

Syllabic-Companding Log Domain Filters

Douglas Frey, *Senior Member, IEEE*, Yannis P. Tsvividis, *Fellow, IEEE*, Giorgos Efthivoulidis, and Nagendra Krishnapura

Abstract—A general theory for companding log domain filters is proposed which combines not only exponential mappings, but also a new translational mapping approach which guarantees suitable operating conditions in any log domain filter. The filter equations resulting from the use of the theory ultimately contain translinear terms which are known to be realizable using translinear techniques. A discussion of the design of the companding filters, regarding the economical use of translinear loops and the convenient selection of system parameters, is offered which leads to first- and second-order circuit designs. Finally, the noise performance of an example design is investigated using a carefully crafted large-signal simulation technique, showing clearly the advantage of the companding filter approach.

Index Terms—Analog filters, companding, log domain filters, system theory.

I. INTRODUCTION

RECENTLY, there has been considerable interest in so-called externally linear, internally nonlinear (ELIN) filters [1], in which we might include classical linear filters as a subset. Because ELIN filters are in general quite nonlinear internally, classic linear filter theory is not directly applicable. Nevertheless, a fairly complete theory for externally linear filters has begun to develop, with notable contributions found in [2]–[4], regarding state space synthesis, in [5] and [6] regarding dynamic translinear analysis, and in [7] and [8] regarding modular analysis, although a number of interesting contributions have been made elsewhere. The current interest in ELIN filters was spawned in part by the pioneering work found in [9] and [10]. Concurrent to some of this work was that in [11] and [12] regarding syllabic companding filters. An interesting way to compare the bulk of the recent work, primarily related to log domain filters, and that of [11] and [12] is to consider log domain filters to be instantaneously companding, which has been the choice of several researchers. Unfortunately, such a characterization clouds the concepts if one chooses to investigate syllabic companding log domain filters, as in [13] and [14]. In fact, the goal of this paper is to propose a general approach to the design of syllabic companding log domain filters, which we shall assume are another subclass of ELIN

filters. The reader may find the discussion in [15] interesting regarding companding in log domain filters. The value of syllabic-companding log domain filters lies in their potential for high speed and/or low power, low circuit complexity, and wide dynamic range.

In a channel with limited dynamic range, such as an active filter, there is always a tradeoff to be made between signal-to-noise ratio (SNR) and headroom. Specifically, the SNR is improved by preamplification, which minimizes the effective contribution of the channel noise. Even if this gain is unwanted, it may be removed by attenuation after the noisy channel, thereby restoring the signal level while attenuating the channel noise. The price paid for this is that large inputs that would have been acceptable as inputs to the channel will now be distorted by the channel due to the preamplification. A straightforward resolution for this dilemma has been known and adopted for many years and is known as companding. This method varies the preamplifier gain as a function of signal level in such a way as to keep the signals applied to the channel near the top end of its dynamic range at all times, thereby improving the SNR without sacrificing headroom. It is important to realize, however, that in order for the overall system response to be unaffected by the variations in preamplifier gain, a post-amplifier must be used to compensate for the front end gain variations.

A very well-known example of a successful companding system is the Dolby system [16]. An important aspect of the Dolby system is that it is a “syllabic” companding system. Namely, it varies the preamplifier gain in accordance with the short term average—that is, on a syllabic time scale—of the signal as opposed to the instantaneous level of the signal. This is crucial to the performance of the system, since the instantaneous level of the signal at each point in a system is a complex function of the frequency and phase response of the subsystems. In systems having the property that the magnitude response and group delay is fairly flat for each subsystem, instantaneous companding can be implemented quite successfully. However, most channels, and certainly active filters, do not possess this property. Therefore, the right amount of instantaneous gain needed to optimize one part of a system may not work for another part. While one may envision systems with many variable gain blocks, in practice only one main gain control signal is used. Hence, a measure of signal intensity that is insensitive to at least group delay variation in the passband is desirable. RMS detection is a widely used option.

Unfortunately, rms detection is not a perfect solution to the problem of level detection in companding systems. This is because it inherently introduces delay between the actual signal

Manuscript received February 1999; revised March 2001. This work was supported in part by the National Science Foundation under Grant MIP-9703 773 and Grant MIP-9902781. This paper was recommended by Associate Editor B. Nauta.

D. Frey is with the Electrical Engineering and Computer Science Department, Lehigh University, Bethlehem, PA 18015 USA (e-mail: drf3@lehigh.edu).

Y. P. Tsvividis and N. Krishnapura are with the Department of Electrical Engineering, Columbia University, New York, NY 10027 USA.

G. Efthivoulidis is with the Department of Electrical and Computer Engineering, National Technical University of Athens, Greece.

Publisher Item Identifier S 1057-7130(01)05225-9.

level at a given time and its rms level estimate. As long as the level of the signal varies slowly compared to the time constant of the rms detector, no problems arise. However, when the input signal level suddenly increases, it can overload the system until the rms detector responds. In practical syllabic-companding systems, such as the Dolby system, extra circuitry is used to minimize the transient overload to a point where it is acceptable to the user. In this way, users get the benefit of a perceived increase in dynamic range. The goal of this paper is to gain the advantages of syllabic companding for active, especially log domain, filters. It is recognized, however, that syllabic-companding filters will suffer the same transient overload problems encountered in any syllabic-companding system.

The idea of syllabic-companding filters has been proposed in [11]. The associated technique systematically varies the gain of an input applied to a filter along the lines of the above discussion; however, changes must be made internal to the filter to correct for changes in the dynamics of the filter due to the time-varying input gain. Syllabic-companding filtering using this principle has been described in [12] and has been used in conjunction with log domain filtering in [13]. In the latter implementation, the bias currents of the main filter were kept constant. However, as suggested in preliminary form in [14], companding can also be done by dynamically adjusting the bias currents in such a way as to minimize the standing currents required during intervals of relatively constant signal strength. There it has been suggested that one may view the variation of bias currents in response to signal level as the dual operation to gain change. In particular, one may adjust the gain of the front end of the system to ensure that the signal is at the top end of the dynamic range, or equivalently, one may vary the top end of the dynamic range of the system to match the signal level. As long as the noise floor of the system moves appropriately with variations in the top end of the dynamic range, then the SNR will be improved as in companding using explicit gain variations. Of course, the dynamic bias current adjustment must be done in such a way that the external behavior of the filter remains linear.

This paper will propose a complete theory, and an accompanying design approach, for the realization of companding log domain filters based on our dynamic biasing approach. Section II begins by introducing the idea of translational mapping to the state space equations for a filter. This idea is the key to the systematic design of companding filters from our perspective. It is also valuable in improving earlier methods for the design of log domain filters. We follow this with a discussion of parameter selection. In Section III, the exponential mappings for log domain filters are applied to the linearly mapped, via translation, state equations. This is shown to yield nodal equations for companding log domain filters containing various translinear terms, which are realizable using translinear loops. A general set of equations for a companding filter are shown and two special cases—that is, a first-order and a second-order filter—are presented to aid in developing the concepts. Finally, in Section IV, carefully developed SPICE simulation results, using transient analysis, are given regarding the noise behavior of one of the example circuits proposed in Section III.

