INTERNALLY VARYING ANALOG CIRCUITS MINIMIZE POWER DISSIPATION

Significant Power Dissipation Savings Can Be Obtained by Allowing the Internal Circuit Attributes to Vary Dynamically As Needed for the Task At Hand

Y. Tsividis, N. Krishnapura, Y. Palaskas, and L. Toth

Low energy drain is key to the success of portable, battery-operated equipment. If the energy drain per task can be slashed down, battery life can proportionately be increased. Most conventional analog circuits are designed worst-case, for adequate performance while handling the most demanding tasks. This requires large power dissipation; when these circuits are called to perform less demanding tasks, they continue to dissipate the same large power unnecessarily. In other words, in such cases we are *stuck* with a fixed circuit.

The solution is to allow the *internal* attributes of the circuit to vary, depending on the task at hand. In this article we review several recently proposed techniques that make this possible. We call these circuits *internally varying*. (We cannot call these circuits "time varying," as in some systems context this term implies time-varying input-output behavior, which is not necessarily our goal.) The variables that control the internal variations can be independent variables controlled by the larger system, in which case we have an *adjustable* or *programmable* internally varying circuit; or, they can be derived from the signals being handled by the circuit, in which case we have an *adaptive* internally varying circuit.

As discussed below, internally varying circuits have in general a time-dependent power dissipation $p(t)$; the total energy drain W between times t_1 and t_2, is given by:

$$W = \int_{t_1}^{t_2} p(t)dt. \tag{1}$$

8755-3996/03/$17.00 ©2003 IEEE

By varying the circuit's internal structure and making it optimum for the task at hand, p is minimized for each task, and over time W can be made much smaller than what it would have been for conventional circuits. The battery will last much longer.

Varying a circuit structure can cause undesired disturbances at the output. In certain cases, such disturbances do not occur to a significant extent or are not an issue. Reducing the power dissipation in such cases is relatively easy; for example, it has been done for a very long time in power amplifiers, which sometimes are called to transmit at less than maximum power. Such cases will not be considered in this article. Rather, we will consider general and demanding situations where the internal attributes of a *dynamical* analog circuit are varied, while the circuit is processing a signal. This can cause strong transients at the output and can interfere with proper operation. A significant part of the techniques we discuss deal with how to *prevent or eliminate such transients*. Although much of our discussion is of a general nature that can be applied to many analog circuits, we will use as a specific example the particularly tough case of filters.

We first give some definitions and review the reasons for large power dissipation in conventional signal processing circuits. We then discuss the dynamic variation of several internal aspects of such circuits, including gain distribution, impedance levels, bias levels, and structure.

Signal-to-Noise Ratio and Usable Dynamic Range: Two Distinct Quantities

We consider for simplicity a system with unity voltage gain. For a given in-band signal with RMS value V_s and noise with RMS value V_n, the signal-to-noise ratio (SNR) is defined in terms of the squares of these quantities, and thus, in systems parlance, is a "power ratio":

$$\mathrm{SNR} = \frac{V_s^2}{V_n^2}. \tag{2}$$

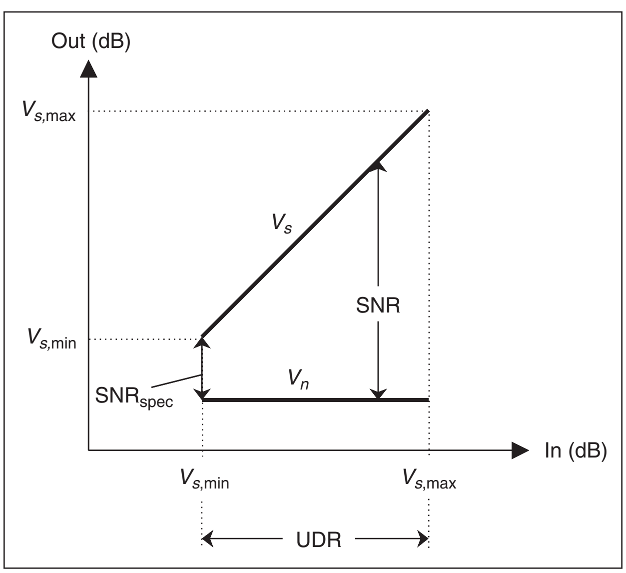

1. Output RMS signal and noise versus input RMS signal for a unity-gain signal processor.

A simple specification often encountered is that this quantity be no less than a certain value, $\mathrm{SNR}_{\mathrm{spec}}$. The range of signal values over which this happens will be defined as the *usable dynamic range* (UDR):

$$\mathrm{UDR} = \frac{V_{s,\max}^2}{V_{s,\min}^2}. \tag{3}$$

We assume that $V_{s,\max}$ is made as large as possible; the value of this quantity is such that an appropriate distortion measure (e.g., harmonic or intermodulation distortion) does not exceed a maximum allowed value. The value of $V_{s,\max}$, together with the desired UDR, sets $V_{s,\min}$ from (3). We assume that V_n is kept below this level so that SNR from (2) is just equal to $\mathrm{SNR}_{\mathrm{spec}}$ at that point:

$$V_n^2 = \frac{V_{s,\min}^2}{\mathrm{SNR}_{\mathrm{spec}}} = \frac{V_{s,\max}^2}{\mathrm{SNR}_{\mathrm{spec}} \times \mathrm{UDR}}. \tag{4}$$

Since for conventional circuits the noise floor is constant, this guarantees that $\mathrm{SNR} > \mathrm{SNR}_{\mathrm{spec}}$ at larger signal levels, as shown in Figure 1. In this figure we show $\mathrm{SNR}_{\mathrm{dB}} = 10\log(\mathrm{SNR})$ and $\mathrm{UDR}_{\mathrm{dB}} = 10\log(\mathrm{UDR})$, and their relation to $V_{s,\mathrm{dB}}$ and $V_{n,\mathrm{dB}}$, the latter two obtained as 20 times the log of their ratio to an arbitrary reference. The subscript "dB" in each case has been dropped in order not to clutter the figure. Ratios such as in (2) and (3) correspond to linear distances in the dB plots in Figure 1.

Notice that the maximum SNR is in general not equal to the usable dynamic range; it becomes equal to it only if $\mathrm{SNR}_{\mathrm{spec}}$ is 0 dB, which is a very unlikely specification in practice. From the plot, at $V_s = V_{s,\max}$ we obtain $\mathrm{SNR}_{\max} = \mathrm{SNR}_{\mathrm{spec}} + \mathrm{UDR}$. Thus, at large signal levels, UDR represents a "waste" above the required $\mathrm{SNR}_{\mathrm{spec}}$, which leads to major problems in the design of conventional circuits, as is now discussed.

Power Dissipation and Chip Area Waste in Fixed Circuits

The tradeoffs involved among SNR, chip area, and power dissipation in conventional continuous-time active circuits have been studied in the literature; see, for example, [1]-[5]. In their simplest form [2], these tradeoffs can be illustrated using a first-order RC low-pass circuit with bandwidth $1/(2\pi RC)$, in which the resistor noise power spectral density, equal to $4kTR$, is processed by the equivalent noise bandwidth of the circuit, $1/4RC$; thus, a mean-square noise of $4kTR \times [1/(4RC)] = kT/C$ appears at the output [6]. Therefore, low noise requires large capacitance (large chip area) and, for a given bandwidth, low resistance; the resulting large admittance levels require large currents in order to be driven at large signal swings, which results in large power dissipation. Finally, since these admittance levels are also proportional to band-edge frequency, so is power dissipation.

