it cannot reach high SNRs. It would, however, be suitable for building pulse equalisers similar to Filter 2. The comparison is not really fair, because Filter 2 is a 7th-order filter where all cascading problems have been solved. We believe, however, that if an experienced analogue-IC designer used MOSFET–C biquads to build a pulse equaliser, its performance would be similar to the performance of Filter 2.

This raises the question of biquad cascading. Filter 3 is the filter presented in this paper. It has a comparatively low dynamic range since every biquad has a low-frequency gain of two and only the last biquad can use its full signal swing. If a gain of one was used, the maximum current through the filter would increase by a factor of about 5, the power consumption would decrease slightly, and Filter 3 would end up somewhere between Filter 6 and Filter 9. However, using unity gain would also increase the variance of the q_p of the biquadratic stages. As a compromise, giving the highest-q_p biquad a gain of two and the remaining biquads a gain of one results in sufficiently stable poles, gives 9.5 dB more SNR as discussed above, and places Filter 3 into the group of Filters 4, 5, and 6. Compared to these three filters, the main disadvantage of our filter is that it requires a charge pump, the main advantage is its size: it only uses $0.04\,\mathrm{mm}^2$ per pole (including the charge pump), whereas, e.g., Filter 4 (an LC ladder simulation) uses $0.25\,\mathrm{mm}^2$ per pole, and Filter 7 uses $0.12\,\mathrm{mm}^2$ per pole, but uses less power per pole and frequency.

Filter 9 [28] is the best single biquad we could build with the MOSFET–C SAB technique. It has a pole Q of three. With its high SNR, its low power per pole and frequency, its tuning range of 26–36 MHz, and its chip area use of only $0.055\,\mathrm{mm}^2$ per pole (including the charge pump), it is among the best available continuous-time biquadratic filter sections, at least according to the figure of merit discussed here. It is, however, an open (and complex) question how such biquads can be cascaded in an optimum way, and which performance can be achieved with a higher-order filter. The educated guess made in the previous paragraph lets one expect that it is possible to build a 7th-order Bessel filter with a power per pole and frequency of 400 pJ, an SNR at 1% THD of 60 . . . 65 dB, and a chip area of $0.04\,\mathrm{mm}^2$ per pole.

It is open to debate whether a comparison by a simple figure of merit is meaningful at all. If it is, we have shown that our filters can achieve a performance similar to the performance of typical Gm–C video-frequency filters while using far less chip area. We think, however, that a figure like Fig. 13 should mainly be used as a map showing with which other filters one should compare ones own filters in more detail. Much more important than a comparison with other filters is a discussion of trade-offs.

10. Discussion

Several trade-offs that are important during the design of our filters were already discussed above. This final section covers a few important trade-offs from a wider perspective; the aim is to give the reader an impression of what can be done with MOSFET–C SABs.

Amplifier input resistance and output capacitance. As described above, the maximum achievable pole frequency of an SAB is determined by the required stopband attenuation and by the input resistance and output capacitance of the amplifier. Since the output capacitance cannot be decreased much without reducing the voltage swing (and with it the signal swing), the only viable alternative is to reduce the input resistance. One way to do this is to simply increase the supply current of the current amplifier input stage. However, since this current is mirrored to all other stages, this makes it necessary to enlarge the current mirrors, which again increases the output capacitance and limits the use of this method.

Another idea is to reduce the input resistance by using local feedback with a very high unity-gain bandwidth. Then the local feedback amplifier would consume the major part of the total power, which is possibly the only way to considerably increase the maximum possible pole frequency by trading power off for speed.

Signal swing, THD, and SNR. It was explained above how the signal swing in charge-pumped MOSFET–C SABs should be set in order to maximise the SNR of the filter at a certain level of THD. However, there is little correlation between the level of THD and the SNR in a certain filter. Because the THD is mostly caused by clipping, it increases quickly when a certain input current is exceeded. Therefore the SNR for -40 dB THD normally is only 2 dB larger than for -60 dB THD.

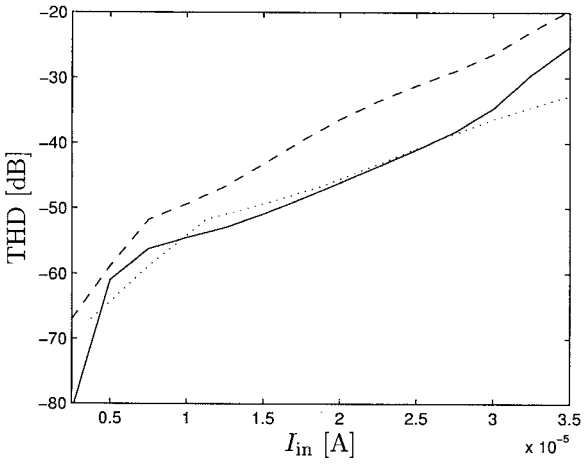

Fig. 14. Simulated THD of a conventional MOSFET–C SAB (solid), of a MOSFET-only SAB (dashed), and the latter curve with I_{in} scaled by 1.5 (dotted).