II. A TRANSLATIONAL MAPPING APPROACH

A. Introduction

In light of the discussion above, we now offer a systematic formulation for systems explicitly incorporating dc bias currents to establish the dynamic range and, in particular, syllabic-companding log domain filters. We assume that state equations for a given filter are specified in the following standard form:

$$\frac{d}{dt}\bar{x} = A\bar{x} + \bar{b}u; \quad y = \bar{p}^T\bar{x} + du \quad (1)$$

where the input u and the output y are assumed scalars, $\bar{x} = (x_1, x_2, \dots, x_N)^T$ is the state vector, A is the N by N state matrix, \bar{b} and \bar{p}^T are N dimensional vectors, and d is a scalar. The system of (1) possesses a transfer function, $H(s)$, dependent upon the parameters in (1) given by the formula

$$H(s) = \bar{p}^T(sI - A)^{-1}\bar{b} + d. \quad (2)$$

In anticipation of the log domain realizations to be discussed later, we will require that the input be a strictly positive signal so that it may be applied to a diode, or equivalent combination of transistors, which provides the logarithm of the input. Typically, this input signal would be the sum of an original ac input signal and an added offset. Since the output is typically the current in an output transistor, we recognize that an always positive input u will have to result in an always positive output y . In the process, each of the state variables will have to be strictly positive for all times, since these variables correspond to currents in transistor junctions as well.

Ensuring always positive state variables and output is not always a trivial matter. Consider the case where (1) describes an ideal bandpass filter, with output y equal to x_1 , as shown explicitly in the second-order example of Section III. Suppose that a sinusoidal ac input is to be applied to the filter, and that for the purpose of log domain filtering, a dc component is added to the ac input sufficient in size to guarantee that the composite input u is always strictly positive. Since this is a linear system, the superposition principle states that the output corresponding to this composite input will consist of a dc and an ac component, where each component is determined by the gain of the system at the respective frequencies. Specifically, the output ac component will have a magnitude and phase determined by the transfer function $H(s)$, evaluated at its frequency. On the other hand, the dc component of the output is determined by the dc gain of the system. Since we have taken the case of a bandpass filter, the ac gain will be nonzero, while the dc gain will be equal to zero. Hence, adding dc to the input in this case does not guarantee that there will be sufficient dc in the output. Clearly in this case the output, which equals one of the state variables, will be a purely ac signal—that is, a signal having positive and negative values over time—which is unacceptable in log domain filters.

B. The Translational Mapping

In order to address this problem, let us consider a translational mapping applied to (1) which will allow us to guarantee that all state variables will remain positive at all times. Suppose that

we adopt the change of variables in the state equations of (1) where specified below as

$$\begin{aligned} x_i &= X_i - X_{0i}, \quad 1 \leq i \leq N \\ u &= U - U_0 \\ \bar{X} &= (X_1, X_2, \dots, X_N)^T \\ \bar{X}_0 &= (X_{01}, X_{02}, \dots, X_{0N})^T \end{aligned} \quad (3)$$

where each of X_{0i} and U_0 are (possibly) time-varying offsets chosen such that the new variables X_i and U are always positive. With these substitutions, (1) becomes

$$\frac{d}{dt} \bar{X} = A\bar{X} + \bar{b}U + \frac{d}{dt} \bar{X}_0 - A\bar{X}_0 - \bar{b}U_0 \quad (4a)$$

$$y = \bar{p}^T \bar{X} + dU - \bar{p}^T \bar{X}_0 - dU_0. \quad (4b)$$

Observe that the transfer function from the original input u to the output y is unchanged; however, the translated internal variables X_i and U are now always positive regardless of the overall transfer function. Therefore, the internal variables are completely compatible with the exponential mappings used for the realization of log domain filters [17]. The extra terms introduced into the equations of (4), relative to those of (1), by the translational mappings appear as additional inputs and, as such, will ultimately be realized in a way similar to the nominal input U .

C. The Noncompanding Case

Before exploring the central purpose of this paper regarding companding filter design, it is both of interest and helpful in understanding to consider the special case where the offsets X_{0i} and U_0 specified above are constants, which corresponds to the noncompanding case typically covered in the literature thus far. The choice of the offsets for this case is very instructive. Let us assume that the input offset U_0 is chosen equal to the largest peak excursion of the input, plus a safety margin, such that the composite input is always positive and greater than zero by some minimum value. This will guarantee that the input logging device in the log domain filter will always run with at least a sufficient current to guarantee good performance, e.g., good bandwidth. Continuing, let us choose the state variable offsets X_{0i} to be multiples of the input offset, determined by the maximum ac gain that a signal may have at any frequency. Such gains are found, for example, by using the intermediate transfer functions discussed in [18]. In this way, the composite state variables X_i will be bounded away from zero by some amount which will not only guarantee that they are strictly positive, but also that the currents in the associated transistors in the eventual log domain filters will never get too close to zero. Denoting the respective scale factors for each state variable by λ_i , we have the following state space formulation for the constant-translated system:

$$\frac{d}{dt} \bar{X} = A\bar{X} + \bar{b}U - (A\bar{\lambda} + \bar{b})U_0 \quad (5a)$$

$$y = \bar{p}^T \bar{X} + dU - (\bar{p}^T \bar{\lambda} + d)U_0 \quad (5b)$$

$$\bar{\lambda} = (\lambda_1, \lambda_2, \dots, \lambda_N)^T$$

$$x_i = X_i - \lambda_i U_0$$

$$u = U - U_0. \quad (6)$$

The net result of the translational mapping, as shown in (5), has been to introduce an additional constant input, proportional to U_0 , into each of the equations. These inputs are completely analogous to the ‘‘dummy’’ inputs introduced in [3] to permit a suitable dc equilibrium solution. Here, however, the introduction of the variables is done in a way that guarantees that no internal variables may be over driven.

Since each of the terms in the equations may result in some additional circuitry in the ultimate circuit realization, it is usually desirable to suppress the new terms where possible. Using the state space transformations described elsewhere, this can be systematically achieved. Specifically, if a nonsingular state space transformation matrix M is used as in [2] to produce a new set of equations, we may in the process eliminate the vector $A\bar{\lambda} + \bar{b}$. Using the results shown in [2], for example, on the impact of this transformation on A and \bar{b} , we have

$$\begin{aligned} A &\rightarrow MAM^{-1} \\ \bar{b} &\rightarrow M\bar{b} \\ (A\bar{\lambda} + \bar{b}) &\rightarrow M(AM^{-1}\bar{\lambda} + \bar{b}). \end{aligned} \quad (7)$$

Therefore,

$$M(AM^{-1}\bar{\lambda} + \bar{b}) = 0 \Rightarrow \bar{\lambda} = -MA^{-1}\bar{b}. \quad (8)$$

The final result in (8) states that there are many ways to choose M such that the vector of constants $A\bar{\lambda} + \bar{b}$, added to the state equations, as shown in (5), due to the translation of the variables will be identically zero. This constraint is a generalization of that given in [2] for allowing a suitable dc equilibrium. By using (8), however, not only is a suitable dc equilibrium guaranteed, but also the needed headroom for each of the state variables in consideration of the intermediate transfer functions [18]. The result in (8) is identical to the analogous one in [2] for the special case where each of the ac gains λ_i is assumed to be equal to 1. In all cases below, we shall assume that the transformation, using a nonsingular matrix M , has already been applied, leaving a convenient state space description of the form of (1). In general, one may arrive at a convenient form after some trial and error, although the examples given in Section III are typical cases where little or no manipulation is required.