Similar results can be derived for larger circuits. For illustration purposes, we use here the case of a second-order bandpass filter with a high quality factor Q (ratio of center frequency to bandwidth), implemented with a resonator fed by an input device; two implementations of this circuit are shown in Figure 2 [7]-[10]. The voltage gain is 1 at the center frequency. We will concentrate on the circuit on the left, which is implemented with transconductors and capacitors. The center frequency of this circuit is $G_m / (2\pi C)$. The tradeoffs between noise, capacitance, and power dissipation for this circuit have been studied quite early by Bloom and Voorman [1]. A detailed calculation of noise shows that kT / C again appears as above, along with some new factors [1]:

$$V_n^2 \approx 2Q\gamma \frac{kT}{C}. \tag{5}$$

Here the factor of 2 accounts for the contributions of the two large transconductors (which can be shown to be dominant), and Q appears due to the fact that there is a gain of Q between the equivalent input noise of these transconductors and the output. The quantity γ is the "excess noise factor" of the transconductors (typically 2 to 3). The value of C must be chosen to give a sufficiently low noise from (5), to satisfy the specs illustrated in Figure 1. Solving (5) for C, and using (4) in the result, we obtain:

$$C \approx \frac{2Q\gamma kT}{V_{s,\max}^2} \times \mathrm{SNR}_{\mathrm{spec}} \times \mathrm{UDR}. \tag{6}$$

Thus, large UDR and small signal swings (necessary to achieve low distortion) require large capacitance. This is the reason that, in many integrated filters, especially ones with high-Q poles, the chip area is dominated by the capacitors. In fact, *for each 3 dB improvement in UDR* (a factor of 2 increase in this power ratio), *the capacitance area must be doubled*.

Figure 3(a) shows a simplified schematic of a transconductor that will be used as an example in this article. Such a transconductor would be appropriate for use in the balanced version of the circuit in Figure 2(a). Assuming that the transistor transconductance is much larger than the resistor conductance, the transconductance of the circuit is approximately equal to $1/R$. [The resistor can be replaced by a MOS transistor in the triode region, which behaves linearly as long as the input signal is balanced [11]. The gate voltage of this device can be used for limited transconductance tuning (e.g., to counteract tolerances and temperature variations).] The input-output characteristic is as shown in Figure 3(b). The value $V_{s,\max}$ used above is marked in the figure, and it has to be at a point well below clipping if low distortion is desired. This means that the maximum output current will be well below $I_{\mathrm{BIAS}}/2$. We should note here that, for a given transconductor nonlinearity, distortion in filters depends not only on signal level but also on the quality factor Q, generally increasing with the latter [12]. Thus, $V_{s,\max}$ needs to be kept especially low in high-Q filters.

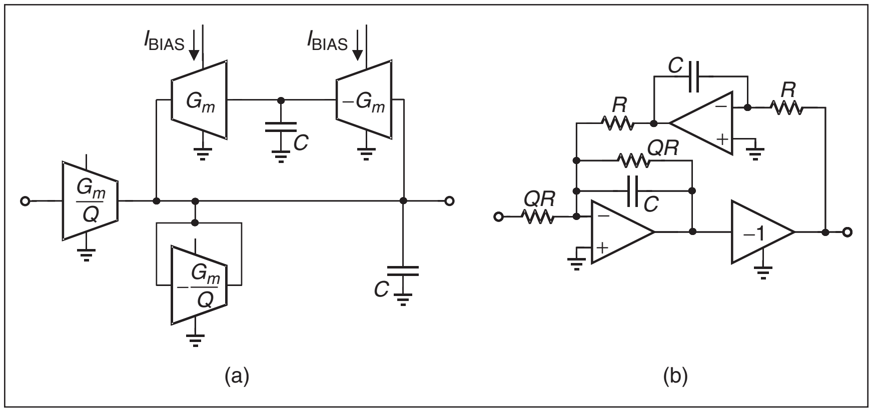

2. Unity-gain second-order bandpass filters: (a) Transconductor-C (gyrator-C); (b) active RC (Tow-Thomas).

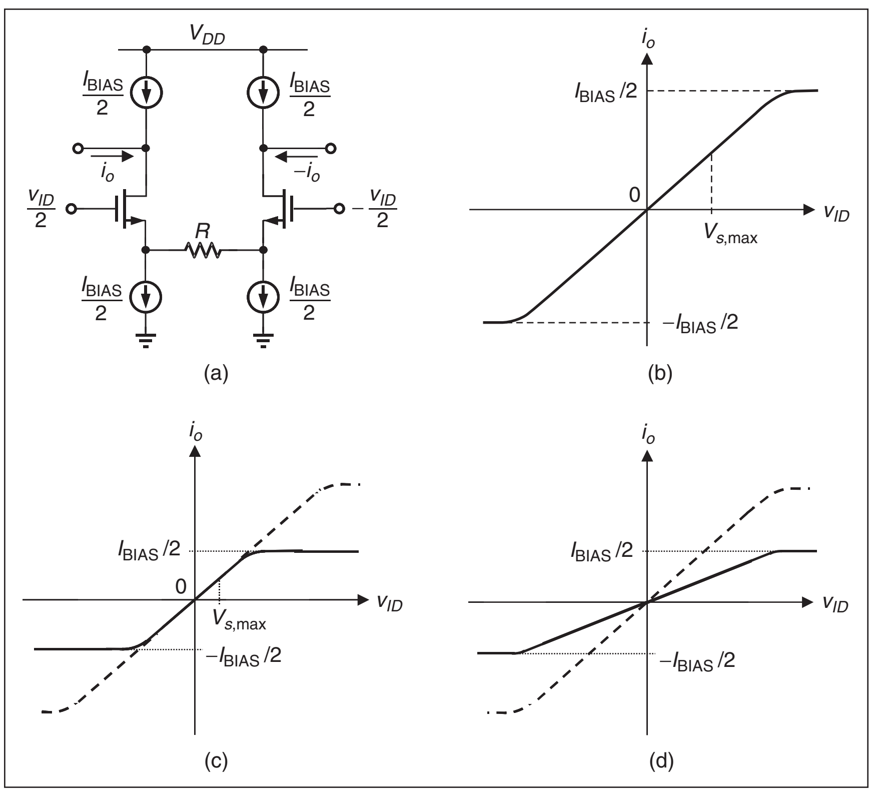

3. Balanced input-output tranconductor; (a) simplified circuit; (b) output current versus input differential voltage (the slope is equal to the transconductance, G_m); (c) effect of varying the bias current while keeping G_m constant; (d) effect of varying G_m by scaling all elements in (a).

For future reference, we show in Figure 3 what happens if the bias current is changed [Figure 3(c)] and if the transconductance is changed (e.g., by changing the number of several transconductor units placed in parallel) [Figure 3(d)].