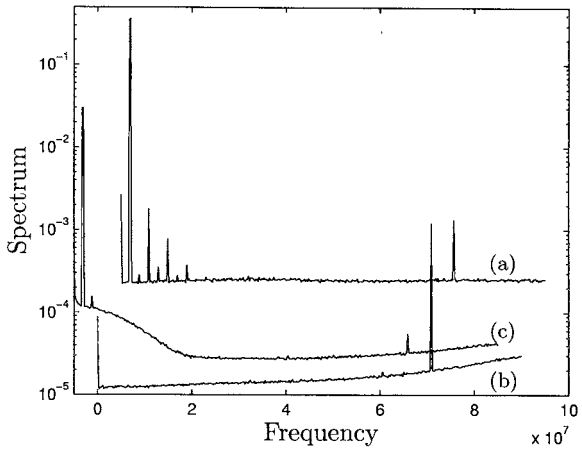

Fig. 15. Clock feed-through at the output pins of a biquad with incorrect signal line layout (a), the same with the biquad switched off such that only the passive clock feed-through caused by improper layout is visible (b), and in the 7th-order filter described in this paper (c). The signal frequency in (a) and (c) is the same (2 MHz), but not the magnitude. Note that the three spectra are shifted by 5 MHz in frequency such that all peaks remain visible. The different noise floors come from different input attenuator settings of the spectrum analyser.

Double-poly or single-poly process. One may also use gate-oxide capacitors instead of poly-poly capacitors to implement the signal capacitances. The resulting filter is then compatible with standard digital CMOS processes. Fig. 14 shows the THD simulated at $f_p/2$ of a conventional MOSFET–C SAB (solid), of a MOSFET-only SAB (dashed) where the poly-poly capacitors have simply been replaced by gate-oxide capacitors, and the latter curve with I_{in} scaled by 1.5 (dotted).

This shows that the replacement reduces the maximum possible input current to 65%, which amounts to a loss of only 4 dB of SNR. Note that this number depends both on the operating point voltage and on the process parameters. The capacitor block becomes only slightly larger (by about 25%). This is also not so in general, in fact, the gate-oxide capacitance can be *much larger* than the poly-poly capacitance in modern deep sub-micron processes, such that the capacitor block actually becomes smaller when gate capacitors instead of poly-poly capacitors are used. Nevertheless, it can reasonably be expected that MOSFET-only SABs generally will not have a much lower SNR than MOSFET–C SABs.

Charge pump or no charge pump? The advantages of driving the MOSFET resistor gates with a charge pump are so great that it should be done if possible. Also, the clock feed-through to the output of our filters is small enough for most applications, if the layout is made correctly. In one of the test biquads, we made the mistake of laying the filter output lines over the guard ring of the charge pump and then in parallel to the supply line of the charge pump. The result was unacceptably large clock feedthrough, as demonstrated by the measurements presented in Fig. 15.

There are two things that could prevent the use of a charge pump.

First, although our filters reject the substrate noise generated by the charge pump quite well, it must be made sure that the same is also true for all other signal-processing circuits on the chip. This may be a problem on purely analogue ICs, but is not really an issue on mixed-signal ICs, because there the substrate noise of the digital part dominates anyway.

Second, the charge pump described above is constructed such that although its output voltage can reach 5 V, no terminal voltage difference on any elements will exceed 3.3 V. Theoretically, no break-down will occur even if the process used does not support 5 V (the process we used is actually a 5-V process). The same is true for the MOSFET–C SABs. However, over-peaking during the transients (start-up) might change this, and it must be made sure by careful simulations that the charge pump is compatible with the process at hand.

Acknowledgment

The major part of this work has been completed as a doctoral thesis at the Signal an Information Processing Laboratory (ISI) of the Swiss Federal Institute of Technology, Zürich (ETH Zürich).

I thank the former head of the ISI, Prof. em. George S. Moschytz, for the great support he gave me there, Felix Frey for helping me with the development of the measurement electronics, Drahoslav Lím and Markus Helfenstein for the continuing support with all aspects of analog IC design, and the ETH Zürich for funding this research project.

Notes

1. The gain α_V has to be low, as will become apparent later in the paper, and therefore its variance also depends on matching. In a voltage amplifier, it often depends on the ratio of two resistors, and in a current amplifier, it depends on the matching of current mirror transistors.
2. The word *considerably* is not quantified here because building a 3.3–V MOSFET–C filter in the process we used was barely possible at all without a charge pump. C.f. the discussion in Section 9.