D. The Companding Case

Now let us move on to the case where the offsets in the translational mappings are time-varying, corresponding to the syllabic-companding case, which is the main case we wish to consider. We may proceed exactly as before in relating the state variable offsets X_{0i} to the input offset, U_0 through a vector of gain constants $\bar{\lambda}$ with the result

$$\frac{d}{dt} \bar{X} = A\bar{X} + \bar{b}U + \bar{\lambda} \frac{d}{dt} U_0 - A\bar{\lambda} U_0 - \bar{b}U_0 \quad (9a)$$

$$y = \bar{p}^T \bar{X} + dU - (\bar{p}^T \bar{\lambda} + d)U_0. \quad (9b)$$

Note that, while the output equation is unchanged, the state equations now include a vector of new inputs proportional to the time derivative of U_0 . The likelihood of such a quantity being available directly as a voltage or a current is small; however, by the use of an additional first-order filter, which we shall call “the auxiliary filter,” we can eliminate this problem. Specifically, suppose that the offset U_0 was obtained by a first-order lowpass filtering operation on another input U_R . In practice, U_R would be the output of a level detector operating on the ac input to the system plus some relatively small dc component added to ensure that U_R , and hence U_0 , would remain above some minimum acceptable value. The level detector employed would typically comprise a rectifier, a peak and/or rms detector, and some scaling to account for crest factor. Due to the first-order lowpass filtering, the time-varying offset U_0 would obey a differential equation of the form

$$\frac{d}{dt}U_0 = -\omega_R U_0 + \omega_R U_R. \quad (10)$$

We will refer to the filter described by (10) as the “auxiliary filter” in the discussion below, since it leads to additional circuitry in the final companding filter implementation. Using the result in (10) to eliminate the derivative of U_0 in (9), we have

$$\begin{aligned} \frac{d}{dt}\bar{X} &= A\bar{X} + \bar{b}U + \bar{\lambda}(-\omega_R U_0 + \omega_R U_R) \\ &\quad - A\bar{\lambda}U_0 - \bar{b}U_0 \\ &= A\bar{X} + \bar{b}U - (\omega_R \bar{\lambda} + A\bar{\lambda} + \bar{b})U_0 + \omega_R \bar{\lambda}U_R \\ &= A\bar{X} + \bar{b}U - \bar{b}_0 U_0 + \bar{b}_R U_R \end{aligned} \quad (11a)$$

$$y = \bar{p}^T \bar{X} + dU - (\bar{p}^T \bar{\lambda} + d)U_0 \quad (11b)$$

where

$$\bar{b}_0 = \omega_R \bar{\lambda} + A\bar{\lambda} + \bar{b}; \quad \bar{b}_R = \omega_R \bar{\lambda}. \quad (12)$$

This result shows that, for the companding case, we introduce as many as two additional new inputs to each state equation; however, each of these new inputs is proportional to an undifferentiated system signal. We will refer to the filter described by (11) as the “main filter” in the discussion below.

As in the case of constant offsets, the filter described by (11), having the auxiliary filter described by (10), may be simplified in its final implementation by using certain degrees of freedom which are available. First, one may employ a transformation using a nonsingular matrix, as in (7) and (8), to eliminate some or all of the components of the vector \bar{b}_0 . Alternatively, one may vary the components in the vector $\bar{\lambda}$ to reduce the number of nonzero components of \bar{b}_0 . Note that, while there is a minimum value allowed for each of the components of $\bar{\lambda}$ to ensure enough headroom for the respective state variables, any of the components of $\bar{\lambda}$ may be increased above this minimum without causing clipping problems. On the other hand, increasing the components of $\bar{\lambda}$ will increase the standing currents in transistors in the ultimate log domain circuit realization, causing an increase in noise. This increase in noise, however, may be compensated for by the decrease in circuit complexity resulting from fewer terms in the state equations.

The other design parameter which may be varied to eliminate nonzero terms in the main filter state equations is ω_R —that is, the cutoff frequency of the auxiliary filter. However, this parameter must be varied in consideration of system specifications. This relates to the fact that syllabic companding, as explained in the introduction, requires that the level of the system input signal u be monitored and used to control the companding gains. Specifically, the offset U_0 produced at the output of the auxiliary filter must correspond to the desired control signal for proper companding. Since level detectors used in syllabic-companding systems inherently employ lowpass filtering, it is reasonable to assume that the auxiliary filter, in our case, may provide that function. Certainly, this would be economical, since the auxiliary filter would now perform the joint functions of removing unwanted derivative terms in the main filter and of envelope (as it is often called) filtering for the syllabic-companding operation. Unfortunately, however, it is difficult to achieve these functions simultaneously in practice. This is because the cutoff frequency of the envelope filter is usually set much lower than the critical frequencies of the main filter to avoid undesirable “pumping effects”—that is, signal-related noise modulation—in the overall companding system. Using typical scenarios, it has been our observation that the values of ω_R , needed to implement a desirable envelope filter, do not typically enable the elimination of unwanted terms in the main filter equations. As a result, we will take the approach in this paper that ω_R is a free parameter, allowing the simplification of the main filter. We will assume, therefore, that the overall envelope filtering, which is necessary to achieve the desired syllabic-companding operation, will be provided in the level detection circuitry that produces the signal U_R designed in consideration of the filtering action of the auxiliary filter.

III. LOG DOMAIN FILTER REALIZATION

A. Introduction

Given the above preliminaries, we are now in a position to design log domain filter realizations of syllabic-companding filters. In order to facilitate the design, we will offer a new formulation which combines the exponential mappings introduced in [2] with the translinear ideas described elsewhere—for example, in [5], [6], and [19]. First, observe that the translational mappings in Section II provide a simple systematic way to guarantee that all of the system variables—that is, U , and each of the state variables—will remain strictly positive at all times. Therefore, if exponential mappings are now used to substitute for all of the system variables, then these mappings will be well defined. Such mappings have been shown—for example, in [2] and [3]—to produce the circuit equations for log domain filters, where all system variables are assumed to be currents. Since the synthesis of log domain filters is our goal, consider the following mappings for the system variables in (10) and (11):

$$\begin{aligned} U &= I_s e^{V_0/V_t} \\ X_i &= I_s e^{V_i/V_t}, \quad \forall 1 \leq i \leq N \\ U_0 &= I_s e^{V_B/V_t} \\ U_R &= I_s e^{V_R/V_t} \end{aligned} \quad (13)$$

where I_s and V_t are constants typically equal to the reverse saturation current and thermal voltage, respectively, of a bipolar transistor. Using (13), we have

$$\begin{aligned} C \frac{d}{dt} V_i &= \frac{C V_t}{X_i} \frac{d}{dt} X_i \\ C_R \frac{d}{dt} V_B &= \frac{C_R V_t}{U_0} \frac{d}{dt} U_0 \end{aligned} \quad (14)$$

where the constants C and C_R have units of capacitance such that their multiplication with the derivatives of the voltages—that is, V_0, V_1 , etc.—yields currents.

B. A First-Order Example

Equations (13) and (14) may be substituted into (10) and (11), yielding the equations necessary for a general log domain companding filter, excluding the level detection circuitry. Before giving this general formulation, however, let us consider a simple first-order case which highlights the important ideas with a minimum of complexity. Specifically, let us begin with a first-order lowpass filter, characterized by the following equations:

$$\begin{aligned} \frac{d}{dt} x &= -\omega_0 x + k \omega_0 u \\ y &= x, \quad A = -\omega_0, \bar{b} = k \omega_0, \bar{p} = 1, d = 0. \end{aligned} \quad (15)$$

Notice that the operators A, \bar{b}, \bar{p} , and d from (1) are simply scalars in this case and have been given parenthetically in (15). Because this is a first-order filter with dc gain equal to k , the maximum gain to the state variable x from the input u is k . Any value of the gain constant λ (scalar in this case) that equals or exceeds this value will ensure adequate headroom in the final implementation. Thus, adopting the offsets U_0 and λU_0 for the input u and the state variable x , respectively, in (3), we obtain new system equations where all system variables are strictly positive as long as U_0 is at least as large as the maximum peak value of u . Proceeding along the lines above, we have