The transconductor bias currents in Figure 2(a), assuming class A operation and low distortion, must be larger than the maximum current these elements are expected to drive (which, for the two larger transconductors, is approximately the capacitor admittance at the center frequency times the maximum peak voltage) by a factor that we will denote by b (e.g., 3 to 6):

$$I_{\text{BIAS}} = b \times 2\pi f C \times \sqrt{2} V_{s,\max}. \tag{7}$$

The total bias current of the filter in Figure 2(a) is approximately twice this amount, since the two large transconductors draw most of the current. This, multiplied by V_{DD}, gives an expression for the power dissipation P. If, in this relation, C is replaced by (6), we obtain:

$$P \approx akTQf \times \frac{V_{DD}}{V_{s,\max}} \times \text{SNR}_{\text{spec}} \times \text{UDR} \tag{8}$$

where the factor a is equal to $8\sqrt{2}\pi b\gamma$ and in practice can have a value of several hundred. We thus see that P, in addition to being proportional to the center frequency, is proportional to the UDR; *improving the required UDR by 3 dB requires doubling the power dissipation!* The factor $V_{DD} / V_{s,\max}$ results from the above derivation and reminds us that, if we must keep the signals well below the supply voltage for low distortion, we will pay a price in power. We remind the reader that distortion effects are more severe in high-Q filters, and thus the allowable $V_{s,\max}$ will be smaller for higher Q. For this reason (6)-(8) cannot be used to compare two filters with different Q factors, unless the value of $V_{s,\max}$ has been determined for each of them.

Similar results (with $\gamma = 1$) are found for the circuit of Figure 2(b), assuming that resistor noise dominates op-amp noise, which is often the case.

For high-order circuits, results similar to the above are found [3], [4]. Here the problems are compounded by the fact that each circuit block contributes noise, and this noise is especially large for filters containing high-Q poles, such as required, for example, for very narrow passband and/or very sharp transition characteristics. One can use the same formulas as above, with Q corresponding to the highest-Q pole pair, and appropriately changing the value of the constant a in (8), depending on filter order, architecture, and node voltage scaling strategy [13], [14]. Again, *the total capacitance and the power dissipation are proportional to the required UDR*, but the constant of proportionality is larger. In the rest of this article we will continue to use the above relations for illustrative purposes.

The above simple estimates suffice to demonstrate the practical limitations of conventional on-chip active filters. For example, consider a second-order section with Q of 20 and a center frequency of 10 MHz, SNR_{spec} = 40 dB, and operating with V_{DD} = 3 V and $V_{s,\max}$ = 0.5 V. Assume $\gamma = 3$ and $a = 200$. We will contrast two cases, one with low UDR [Figure 4(b)] and one with high UDR [Figure 4(a)]. If UDR = 20 dB, the required C from (6) is 2 pF, and the resulting P from (8) is 1 mW. If, instead, the spec calls for UDR = 60 dB, the capacitance becomes 20,000 pF and the power 10 W, either of which would rule out an integrated implementation.

It can thus be appreciated that, in integrated filter design, we can run out of pF and of mW pretty quickly. The reason can be traced to the fixed nature of conventional circuits. In the rest of this article, we review several techniques through which the internal attributes of the circuits can be varied, thus making possible drastic savings in chip area and/or power dissipation.

Technique 1: Companding (Dynamic Gain Scaling)

Assume that a filter must satisfy specs as in Figure 4(a), requiring a very wide UDR, denoted by UDR_1 in that figure. We can process this signal using a filter with a much narrower range UDR_2 as indicated in Figure 4(b), if we can first compress the signal range UDR_1 into UDR_2, process the signal with a filter designed to handle only this narrow range, and expand the range at the output back to UDR_1 (the extent to which this is possible may be limited if large "blockers" are present—see below). This process of compressing and expanding is called "companding." It has been used in telecommunications and audio recording for a very long time [15], [16], but its use in signal processing [17] is more difficult, as will be explained shortly.

Companding can be done as shown in Figure 5. Assume for simplicity that the filter has unit gain, that the signal is in the passband, and that the noise of the input and output gain elements is negligible in comparison to the filter noise. The input

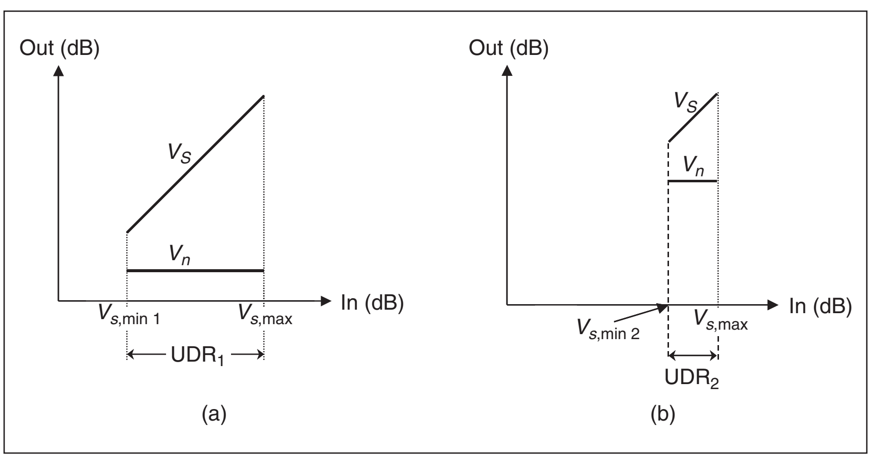

4. Output RMS signal and noise versus input RMS signal for signal processors with (a) wide and (b) narrow usable dynamic range.

and output gain blocks are controlled together, and their gains are inverses of one another (see below).

> Varying a circuit structure can cause undesired disturbances at the output. These techniques deal with how to prevent or eliminate such transients.

We consider two representative cases in Figure 6. In (a), the input signal is significantly smaller than what the filter can handle; g is thus made larger than unity (5 in this example), to increase the signal so that it remains well above the filter's noise floor. The output gain element attenuates the signal back to its original level and at the same time attenuates the noise by the same amount. In (b), the input is assumed to be as large as the filter can handle, so g is made equal to 1. The signal is sufficiently above noise, so that the output SNR is satisfactory. The output gain is also equal to 1 in this case, and both the signal and the noise pass to the output of the system essentially unchanged. In some cases, even larger signals can be handled at the system input, in which case g will have to be made smaller than 1. In all cases the same, narrow-UDR filter is used, since the signal presented to the filter is always kept large and therefore it does not require a low noise floor. Thus, the filter can be designed economically with small capacitors and low power dissipation. As expected from the above calculations, power savings by a factor of UDR_1 / UDR_2 can be achieved in the filter; for example, if the dynamic range is narrowed by 40 dB, the required filter dissipation can in principle be decreased 10,000 times! This savings is, of course, not attainable in practice for the overall system because the supporting circuitry will also require some power; the savings can be drastic nevertheless.

It can be seen from Figure 6 that, in all cases, the output SNR is approximately the same. This tendency toward flatness of the SNR characterizes all companding systems (see, for example, the measurements in [18]).