References

1. Brown, J. E. C., Hurst, P. J., Rothenberg, B. C. and Lewis, S. H., "A CMOS adaptive continuous-time forward equalizer, LPF, and RAM-DFE for magnetic recording." *IEEE Journal of Solid-State Circuits* 34(2), pp. 162–169, 1999.
2. Burger, T. and Huang, Q., "A 100 dB, 480 MHz OTA in 0.7 μm CMOS for sampled-data applications," in *Proceedings of the IEEE Custom Integrated Circuit Conference,* San Diego, California, 1996, pp. 101–104.
3. Celma, S., Sabadell, J., Aldea, C. and Martínez, P. A., "Video-frequency current-voltage mode integrator." *Electronics Letters* 35(10), pp. 773–775, 1999.
4. Chang, Z. Y., Haspeslagh, D. and Verfaillie, J., "A highly linear CMOS G_m-C bandpass filter with on-chip frequency tuning." *IEEE Journal of Solid-State Circuits* 32(3), pp. 388–397, 1997.
5. Dehaene, W., Steyaert, M. S. J. and Sansen, W., "A 50-MHz standard CMOS pulse equalizer for hard disk read channels." *IEEE Journal of Solid-State Circuits* 32(7), pp. 977–988, 1997.
6. Duisters, T. A. F. and Dijkmans, E. C., "A −90-dB THD rail-to-rail input opamp using a new local charge pump in CMOS." *IEEE Journal of Solid-State Circuits* 33(7), pp. 947–955, 1998.
7. Gopinathan, V., Tarsia, M. and Choi, D., "Design considerations and implementation of a programmable high-frequency continuous-time filter and variable-gain amplifier in submicrometer CMOS." *IEEE Journal of Solid-State Circuits* 34(12), pp. 1698–1707, 1999.
8. Gopinathan, V., Tsividis, Y., Tan, K.-S. and Hester, R., "A 5 V 7th-order elliptic analog filter for digital video applications," in *Proceedings of the IEEE International Solid-State Circuits Conference,* San Francisco, 1990, pp. 208–209.
9. Helfenstein, M., "Analysis and design of switched-current networks." Ph.D. Thesis, Swiss Federal Institute of Technology (ETH), Zürich, Switzerland. (Diss. ETH No. 12257), 1997.
10. Huang, Q., "A MOSFET-only continuous-time bandpass filter." *IEEE Journal of Solid-State Circuits* 32(2), pp. 147–158, 1997.
11. Hung, C.-C., Halonen, K. A. I., Ismail, M., Porra, V. and Hyogo, A., "A low-voltage, low-power CMOS fifth-order elliptic GM-C filter for baseband mobile, wireless communication." *IEEE Transactions on Circuits and Systems for Video Technology* 7(4), pp. 584–593, 1997.
12. Kim, C.-S., Cho, G.-O., Kim, Y.-H. and Song, B.-S., "A CMOS 4 × speed DVD read channel IC." *IEEE Journal of Solid-State Circuits* 33(8), pp. 1168–1178, 1998.
13. Kosunen, M., Koli, K. and Halonen, K., "A 50 MHz 5th order elliptic LP-filter using current mode gm-C topology," in *Proceedings of the IEEE International Symposium on Circuits and Systems*. Monterey, California, 1, pp. 512–515, 1998.
14. Lidgey, J., Toumazou, C., Payne, A., Wadsworth, D. C., Pookaiyaudom, S. and Bruun, E., "Tutorial 10: Current-mode analog signal processing; Part 2: Current Conveyors," in *Proceedings of the IEEE International Symposium on Circuits and Systems,* London, 1994, pp. 569–641.
15. Mehr, I. and Welland, D. R., "A CMOS continuous-time G_m–C filter for PRML read channel applications at 150 Mb/s and beyond." *IEEE Journal of Solid-State Circuits* 32(4), pp. 499–513, 1997.
16. Moschytz, G. S., "Single-amplifier active filters: A review." *Scientia Electrica* 26(1), pp. 1–46, 1980.
17. Moschytz, G. S. and Carlosena, A., "A classification of current-mode single-amplifier biquads based on a voltage-to-current transformation." *IEEE Transactions on Circuits and Systems–II: Analog and Digital Signal Processing* 41(2), pp. 151–156, 1994.
18. Nagari, A. and Nicollini, G., "A 3 V 10 MHz pseudo-differential SC bandpass filter using gain enhancement replica amplifier." *IEEE Journal of Solid-State Circuits* 33(4), pp. 626–630, 1998.
19. Pavan, S., Tsividis, Y. P. and Nagaraj, K., "Widely programmable high-frequency continuous-time filters in digital CMOS technology." *IEEE Journal of Solid-State Circuits* 35(4), pp. 503–511, 2000.
20. Punzenberger, M. and Enz, C. C., "A 1.2-V low-power BiCMOS class AB log-domain filter." *IEEE Journal of Solid-State Circuits* 32(12), pp. 1968–1978, 1997.
21. Punzenberger, M. and Enz, C. C., "A compact low-power BiCMOS log-domain filter." *IEEE Journal of Solid-State Circuits* 33(7), pp. 1123–1129, 1998.
22. Rao, N., Balan, V. and Contreras, R., "A 3-V, 10–100-MHz continuous-time seventh-order 0.05° equiripple linear phase filter." *IEEE Journal of Solid-State Circuits* 34(11), pp. 1676–1682, 1999.
23. Rezzi, F., Bietti, I., Cazzaniga, M. and Castello, R., "A 70-mW seventh-order filter with 7–50 MHz cutoff frequency and programmable boost and group delay equalization." *IEEE Journal of Solid-State Circuits* 32(12), pp. 1987–1999, 1997.