Main Filter

$$\begin{aligned} \frac{d}{dt} X &= -\omega_0 X + k \omega_0 U - b_0 U_0 + b_R U_R \\ b_0 &= \lambda \omega_R - \lambda \omega_0 + k \omega_0 \\ b_R &= \lambda \omega_R \\ y &= X - \lambda U_0 \end{aligned}$$

Auxiliary Filter

$$\frac{d}{dt} U_0 = -\omega_R U_0 + \omega_R U_R. \quad (16)$$

Observe that X is a scalar in (16). Following the lines of the discussion in Section II, we will assume that the flexible filter parameters λ and ω_R are chosen such that $b_0 = 0$, thereby simplifying the main filter equations. With this assumption, and by using the relations in (14), we obtain

Main Filter

$$\begin{aligned} C \frac{d}{dt} V &= -I_0 + k I_0 \frac{U}{X} + \lambda \frac{C}{C_R} I_R \frac{U_R}{X} \\ y &= X - \lambda U_0 \end{aligned}$$

Auxiliary Filter

$$C_R \frac{d}{dt} V_B = -I_R + I_R \frac{U_R}{U_0} \quad (17)$$

where the currents I_0 and I_R are related to the original system parameters as follows:

$$I_0 = C V_t \omega_0; \quad I_R = C_R V_t \omega_R = \frac{\omega_R C_R}{\omega_0 C} I_0. \quad (18)$$

The differential equations in (17) represent nodal equations in the final circuit implementation, where currents in grounded capacitors are equal to a combination of a current source and one or more “translinear” terms—that is, terms containing products and ratios of currents.

In order to finish the design, we offer the following approach to realizing the “translinear” terms in (17). In each case, these terms are the product of two currents divided by a third. Such a current is easily achieved by a transistor connected in a suitable translinear loop [19] as shown with the help of Fig. 1, where the needed extra biasing circuitry is not shown. We assume for the present that $m = 1$. To begin the explanation, we also assume that the voltages labeled V_{dc} are equal to zero, and that the scenario of the figure regards system variables X_j and X_k , which are arbitrary and could represent, for example, U and X , respectively, in the present case. In consideration of the mappings in (13), these system variables are currents exponentially related to respective voltages V_j and V_k . In a physical circuit transistors carrying the currents X_j and X_k may or may not actually exist. Nevertheless, the mappings suggest that if transistors were connected to the voltages V_j and V_k , as indicated by the dotted lines, then they would carry the respective currents X_j and X_k . As a result, we may imagine the transistors connected via the dotted lines to be a part of the circuit even if they are not physically present. We assume that base currents are negligible so that such a connection would cause no loading in any case. Still assuming the voltage V_{dc} to be zero, we have the translinear loop consisting of the three transistors and the diode. It is a simple matter to show [19] that the emitter current I_E is exactly given by $I_{bias} X_j / X_k$. Hence, the translinear terms in (17) are easily realized via a diode level shift—that is, by a diode carrying a current of I_{bias} amps—and an emitter follower bridging the voltage nodes related to the respective system variables as shown in the figure, sourcing the translinear term current into the node labeled V_k .

Notice if we now allow the battery voltages, labeled V_{dc} , to be nonzero, the current I_E is unchanged regardless of the value of V_{dc} . This says that if all of the voltages related to the system variables in (13) were augmented by the same constant V_{dc} , the translinear terms in the equations in (17) would be unchanged. Assuming V_{dc} is constant, none of the derivative relations in (14) would be affected either; therefore, (17) is valid regardless of the possibility of across the board level shifting. This fact is useful in allowing greater flexibility in the design of actual log domain filters and can be easily incorporated into the mappings of (13) without changing the log domain filter equations of (17). Specifically, (13) becomes

$$\begin{aligned} U &= I_s e^{(V_0 - V_{dc})/V_t} \\ X_i &= I_s e^{(V_i - V_{dc})/V_t}, \quad \forall 1 \leq i \leq N \\ U_0 &= I_s e^{(V_B - V_{dc})/V_t} \\ U_R &= I_s e^{(V_R - V_{dc})/V_t}, \end{aligned} \quad (19)$$

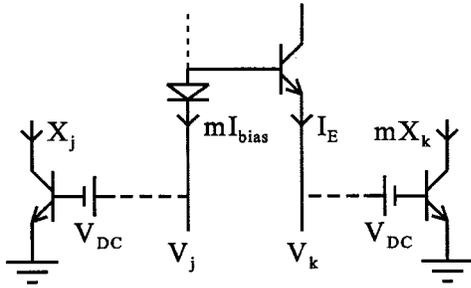


Fig. 1. Simplified loop showing the implementation of a positive translinear term.

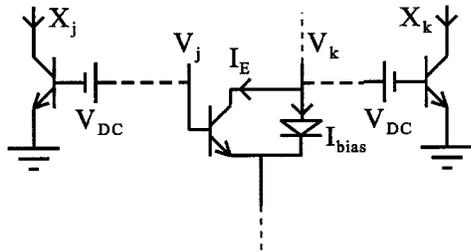


Fig. 2. Simplified loop showing the implementation of a negative translinear term.

Nonzero values of V_{dc} will be used in the second-order filter case discussed below.

The above discussion related to Fig. 1 covers only the case where the translinear term to be implemented is positive. When there is a negative sign, alternate circuitry must be employed as discussed elsewhere [2]. Fig. 2 shows a way to realize a negative translinear term. In this case, there is again a translinear loop formed by the three transistors and the diode, and the current, labeled I_E , being drawn from (yielding the negative sign) the node labeled V_k is again given by $I_{bias}X_j/X_k$, regardless of the value of V_{dc} . It should be noted that the actual circuitry, not shown, needed to bias the diode in this case is more complicated in practice than that needed to bias the diode in the case of Fig. 1. Nevertheless, it is straightforward, and ways to do this have already been proposed in the literature [20]. Thus, we have a basic synthesis procedure for the implementation of any of the translinear terms which appear in (17).

Because we wish to develop methods for synthesizing economical designs, a slight generalization of the above ideas is worth considering. Since each translinear term requires a level shifter, indicated by the diode in Figs. 1 and 2, it is advantageous to use the same level shifter more than once where possible in the final design. This can be done in cases where two translinear terms share the same type of product terms, having only different denominators. While there are many possibilities for this in general systems, we wish to consider a certain case relevant to the companding filters of this paper. Specifically, the auxiliary filter contains a translinear term which shares product terms often occurring in one or more of the translinear terms of the main filter—that is, a bias current times the auxiliary filter input U_R . The difference between these comparative terms in general is that the bias currents are not the same, which necessitates the use of a different level shifter for each. It is possible, however, to scale the bias current—for example, by a factor

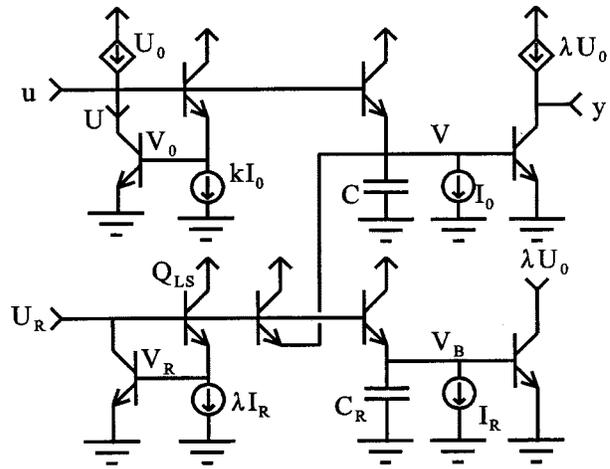


Fig. 3. Log domain implementation of a first-order companding filter.

m —in a translinear term, without altering the value of the term, by also scaling the denominator quantity in the same way—that is, $I_{bias}X_j/X_k = (mI_{bias})X_j/(mX_k)$. This suggests that we may rewrite the auxiliary filter equation of (17) as shown in (20) without changing its basic operation. Specifically,

Auxiliary Filter

$$C_R \frac{d}{dt} V_B = -I_R + mI_R \frac{U_R}{mU_0}. \quad (20)$$

Fig. 1 describes the situation in a way specific to (20) if we let m be an arbitrary positive constant. The net bias current mI_{bias} , flowing through the level shifting diode, can now be varied through the parameter m to be compatible with the corresponding translinear term in the main filter—that is, the translinear term involving the product of I_R and U_R . Because the right-hand side of the differential equation for the auxiliary filter is unchanged, the filtering operation on U_R is unchanged—that is, the cutoff frequency of the auxiliary filter is still ω_R . Observe, however, in Fig. 1 that the “output current” from the rightmost transistor is now mU_0 , instead of U_0 . As long as this is taken into account in the final design, it represents no particular problem. Integer values of m , for example, make this a very simple matter in practice.