The key to the above technique is keeping the signal amplitude at the input to the filter constant at a near-optimum value (large enough to be well above the filter noise, but not so large as to excite the filter nonlinearities to a significant extent). Thus, g should in principle be inversely proportional to the amplitude of the signal at the input of the system. If the system input amplitude is known, an appropriate control can be applied to the g and $1/g$ elements to achieve this; if it is not known, an envelope detector can be used, connected to the input or output of the first gain element [17] (Figure 5), using well-known feedforward or feedback automatic gain control (AGC) techniques [19]. (In the special case where the filter is the IF or baseband filter in a receiver, the output gain block may or may not be necessary; the function of this block should be considered in the context of the AGC strategy of the receiver.) Depending on the signal characteristics, measures other than the envelope (e.g., the rms value) may be used, but using the envelope is safer, especially if the signal's crest factor is large and not precisely known. In general, the required circuits have release time constants corresponding to many signal cycles; this type of companding is referred to as "syllabic" companding, since in speech processing such time constants would be comparable to the duration of word syllables. Issues with the design and use of envelope detectors are similar to those encountered in AGC design [19].

The above technique can be applied also when the system's input and output variables are of a different type than those at the input and output of the filter. As an example, Figure 7 shows a system called to handle current signals [20]. These are converted to voltage by the input transconductor connected as a resistor ($v_1 = i_{IN} / G_{mc}$), are processed by the voltage-in, voltage-out filter in the center, and are converted back to current by the output transconductor. This system corresponds to the one in Figure 5, with $g = 1/G_{mc}$. The transconductance G_{mc} is made variable and proportional to the system's input current envelope by using an envelope detector as above. The input resistor-connected transconductor could be replaced by a variable resistor, but using an input transconductor ensures good matching with the characteristics of the output transconductor.

Both transconductors are required to handle a large, fixed signal voltage. G_{mc} can be varied over wide ranges if transconductor cells are switched in parallel as necessary. The combined bias current will vary along with the combined transconductance, as illus-

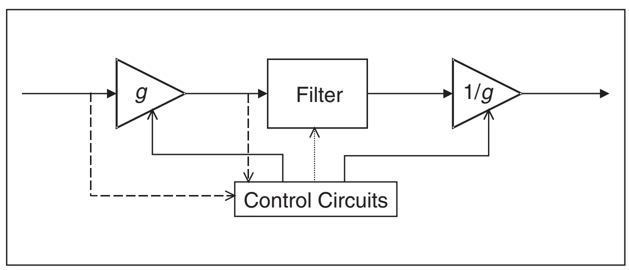

5. Companding signal processor.

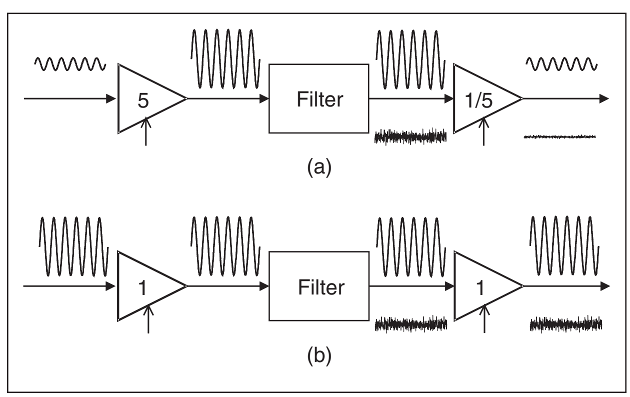

6. The main path of the companding signal processor of Figure 5 with (a) small input signal amplitude and (b) input signal amplitude increased by a factor of 5.

trated in Figure 3(d). Thus, if the signal is small, the total G_{mc} and I_{BIAS} will be small. Large power dissipation will be needed only when the signal is large. Figure 7(b) shows this behavior. Depending on the statistics of the input signal, drastic savings can be possible in the overall energy drain [20].

Leaving Headroom for Blockers

In certain applications (notably wireless), out-of-band signals can be much larger than in-band ones. This situation is depicted in Figure 8(a), where a large low-frequency blocker is superimposed on a small desired signal. Let the peak values of these signals be $V_{peak,blocker}$ and $V_{peak,signal}$, respectively. If the blocker were not present [Figure 6(a)], the system would have given g a value inversely proportional to $V_{peak,signal}$; however,

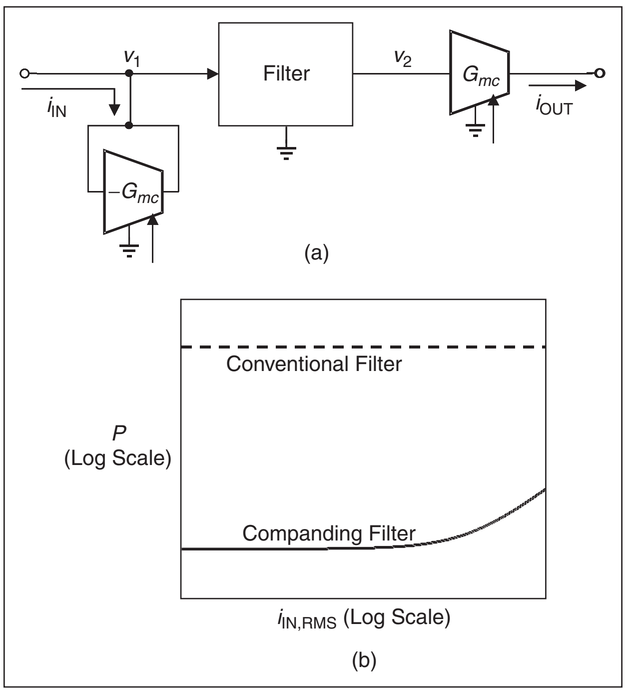

7. (a) The main path of the companding signal processor of Figure 5 in the case of current input and current output. (b) Resulting power dissipation versus input RMS amplitude.

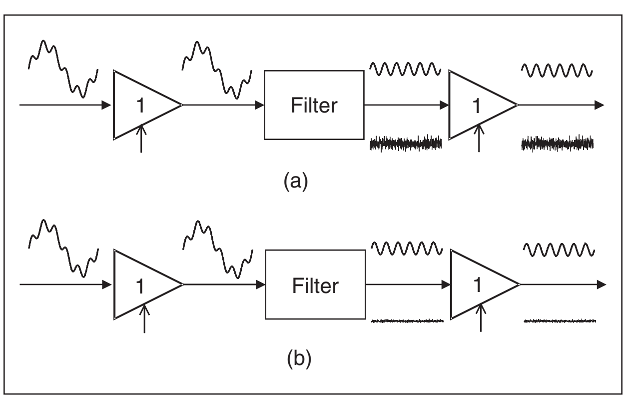

8. (a) The main path of the companding signal processor of Figure 5 with a desired input of the same magnitude as in Figure 6(a), superimposed on a much larger out-of-band "blocker"; output signal-to-noise ratio is deteriorated. (b) The system in (a) redesigned to provide adequate output signal-to-noise ratio in the presence of the blocker.