24. Schmid, H., "Approximating the universal active element." *IEEE Transactions on Circuits and Systems–II: Analog and Digital Signal Processing* 47(11), pp. 1160–1169, 2000a.
25. Schmid, H., "Single-amplifier biquadratic MOSFET–C filters for video frequencies." Ph.D. Thesis, Swiss Federal Institute of Technology, Zürich. (ETH Thesis No 13878; copies can be obtained from h.p.schmid@ieee.org), 2000b.
26. Schmid, H. and Moschytz, G. S., "Fundamental frequency limitations in current-mode Sallen-key filters," in *Proceedings of the IEEE International Symposium on Circuits and Systems* Monterey, California, vol. 1, 1998, pp. 57–60.
27. Schmid, H. and Moschytz, G. S., "Minimum-sensitivity single-amplifier biquadratic filters," in *Proceedings of the European Conference on Circuit Theory and Design*, Stresa, Italy, vol. 2, 1999a, pp. 1027–1030.
28. Schmid, H. and Moschytz, G. S., "A tunable, video-frequency, low-power, single-amplifier biquadratic filter in CMOS," in *Proceedings of the IEEE International Symposium on Circuits and Systems* Orlando, Florida, vol. 2, 1999b, pp. 128–131.
29. Schmid, H. and Moschytz, G. S., "Active-MOSFET-C single-amplifier biquadratic filters for video frequencies." *IEE Proceedings—Circuits, Devices and Systems* 147(1), pp. 35–41, 2000a.
30. Schmid, H. and Moschytz, G. S., "A charge-pump-controlled MOSFET–C single-amplifier biquad," in *Proceedings of the IEEE International Symposium on Circuits and Systems,* Geneva, Switzerland, vol. 2, 2000b, pp. 677–680.
31. Serdijn, W. A., Broest, M., Mulder, J., van der Woerd, A. C. and van Roermund, A. H. M., "A low-voltage ultra-low-power translinear integrator for audio filter applications." *IEEE Journal of Solid-State Circuits* 32(4), pp. 577–581, 1997.
32. Tawfik, M. S. and Senn, P., "A 3.6-MHz cutoff frequency CMOS elliptic low-pass switched-capacitor ladder filter for video communication." *IEEE Journal of Solid-State Circuits* 22(3), pp. 378–384, 1987.
33. Wu, C.-Y. and Hsu, H.-S., "The design of CMOS continuous-time VHF current and voltage-mode lowpass filters with Q-enhancement circuits." *IEEE Journal of Solid-State Circuits* 31(5), pp. 614–624, 1996.
34. Yoo, C., Lee, S.-W. and Kim, W., "A ±1.5-V, 4-MHz CMOS continuous-time filter with a single-integrator based tuning." *IEEE Journal of Solid-State Circuits* 33(1), pp. 18–27, 1998.
35. Zele, R. H. and Allstot, D. J., "Low-power CMOS continuous-time filters." *IEEE Journal of Solid-State Circuits* 31(2), pp. 157–168, 1996.

Hanspeter Schmid received the diploma in electrical engineering in 1994, the post-graduate degree in information technologies in 1999, and the degree Doctor of Technical Sciences in 2000 from the Swiss Federal Institute of Technology (ETH Zürich), Switzerland. He is now an analog IC designer with Bernafon AG, Switzerland, where he develops analog electronics for hearing aids. He is a lecturer at ETH Zürich (Analog Signal Processing and Filtering), a member (currently the Secretary) of the Analog Signal Processing Committee of the IEEE Circuits and Systems Society, and he serves several journals and conferences as a reviewer.