With this synthesis procedure for the translinear terms in (17) and (20), a log domain realization of the companding first-order lowpass filter is straightforward. If we choose $C_R = C$ and $m = \lambda$, which provides a simple, yet very practical, special case, we obtain the composite filter equations

Main Filter

$$C \frac{d}{dt} V = -I_0 + kI_0 \frac{U}{X} + \lambda I_R \frac{U_R}{X}$$

$$y = X - \lambda U_0$$

Auxiliary Filter

$$C \frac{d}{dt} V_B = -I_R + \lambda I_R \frac{U_R}{\lambda U_0}. \quad (21)$$

Fig. 3 shows the realization of the filter equations in (21), where the level shifter comprising transistor Q_{LS} has been used to implement the corresponding translinear terms in the main and auxiliary filters. In fact, each of the translinear loops used to implement the various translinear terms is easily identified by studying the schematic. Appropriate current mirrors, not shown

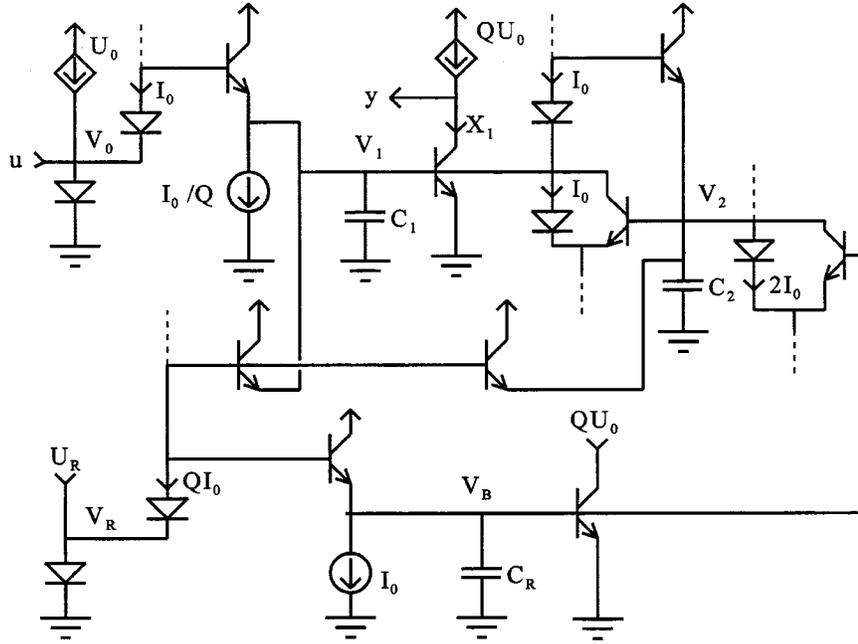


Fig. 4. Simplified realization of the bandpass filter of (28).

in Fig. 3, are needed to reflect the output current λU_0 from the auxiliary filter as the current sources, labeled U_0 and λU_0 , in the main filter. Comparing the analysis here to that in [14], we note that the choices $k = 1.5$, $\lambda = 2$, $C = 10$ pF, and $\omega_R = 0.25\omega_0$ yield the filter proposed there. This first-order case shows how the flexibility in the design parameters may be exploited to simplify the state equation for the main filter, how the relations in (13) yield log domain nodal equations, and how translinear ideas may be used to realize the final filter.

C. The General Case

Having considered a simple example, we are now in position to state the general log domain filter equations associated with the class of syllabic-companding filters we propose. The substitution of the relations in (13) or (19) into (10) and (11) yields

Main Filter

$$\begin{aligned} C_i \frac{d}{dt} V_i &= \sum_{j=1}^N C_i V_t A_{ij} \frac{X_j}{X_i} + C_i V_t \bar{b}_i \frac{U}{X_i} \\ &\quad - C_i V_t \bar{b}_{0i} \frac{U_0}{X_i} + C_i V_t \bar{b}_{Ri} \frac{U_R}{X_i} \\ &= \sum_{j=1}^N I_{ij} \frac{X_j}{X_i} + I_i \frac{U}{X_i} - I_{0i} \frac{U_0}{X_i} + I_{Ri} \frac{U_R}{X_i}, \\ &\quad \forall 1 \leq i \leq N \end{aligned}$$

$$y = \bar{p}^T X + dU - (\bar{p}^T \bar{\lambda} - d)U_0$$

Auxiliary Filter

$$C_R \frac{d}{dt} V_B = -I_R + m I_R \frac{U_R}{m U_0} \quad (22)$$

where subscripts i and j denote individual elements of the respective vectors or matrix in (11). The various currents in (22) are given by

$$\begin{aligned} I_{ij} &= C_i V_t A_{ij} \\ I_i &= C_i V_t \bar{b}_i \\ I_{0i} &= C_i V_t \bar{b}_{0i} \\ I_{Ri} &= C_i V_t \bar{b}_{Ri} \\ I_R &= C_R V_t \omega_R. \end{aligned} \quad (23)$$

Note the translinear terms on the right-hand side of the differential equations as before. Also note that the scaling idea, using the scale factor m , has been incorporated into the writing of the auxiliary filter equation in (22). As in the first-order case above, this can be exploited to reduce complexity in the final design, at the expense of producing a scaled output mU_0 from the auxiliary filter. In practice, it would be prudent for the translated state equations in (11) to have been simplified by the use of the flexible parameters so that only a minimum number of terms appear on the right-hand side of the main filter differential equations.

Synthesis of the general filter equations in (22) is analogous to that for the first-order case. To demonstrate this, we will now consider a second-order example. Suppose we wish to synthesize a bandpass filter characterized by the following dynamical equations:

$$\frac{d}{dt} \bar{x} = \begin{vmatrix} \dot{x}_1 \\ \dot{x}_2 \end{vmatrix} = \begin{bmatrix} -\omega_0/Q & -\omega_0 \\ \omega_0 & 0 \end{bmatrix} \begin{vmatrix} x_1 \\ x_2 \end{vmatrix} + \begin{vmatrix} \omega_0 \\ 0 \end{vmatrix} u \quad (24a)$$

$$y = \begin{vmatrix} 1 & 0 \end{vmatrix} \begin{vmatrix} x_1 \\ x_2 \end{vmatrix} \quad (24b)$$