in the presence of the blocker [Figure 8(a)], g will instead be inversely proportional to $V_{peak,signal} + V_{peak,blocker}$, so it will assume a value which will be smaller by the factor $(V_{peak,signal} + V_{peak,blocker}) / V_{peak,signal}$. Although the blocker is removed at the output of the filter, the signal power at this point has been weakened by the square of the above factor, compared to the case in Figure 6(a), and the SNR has been thus weakened by the same factor, as seen in Figure 8(a). To correct this problem, the design of the system has to be revised so that the noise floor is made smaller by the same factor, too, as shown in Figure 8(b); this means that the filter should be designed with an SNR of

$$SNR_{filter} = SNR_{spec} \times HR \tag{9}$$

where HR is the headroom for blockers, defined by:

$$HR = \left(\frac{V_{peak,signal} + V_{peak,blocker}}{V_{peak,signal}} \right)^2 . \tag{10}$$

Thus, if blockers with total peak value as large as that of the signal are expected, a headroom of 6 dB should be allowed; if blockers with total peak value 10 times that of the signal are expected, the headroom becomes 21 dB. This headroom comes at the expense of capacitance and power dissipation. Thus, all estimates as to how much savings companding affords those quantities must be lowered by the headroom. Appropriate optimization, in which the gain within the filter is made progressively higher towards the output (i.e., as the blocker is progressively rejected), can help [13], [14].

Avoiding Output Transients

The use of companding in filters encounters a difficulty not present in companding transmission systems. Filters are dynamical circuits; changes at the input appear as changes at the output with certain delay. Assume that, in Figure 5, the input signal strength suddenly changes, thus momentarily changing the signal at the input of the filter. The envelope detector senses the system input change and adjusts g to restore the filter input to its original level. The output element's gain, $1/g$, changes accordingly at the same time, which is *before* the change at the filter input has had a chance to propagate to the output. Thus, the output element's gain changes too soon, and an undesirable transient occurs at the system output [17]. In some applications this is not a problem, as these changes may occur before the time comes to process the useful signal (for example, during the preamble slot in certain wireless protocols). If, however, continuous service is desired, this transient must be eliminated or prevented in the first place. Several schemes that have been proposed to accomplish this are now described; some may prove to be more practicable than others.

State variable correction: The above problem can be avoided if the dynamics of the filter are properly taken into account. It has been pointed out [21] that if the input gain in Figure 6 is switched from a value g_{OLD} to a value g_{NEW}, no transient will appear at the output if all state variables are "updated" by the same factor involved in the change of g. Thus, the voltage on a capacitor must be changed at that instant, from a value v_{OLD} to a value $v_{NEW} = v_{OLD} \times (g_{NEW} / g_{OLD})$. This can be accomplished, for example, by charging the capacitors in a filter appropriately at the transition [21]. If, instead, g is not switched but rather is changed continuously, a likewise continuous modification of the state variables can be used; this problem has a general analytical solution [22], but the resulting circuits have proven not to be easy to design [23]. State variable correction is indicated in Figure 5 by the dotted arrow going from the control circuits to the filter.

Zero-crossing switching: Consider a state variable v in the filter. According to the scheme just described [21], this value needs to be updated to $v_{NEW} = v_{OLD} \times (g_{NEW} / g_{OLD})$. If, however, we wait to perform this change at a moment that v crosses the time axis, $v_{OLD} = 0$ and thus $v_{NEW} = 0$, too. Thus, no updating is needed [24]. The implementation of this idea relies on accurate detection of zero crossings, which is complicated in practice, especially for high-frequency signals.

Feed-forward cancellation: Another way to solve the output transient problem is to allow the transient to occur but add to the output a second transient that is precisely the negative of the first [24]. To produce this second transient, a duplicate filter is needed; this filter is fed by the same signal as the main filter, until the time for the gain switching comes; then, the input of the second filter is switched to 0. Because, at the time of switching, this second filter has all its state variables identical to those in the main filter, it can be shown that the transient that occurs at its output, if multiplied by an appropriate gain, is opposite from that in the main filter; adding the two filter outputs results in elimination of that transient, and the signal appears undisturbed by the gain switching.

Combining companding with log domain: Consider again the system of Figure 7(a), using transconductors at the input and output. The simplest transconductor is a single transistor. If the two transconductors are replaced by bipolar transistors, we obtain the system of Figure 9(a). The transconductors are now of course nonlinear, and the total voltage at the filter input is proportional to the logarithm of the instantaneous transistor current. It can be shown that if this voltage is processed by a "log domain" filter core, and the filter's output is processed by an exponential nonlinearity as in Figure 9(a), the input-output behavior is linear [25], [26] (this is a perfectly sound concept, which has been widely misunderstood). Now, to vary the gain of the input "transconductor" in Figure 9(a), one can vary its bias current in proportion to the signal envelope [27], since $g_m = qI / kT$ for a bipolar transistor. If a second filter is used as shown in Figure 9(b), with the same bias but the opposite input signal, the transients due to the bias change are the same at the output of both filters; thus, if the *difference* of the two outputs is taken, this transient cancels out, although the signal components add [28]. Key to this idea is the fact that, since the two systems are *linear* from input to output, both signal and bias are treated by them by one and the same transfer function. A chip using this approach [28] has been demonstrated to give a very large UDR, which approaches the fundamental limits of power dissipation of *passive* RC circuits with the same UDR! (Passive circuits do dissipate some power—from the input signal source.)

Divide-and-conquer: Rather than switching the gain elements in the system of Figure 5, one can use more than one path with different, but fixed, gain elements, as shown in the example of Figure 10(a) [29]. Here the gain elements are fixed. The three paths should be closely matched. Depending on the input signal, the path switched to the output is changed; large signals are processed by the top path, smaller ones by the middle path, and very small ones by the bottom path. Transients at the output are avoided, since the paths have the same outputs, as long as no large distortion occurs. Thus, for example, assume that as the signal is being processed by the middle path, its amplitude begins to decrease, which would result in a decreased *SNR* at the output; the decreased signal level is appropriate for the bottom path, which has a better output SNR thanks to the larger gain at its input (see Figure 6). Other than that, both paths produce the same output signal, so when the output is switched to the lower path, no transients occur. More sophisticated schemes can be used, where some filters are switched off part of the time to save power, while output transients can still be avoided by employing a special switching sequence [30].