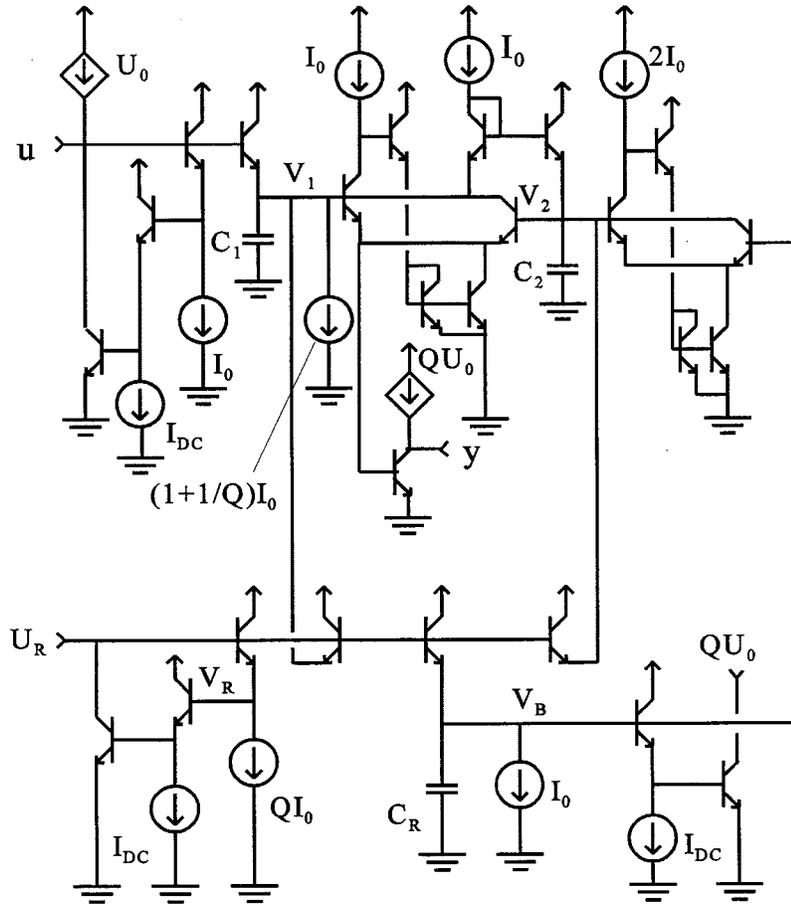


Fig. 5. Final circuit realization of the bandpass filter of (28) based on the simplified diagram of Fig. 5 ($I_{dc} = I_0$).

where a dot over a variable indicates time differentiation. The bandpass transfer function $H(s)$ corresponding to (24) is given by

$$H(s) = \frac{Y(s)}{U(s)} = \frac{s\omega_0}{s^2 + s\omega_0/Q + \omega_0^2}. \quad (25)$$

It is a simple matter to show that the maximum gains from the input to the individual state variables are given approximately by $\lambda_1 = \lambda_2 = Q$. Thus, let us define $\bar{\lambda} = (Q, Q)^T$ which, using (12) and (24), results in the following:

$$\begin{aligned} \bar{b}_0 &= Q \left| \frac{\omega_R - \omega_0}{\omega_R + \omega_0} \right| \\ \bar{b}_R &= Q \left| \frac{\omega_R}{\omega_R} \right|. \end{aligned} \quad (26)$$

Clearly, by choosing $\omega_R = \omega_0$ we can easily eliminate one of the entries of \bar{b}_0 , thereby simplifying the final filter. Adopting this assumption and writing (22) for the special case of the bandpass filter of (24) yields

Main Filter

$$\begin{aligned} C_1 \frac{d}{dt} V_1 &= -\frac{I_0}{Q} - I_0 \frac{X_2}{X_1} + I_0 \frac{U}{X_1} + QI_0 \frac{U_R}{X_1} \\ C_2 \frac{d}{dt} V_2 &= I_0 \frac{X_1}{X_2} - 2I_0 \frac{QU_0}{X_2} + QI_0 \frac{U_R}{X_2} \\ y &= X_1 - QU_0 \end{aligned}$$

Auxiliary Filter

$$C_R \frac{d}{dt} V_B = -I_0 + QI_0 \frac{U_R}{QU_0}. \quad (27)$$

These equations form the basis for a companding log domain filter realization of the bandpass filter. Each of the differential equations represents a nodal equation which can be realized using the ideas above regarding the translinear terms. In order to make this realization economical, we have set m equal to Q , which allows the level shifter used to implement the translinear term in the auxiliary filter to also be used to implement both of the corresponding translinear terms—that is, the terms involving the product $I_0 U_R$ —in the main filter. To achieve this, notice that we have grouped terms involving this product carefully.

Using the notation of Figs. 1 and 2, we obtain the filter realization of Fig. 4 directly from (27). Fig. 4 is, of course, not a final circuit, since the level-shifting diodes require additional circuitry to force the indicated currents. Using the circuits proposed elsewhere [20], we may offer a complete circuit realization as shown in Fig. 5. Here an extra diode drop level shift has been added to all of the system-variable node voltages to allow for the biasing circuitry associated with the negative transconductors. These diode drops are completely analogous to the batteries labeled V_{dc} in Figs. 1 and 2 and are set up by the current sources labeled I_{dc} in Fig. 5. However, because the diode drop in the negative transconductor is used effectively as V_{dc} in developing the bias for the output transistor for the main filter, it

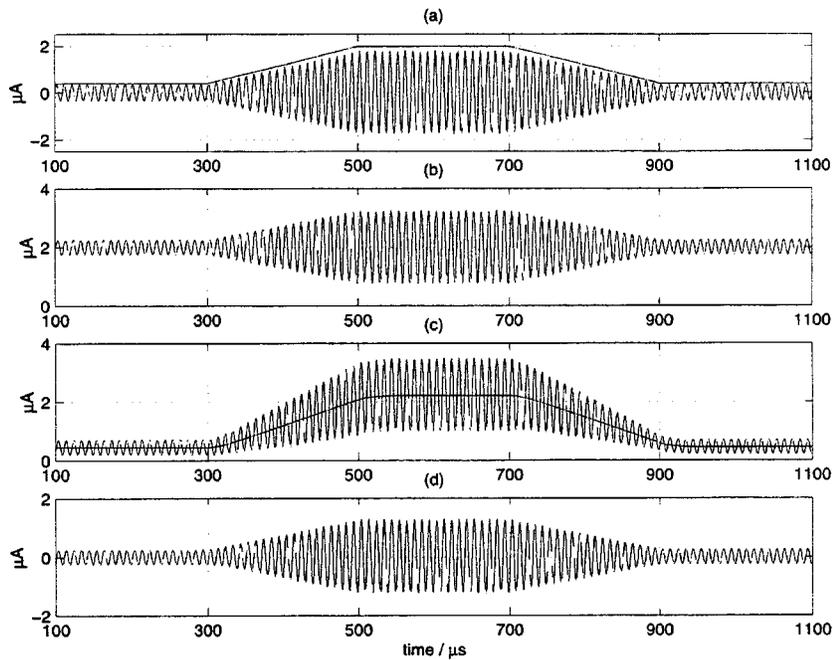


Fig. 6. (a) Input with a changing envelope and corresponding U_R . (b) Output transistor current in the filter with a constant bias. (c) Output transistor current in the filter with dynamic bias and the bias current λU_0 . (d) Output of the filter.

is necessary to set $I_{dc} = I_0$ in Fig. 5. A description of some of the subcircuits used in Fig. 5 is given in [20].

IV. SIMULATION RESULTS AND DISCUSSION

Having given the theory and many details of the design of syllabic-companding filters we are proposing, we wish to finish our presentation with simulation results that show the performance of these filters. The circuit of Fig. 3 is used for this purpose. Because the main reason for the use of these filters, as opposed to noncompanding filters, is to improve the SNR, we focus only on this aspect of performance in this paper. The simulation of noise in companding filters is not a trivial matter, however, as has been pointed out in the literature [21]–[23]. Therefore, a large-signal approach—that is, based on transient analysis—was used to investigate the noise properties of the first-order dynamically biased log domain filter. Substantially ideal transistor models were used for each device throughout the circuit, and transistor current sources and mirrors were used for the indicated biasing. The effects of transistor nonidealities have not been investigated in this theoretical study, which is devoted primarily to the principles of the proposed ideas. The collector shot noise of each transistor was modeled using a gaussian distributed current source from its collector to its emitter. The rms value of this source was made proportional to the square root of the instantaneous collector current. Noiseless and noisy versions of the circuit of Fig. 3 were simulated, using uncorrelated noise sources for each transistor in the noisy case. The output noise was then found by subtracting the two results. The spectral density of the noise was determined by averaging the square of the magnitude of the Fourier transform of the noise taken over 25 runs. Numerically integrating this power spectral density over frequency yields the mean squared value of output noise.

The first job in specifying the dynamically biased filter of Fig. 3 is that of parameter selection. We chose $C = C_R = 100$ pF and the dc gain of the filter—that is, k in (15) and (21)—equal to 1. In order to establish a cutoff frequency of 100 kHz for the main filter, I_0 was set equal to $1.61 \mu\text{A}$. In an attempt to choose an auxiliary filter cutoff frequency that might make it useful in the envelope filtering operation, the cutoff frequency $\omega_R/2\pi$ was set to 10 kHz, yielding $I_R = 0.1I_0 = 0.16 \mu\text{A}$. In order to maintain the constraint that $b_0 = 0$ in (16), λ was set equal to 1.11. Note that, due to this fractional choice for λ , it would be somewhat inconvenient in practice to implement the necessary current mirror taking the output of the auxiliary filter into the current source, labeled U_0 , at the front end of the main filter. This highlights the difficulty of using the auxiliary filter for envelope filtering if small values of capacitance are to be maintained.

Fig. 6 illustrates the principal features of a dynamically biased filter versus a filter with constant bias. Identical inputs [Fig. 6(a)] with varying envelopes were fed into the dynamically biased filter specified above, and a copy of this filter with a constant bias U_0 , both of which have the same transfer function. The *total* current in the output transistor of the filter with a constant bias is shown in Fig. 6(b) and that in the output transistor of a filter with dynamic biasing is shown in Fig. 6(c). Fig. 6(d) shows the output, which is the same for the two filters by design. For small input amplitudes, the reduction in gain from the internal points to the output in a dynamically biased filter as well as the reduced noise from the output transistor itself can be imagined by comparing Fig. 6(b) and (c). This is because the transconductance of the output transistor is proportional to the output current, and this transconductance plays a major role in determining the noise gain from internal nodes to the output.

For the transient noise simulation of the companding filter, the input of the main filter u was a sinusoidal current at the

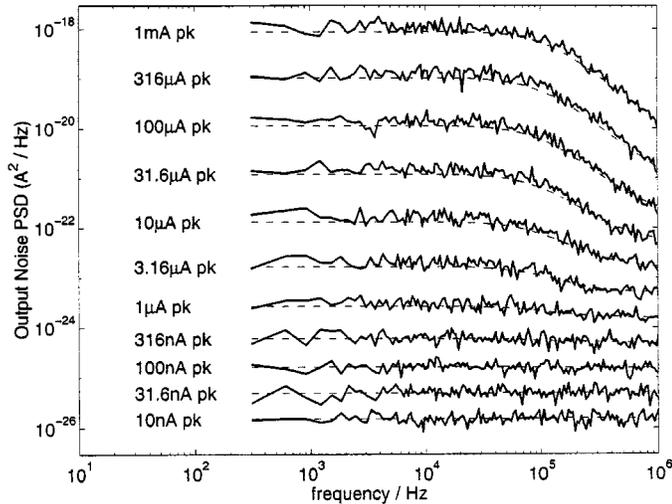


Fig. 7. Output noise PSD of the dynamically biased filter with input amplitude as a parameter: solid line—from transient analysis, dashed line—from ac analysis.

cutoff frequency of 100 kHz, having a constant amplitude. The auxiliary filter input U_R was a constant current with a value of 1.1 times the amplitude of u , mimicking the output of a level detector with a long averaging time compared to the signal frequency. This ensured a 10% safety margin in the collector current swings, which would maintain good bandwidth of the various transistors in practice. The noise spectral density was determined as described above, and this process was repeated for various input amplitudes in the range 10 nA–1 mA. The noise power spectral density (PSD) for the various input levels is shown in Fig. 7. The variation of noise with signal (due to a corresponding change in the bias current) can be seen. Also shown in the figure are the results from small signal noise simulation using the standard “ac” analysis in conjunction with the varying input U_R , applied to the auxiliary filter. Due to the class-A operation of the filter, the results are fairly close to those obtained from large-signal simulation. Fig. 8 shows the integrated (from 0 to 2.5 MHz) noise versus input amplitude.

Finally, Fig. 9 shows the SNR as a function of the input amplitude, computed using the results shown in Fig. 8. For large input currents, the SNR nearly saturates at 60 dB, while for small input currents, the SNR increases by 0.5 dB for every 1 dB increase in the input. This is a departure from conventional class-A linear filters whose SNR increases by 1 dB for every 1 dB increase in the input, until clipping occurs. However, in a conventional class-A filter, the SNR will always be worse than that observed here. For comparison purposes, Fig. 9 also shows the SNR of a fixed bias version of the dynamically biased filter, using a sufficiently large bias to accommodate the largest signals applied. This comparison clearly shows the advantage of companding systems over noncompanding systems. We also note that simulations on a second order dynamically biased filter show the same benefits of companding operation.

The explanation for the noise behavior of the companding filter is illuminating. At small input levels the noise of the dynamically biased filter is dominated by the noise of the logging and exponentiating circuitry, since the core filter circuitry

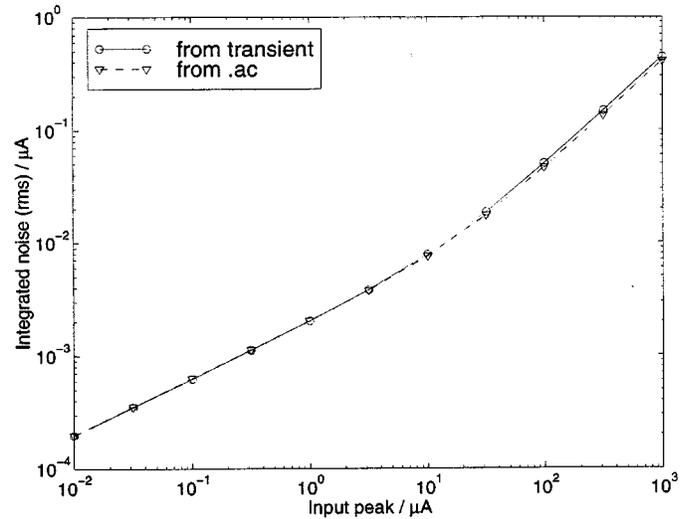


Fig. 8. Integrated output noise versus input amplitude in the dynamically biased filter.