The overall signal-to-noise-plus-distortion ratio (SNDR) of this system is as shown in Figure 10(b). As seen, the total UDR has been divided into three regions, and each filter is only required to handled a limited UDR. The sloping off of each curve

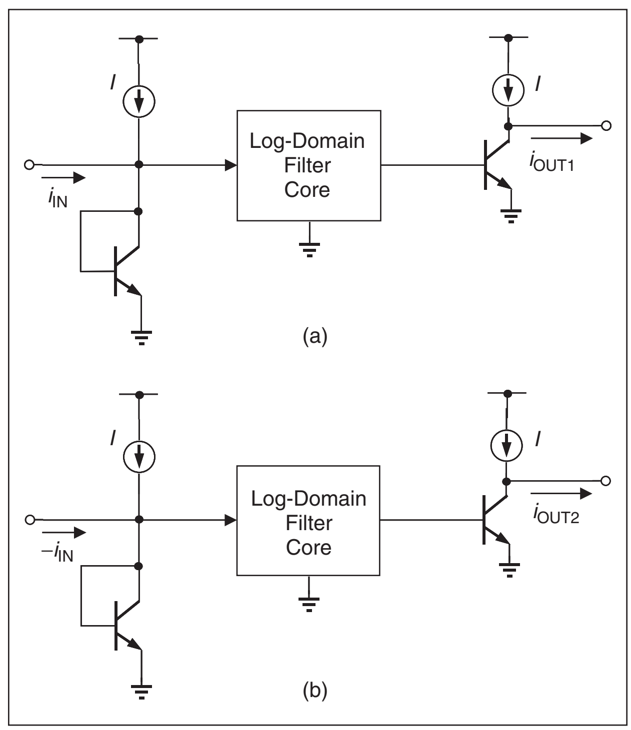

9. The circuit in Figure 7(a), with transconductors replaced by bipolar transistors and the main filter replaced by a log-domain core; two paths, driven differentially, are shown, with the output assumed to also be taken differentially.

toward the right is due to increase in distortion. Consider the example presented earlier, where a total desired UDR of 60 dB resulted in a capacitance of 20,000 pF and a power dissipation of 10 W. Using instead the divide-and conquer technique, each of the filters need only handle a UDR of 20 dB, and it thus requires only 2 pF and 1 mW! Such savings more than make up for the fact that three filters plus supporting circuitry are used. Of course, if blockers are present, the required headroom (see above) will reduce the savings that can be obtained.

Companding techniques can also be applied to switched-capacitor circuits [31].

Technique 2: Dynamic Impedance Scaling

Consider a system designed to handle very small signals in the worst case, as in Figure 4(a). Due to the large UDR, the capacitance values will be large, as seen from (6). However, this worst-case signal situation may not be present continuously; part of the time, this same system may be called to handle large signals with limited UDR or SNR_{max}. This, for example, can be the case in wireless systems, when the desired signal happens to be strong and no strong blocker is present [32]. At those instances, the very low noise floor of this system is unnecessary and can be allowed to increase by scaling C down [33], resulting, for example, in the situation shown in Figure 4(b). Of course, the transconductances or conductances of the filter will have to be scaled together with C, in order to leave the frequency response unchanged. The reason one may want to do such impedance scaling becomes clear from (7): lowering capacitances and transconductances allows one to also lower I_{BIAS}, and thus the total power dissipation can be made small. This can also be seen from the transconductor characteristics in Figure 3(d). In principle, for every 3 dB reduction in UDR, the power dissipation can be cut to half.

One way to achieve the above scenario would be to use two filters, designed with different capacitance values, and place one or the other in service as required [34], turning the other filter off. This, however, causes transients at the output, which may take a long time to die out, especially in the case of filters with high-Q poles; the result may be loss of data.

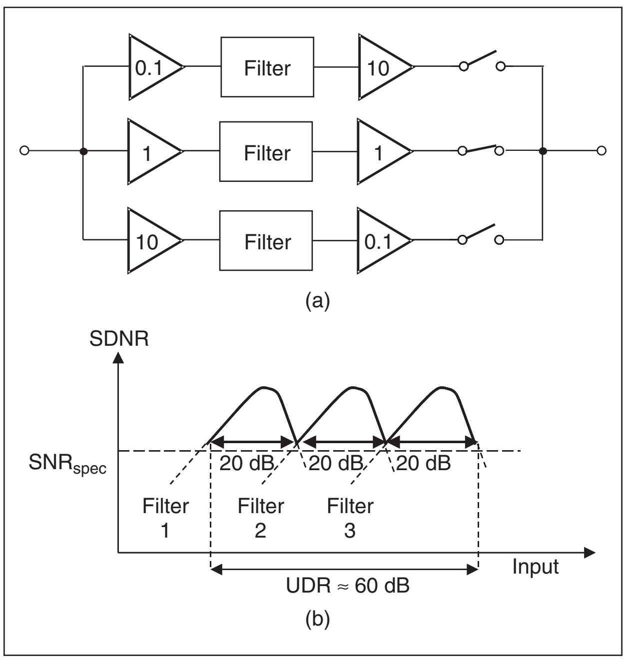

10. (a) "Divide-and-conquer" signal processor; (b) resulting signal-to-noise-plus-distortion ratio (SNDR) versus input RMS value.

Avoiding Output Transients

Delayed output switching: To prevent the system from seeing the output transients, one can use the approach in Figure 11. Before the output is to be switched, both filters are switched on; after the outputs of both filters become the same, the output can be handed off to the second filter, and then the first filter can be turned off [30].

Individual element switching: One can diminish the delay involved above by switching individual filter elements rather than entire filters. An example is shown in Figure 12. Here, when it is acceptable to increase the noise floor, the upper resistor and upper capacitor are switched out of the circuit, and the op amp's bias can be reduced (see also "dynamic biasing" below); the system now performs as in Figure 4(b). When the system is needed to process again signals with a low noise floor, the above two elements should be switched in; in order to avoid output transients, the op-amp bias is first increased; then the upper capacitor is switched between output and ground, so that its voltage becomes equal to that of the lower capacitor; and, finally, the two capacitors are connected in parallel. The system now performs as in Figure 4(a).

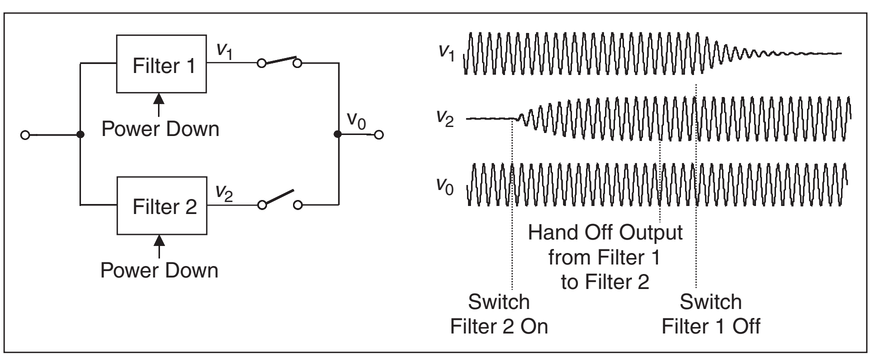

11. Method to switch between two signal processors with the same input-output behavior, without causing output transients.

Technique 3: Dynamic Biasing

The fact that a circuit has been designed to handle large signals does not necessarily mean that such signals will always be present. When they are not (i.e., when $V_{s,max}$ happens to be small) the bias currents can be reduced, as seen from (7). In many cases it is desirable to do this dynamically, without interrupting the filter operation. For example, in the transconductor of Figure 3 all four bias currents can be scaled down together, without causing major output disturbances, provided that the transconductance of the MOSFETs remains sufficiently larger than the conductance of the resistor (if a triode-operated MOSFET is used in lieu of the resistor [11], its gate voltage is assumed to be referenced to its source and drain potentials); a second-order variation of its transconductance will be corrected by the automatic tuning system of the filter. This type of dynamic biasing leads to changes in the transconductor characteristic as shown in Figure 3(c). The biasing in certain op-amp architectures can similarly be varied without causing major disturbances at the output. In general, dynamic biasing decreases bias current values when they are not needed to handle large signal swings. The total current in a conventional and in a dynamically biased class A circuit, under changing signal amplitudes conditions, is illustrated in Figure 13. This technique is sometimes used in power amplifiers, but it is extendable to other circuits. The bias current is adapted to the average value of the signal, in a "syllabic" manner. (Another possibility is to use class AB, class B, or other related circuits in which the supply current is a function of the instantaneous value of the signal. These are not internally varying circuits in the sense considered in this article.)

We note that dynamic biasing is also encountered in the course of applying companding [27], [20], [28] (Figures 7 and 9) and/or dynamic impedance scaling (Figure 12).

Voltage-Mode Versus Current-Mode: A Key Difference

We have seen above that a relaxed SNR can result in power savings, if things are done properly. The mechanisms through which such savings can be accomplished can be very different from case to case. This will be illustrated by an example.

Consider first a voltage-mode circuit, such as a voltage amplifier or the filters in Figure 2. As was seen above, if the impedance levels are raised (thus making possible smaller power dissipation), the noise increases. The maximum signal that can be handled is not affected by impedance scaling. Thus, signal and noise versus power dissipation behave qualitatively as in Figure 14(a).

Consider now a current-mode circuit, such as a current mirror. If the bias current is decreased, the shot noise current (the rms value of which is proportional to the square root of the bias current) decreases as well; thus the behavior of noise with power dissipation is opposite from that in voltage-mode circuit. However, with decreasing bias current the maximum signal swing decreases proportionately; i.e., faster than the rms noise current decreases. The situation is thus as in Figure 14(b). It is seen that the decreased noise made possible by lowering the bias current can only be taken advantage of if the signal current happens to be small.

Technique 4: Dynamic Structure Variation

A fourth possibility is to dynamically vary the very structure of the circuit (e.g., change the architecture of a filter dynamically). A straightforward way to accomplish this is to use two or more filters with different structures and make sure that output transients are avoided using the technique in Figure 11. It may also be possible to use a single filter and vary its structure dynami-

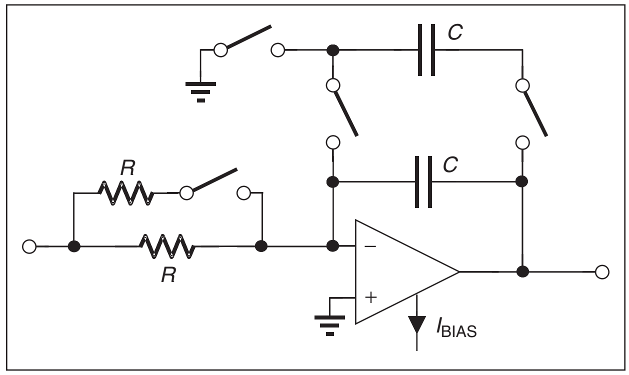

12. An active RC circuit in which dynamic impedance scaling can be performed without causing output transients.

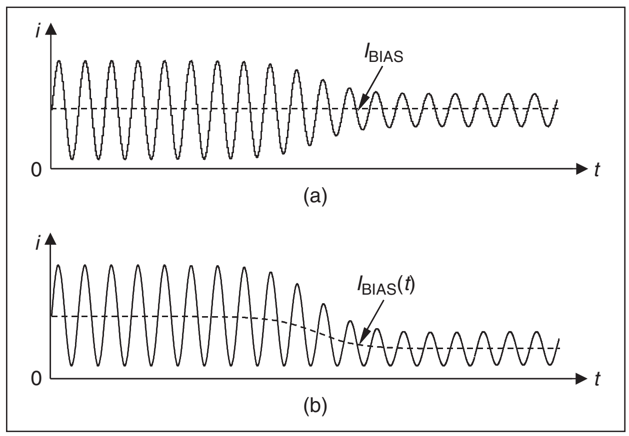

13. Typical bias and total signal currents in (a) a conventional circuit and (b) a dynamically biased circuit.

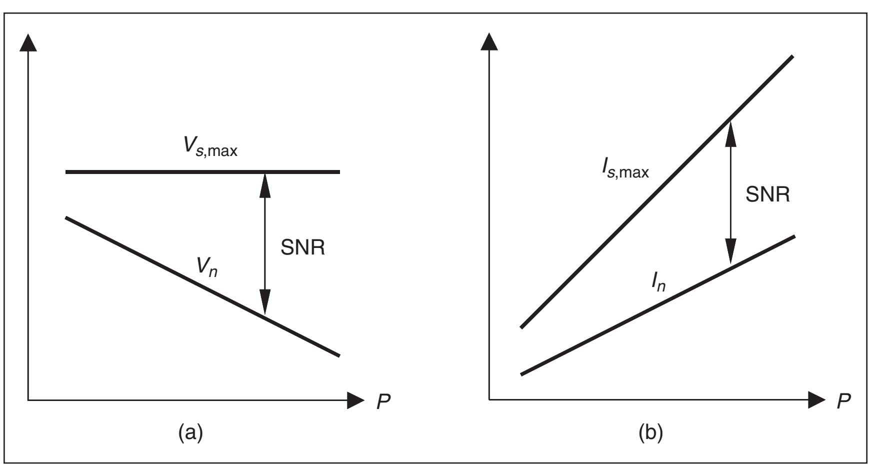

14. Mechanism through which the signal-to-noise ratio is improved with increasing power dissipation in (a) a voltage-mode circuit and (b) a current-mode circuit.

cally, using techniques related to the ones discussed. This approach has not been investigated to any extent yet.

Conclusions

We have reviewed several techniques that make possible the dynamic variation of analog circuits internally, without affecting their input-output characteristics. A mixture of more than one of these techniques is appropriate in some cases. We used filters as a specific example of dynamical analog circuits and placed particular emphasis on avoiding or eliminating transients at the output of such circuits, which would normally have occurred due to such dynamic variations. By allowing for dynamic internal variations, the power dissipation of such circuits can be lowered and can be made to depend on how demanding the task at hand is. This allows for large savings of energy drain over time, thus making possible long battery life in portable equipment.

Acknowledgments

This work was supported by NSF Grant CCR-99-02781.

Y. Tsividis is with Columbia University, New York, New York, USA. *N. Krishnapura, Y. Palaskas*, and *L. Toth* were with Columbia University; they are now with, respectively, Multilink Corp., Somerset, New Jersey, USA; Intel Corp., Hillsboro, Oregon, USA; and the Technical University of Budapest, Budapest, Hungary.

References

[1] D. Blom and J.O. Voorman, "Noise and dissipation of electronic gyrators," *Philips Res. Rep.*, vol. 26, pp. 103-113, 1971.

[2] E. Vittoz, "Future of analog in the VLSI environment," in *Proc. 1990 IEEE ISCAS*, New Orleans, LA, 1990, pp. 1372-1375.

[3] J.O. Voorman, "Continuous-time analog integrated filters," in *Integrated Continuous-Time Filters*, Y.P. Tsividis and J.O. Voorman, Eds. Piscataway, NJ: IEEE Press, 1993, pp. 15-46.

[4] G. Groenewold, B. Mona, and B. Nauta, "Micropower filter design," in *Analog Circuit Design: Low-Power, Low-Voltage, Integrated Filters, and Smart Power*, R.J.v.d. Plassche, W.M.C. Sansen, and J.H. Huijsing, Eds. Boston, MA: Kluwer, 1995, pp. 73-88.

[5] E.A. Vittoz and Y. Tsividis, "Frequency—dynamic range—power," in *Trade-Offs in Analog Circuit Design*, C. Toumazou, B. Gilbert, and G. Moschytz, Eds. Boston, MA: Kluwer, 2002, ch. 11.

[6] A. Papoulis, *Probability, Random Variables, and Stochastic Processes.* New York: McGraw-Hill, 1991.

[7] H.J. Orchard, "Inductorless filters," *Electron. Lett.*, vol. 2, pp. 224-225, June 1966.

[8] J. Tow, "ActiveRC filters—A state-space realization," *Proc. IEEE*, vol. 56, pp. 1137-1139, 1968.

[9] L.C. Thomas, "The Biquad: Part I—Some practical design considerations," *IEEE Trans. Circuits Syst.*, vol. CT-18, pp. 350-357, May 1971.

[10] R. Schaumann and M.E. Van Valkenburg, *Design of Analog Filters.* Oxford, U.K.: Oxford Univ. Press, 2001.

[11] Y. Tsividis, Z. Czarnul, and S.C. Fang, "MOS transconductors and integrators with high linearity," *Electron. Lett.*, vol. 2, no. 5, pp. 245-246, Feb. 1986.

[12] D.K. Shaeffer, A.R. Shahani, S.S. Mohan, H. Samavati, H.R. Rategh, M. del Mar Henderson, M. Xu, C.P. Yue, D.J. Eddleman, and T.H. Lee, "A 115-mW, 0.5-um CMOS GPS receiver with wide dynamic range active filters," *IEEE J. Solid-State Circuits,* vol. 33, pp. 2219-2231, Dec. 1998.

[13] T.B. Cho, G. Chien, F. Brianti, and P.R. Gray, "A power-optimized CMOS baseband channel filter and ADC for cordless applications," *Dig. Tech. Papers, 1996 Symp. VLSI Circuits*, Honolulu, HI, June 1996, pp. 64-65.

[14] G. Palaskas and Y. Tsividis, "An optimization technique for maximizing the dynamic range of continuous time active filters," to be published.

[15] R.C. Mathes and S.B. Wright, "The compandor—An aid against static in radio telephony," *Bell Syst. Tech. J.*, vol. XIII, pp. 315-332, July 1934.

[16] R.M. Dolby, "Signal compressors and expanders," U.S. Patent 3,345,416, Oct. 1974.

[17] Y.P. Tsividis, V. Gopinathan, and L. Toth, "Companding in signal processing," *Electron. Lett.*, vol. 26, pp. 1331-1332, Aug. 1990.

[18] Y.P. Tsividis, P.R. Gray, D.A. Hodges, and J. Chacko, Jr., "A segmented u-255 law PCM voice encoder utilizing NMOS technology," *IEEE J. Solid-State Circuits*, vol. SC-11, pp. 740-747, Dec. 1976.

[19] D.V. Mercy, "A review of automatic gain control theory," *Radio Electron. Eng.*, vol. 51, pp. 579-590, Nov./Dec. 1981.

[20] Y. Tsividis, "Minimizing power dissipation in analog signal processors through syllabic companding," *Electron. Lett.*, vol. 35, pp. 1805-1807, Oct. 1999.

[21] E. Blumenkrantz, "The analog floating point technique," in *Proc. 1995 IEEE Symp. Low-Power Electronics*, San Jose, CA, Oct. 1995, pp. 72-73.

[22] Y. Tsividis, "Externally linear, time-invariant systems and their application to companding signal processors," *IEEE Trans. Circuits and Systems II*, vol. 44, pp. 65-85, Feb. 1997.

[23] G. Palaskas and Y. Tsividis, "Design considerations and experimental evaluation of a syllabic companding audio frequency filter," in *Proc. 2002 Int. Symp. Circuits and Systems*, Phoenix, AZ, 2002, vol. 3, pp. 305-308.

[24] L. Toth, G. Palaskas, and Y. Tsividis, "Two 'noninvasive' techniques for syllabic companding in signal processors," *IEEE Trans. Circuits and Systems*, vol. 48, pp. 1085-1098, Dec. 2001.

[25] E. Seevinck, "Companding current-mode integrator: A new circuit principle for continuous-time monolithic filters," *Electron. Lett.*, vol. 26, pp. 2046-2047, Nov. 1990.

[26] D.R. Frey, "Log-domain filtering: An approach to current-mode filtering," *Proc. Inst. Elect. Eng. –G*, vol. 140, pp. 406-416, Dec. 1993.

[27] D.R. Frey and Y.P. Tsividis, "Syllabic companding log domain filter using dynamic biasing," *Electron. Lett.*, vol. 33, pp. 1506-1507, 1997.

[28] N. Krishnapura and Y. Tsividis, "Noise and power reduction in filters through the use of adjustable biasing," *IEEE J. Solid-State Circuits*, vol. 36, pp. 1912-1920, Dec. 2001.

[29] G. Palaskas and Y. Tsividis, "A 'divide and conquer' technique for the design of wide dynamic range continuous-time filters," in *Proc. 2001 Int. Symp. Circuits and Systems*, Sydney, Australia, May 2001, vol. 1, pp. 252-255.

[30] Y. Tsividis, "Self-tuned filters," *Electron. Lett.*, vol. 17, no. 12, pp. 406-407, June 1981.

[31] N. Krishnapura, Y. Tsividis, and K. Suyama, "Companding switched-capacitor filters," in *Proc. 1998 Int. Symp. Circuits and Systems*, Monterey, CA, May 1998, vol. 1, pp. 480-483.

[32] J.C. Rudell, J.-J. Ou, R.S. Narayanaswami, G. Chien, J.A. Weldon, L. Lin, K.-C. Tsai, L. Tee, K. Khoo, D. Au, T. Robinson, D. Gerna, M. Otsuka, and P.R. Gray, "Recent developments in high integration multi-standard CMOS transceivers for personal communication systems," in *Proc. 1998 Int. Symp. Low Power Electronics and Design*, Monterey, CA, 1998, pp. 149-154.

[33] P.R. Gray and R.G. Meyer, "Future directions in silicon ICs for RF personal communications," in *Proc. 1995 IEEE Custom Integrated Circuit Conf.*, San Diego, CA, May 1995, pp. 83-90.

[34] F. Behbahani, A. Karimi-Sanjaani, W.-G. Tan, A. Roithmeier, J.C. Leete, K. Hoshino, and A.A. Abidi, "Adaptive analog IF signal processor for a wide-band CMOS wireless receiver," *IEEE J. Solid-State Circuits*, vol. 36, pp. 1205-1217, Aug. 2001.

CD■