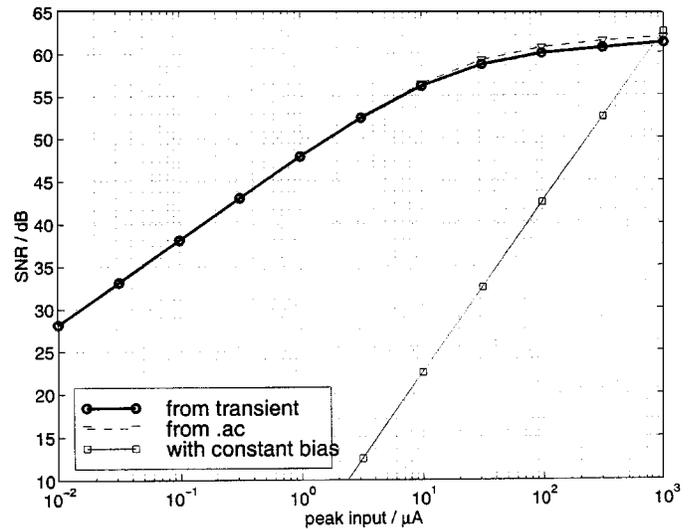


Fig. 9. SNR versus input amplitude in the dynamically biased filter.

is running at considerably higher bias levels. The square root dependence of shot noise on signal level, therefore, explains the 0.5-dB increase in SNR for each 1 dB increase in input level. When the input signal levels are large, the core filter circuitry runs at lower bias than the logging and exponentiating circuitry, making the core circuitry the dominant noise source. Since the dynamically biased filter maintains essentially constant levels in the core filter, the SNR remains constant when the core filter noise dominates at large input levels.

V. CONCLUSION

This paper has presented a complete theory for the design of companding log domain filters, where variations in the bias currents are used to achieve the companding operation. It was shown how translational mappings on the state equations can be used to generate system equations whose internal variables are guaranteed to be positive, which is useful in the design of

any log domain filter. Then a systematic application of the techniques described elsewhere in the literature was used to generate nodal equations for arbitrary companding log domain filters involving translinear terms. It was then shown how these translinear terms could be implemented in an economical way, leading to two generic circuit topologies as special cases. Finally, the noise performance of a first order dynamically biased log domain filter was investigated via a carefully executed large signal simulation technique. The simulation results show clearly the advantage of using our filters in cases where SNR is an important design constraint.

REFERENCES

- [1] Y. Tsividis, "Externally linear, time-invariant systems and their application to companding signal processors," *IEEE Trans. Circuits Syst. II*, vol. 44, pp. 65–85, Feb. 1997.
- [2] D. R. Frey, "Log domain filtering: An approach to current mode filtering," in *Proc. Inst. Elect. Eng.*, vol. 140, Dec. 1993, pp. 406–416.
- [3] —, "Exponential state space filters: A generic current mode design strategy," *IEEE Trans. Circuits Syst. I*, vol. 43, pp. 34–42, Jan. 1996.
- [4] D. Frey and A. Tola, "A state space formulation for externally linear class AB dynamical circuits," *IEEE Trans. Circuits Syst. I*, vol. 46, pp. 306–314, 1999.
- [5] J. Mulder, A. C. van der Woerd, W. A. Serdijn, and A. H. M. van Roermund, "General current mode analysis method for translinear filters," *IEEE Trans., Circuits Syst. I*, vol. 44, pp. 193–197, Mar. 1997.
- [6] E. M. Drakakis, A. J. Payne, and C. Toumazou, "Log-domain filters, translinear circuits and the Bernoulli cell," in *Proc. IEEE Int. Symp. Circuits and Systems (ISCAS97)*, vol. 1, Hong Kong, 1997, pp. 501–504.
- [7] F. Yang, C. Enz, and G. van Ruymbeke, "Design of low-power and low-voltage log-domain filters," *Proc. IEEE Int. Symp. Circuits and Systems (ISCAS96)*, vol. 1, pp. 117–120, 1996.
- [8] D. Perry and G. W. Roberts, "The design of log-domain filters based on the operational simulation of LC ladders," *IEEE Trans. Circuits Syst. II*, vol. 43, pp. 763–774, Nov. 1996.
- [9] R. W. Adams, Filtering in the log domain, New York, NY, May 1979. Preprint #1470, presented at the 63rd AES Conference.
- [10] E. Seevinck, "Companding current-mode integrator: A new circuit principle for continuous-time monolithic filters," *Electron. Lett.*, vol. 26, no. 24, pp. 2046–2047, Nov. 1990.
- [11] Y. Tsividis, V. Gopinathan, and L. Tóth, "Companding in signal processing," *Electron. Lett.*, vol. 26, pp. 1331–1332, Aug. 1990.
- [12] Y. Tsividis and D. Li, "Current-mode filters using syllabic companding," *Proc. IEEE Int. Symp. Circuits and Systems (ISCAS96)*, vol. 1, pp. 121–124, 1996.
- [13] J. Mulder, W. A. Serdijn, A. C. van der Woerd, and A. H. M. van Roermund, "A syllabic companding translinear filter," *Proc. 1997 IEEE Int. Symp. Circuits and Systems (ISCAS97)*, vol. 1, pp. 101–104, 1997.
- [14] D. Frey and Y. Tsividis, "A syllabically companding log domain filter," *Electron. Lett.*, vol. 33, no. 18, pp. 1506–1507, 1997.
- [15] D. Frey, "On instantaneous vs. syllabic companding in log domain filters," *Proc. IEEE Int. Symp. Circuits and Systems (ISCAS99)*, June 1999.
- [16] R. Dolby, "An audio noise reduction system," *J. Audio Eng. Soc.*, vol. 15, no. 4, Oct. 1967.
- [17] D. Frey, "State space synthesis and analysis of log domain filters," *IEEE Trans. Circuits Syst. II*, vol. 45, no. 9, pp. 1205–1211, Sept. 1998.
- [18] W. M. Snelgrove and A. S. Sedra, "Synthesis and analysis of state-space active filters using intermediate transfer functions," *IEEE Trans. Circuits Syst.*, vol. CAS-33, pp. 287–301, Mar. 1986.
- [19] B. Gilbert, *Analogue IC Design: The Current-Mode Approach*, C. Toumazou, F. J. Lidgley, and D. G. Haigh, Eds. London, U.K.: Peregrinus, 1990. Current-mode circuits from a translinear viewpoint: A tutorial.
- [20] D. R. Frey, "Log domain filtering for RF applications," *IEEE J. Solid-State Circuits*, vol. 31, pp. 1468–1475, Oct. 1996.
- [21] M. Punzenberger and C. C. Enz, "A 1.2 V low-power BiCMOS class AB log-domain filter," *IEEE J. Solid-State Circuits*, vol. 32, pp. 1968–1978, Dec. 1997.
- [22] J. Mulder, M. H. L. Kouwenhoven, and A. H. M. van Roermund, "Signal x noise intermodulation in translinear filters," *Electron. Lett.*, vol. 33, pp. 1205–1207, July 1997.
- [23] L. Tóth, Y. P. Tsividis, and N. Krishnapura, "On the analysis of noise and interference in instantaneously companding signal processors," *IEEE Trans. Circuits Syst. II*, vol. 45, pp. 1242–1249, Sept. 1998.



Douglas Frey (S'74–M'76–SM'88) received the B.S., M.S., and Ph.D. degrees from Lehigh University, Bethlehem, PA, in 1973, 1974, and 1977, respectively.

He was with Bell Laboratories, Holmdel, NJ, in 1977 before accepting a teaching position at Lehigh University as Assistant Professor. In 1998, he was promoted to Full Professor of Electrical Engineering and continues to be a faculty member in the Electrical Engineering and Computer Science Department. His main areas of interest are in circuit design and nonlinear circuits and systems theory. He has developed and taught a number of courses in these areas. His research includes work on log domain filters, analog adaptive filters, nonlinear digital fixed and adaptive filters, and chaotic encoding. He is also an engineering consultant, in which capacity he has been involved in discrete, semi-custom, and full-custom integrated circuit design, including consumer electronics. He has been a consultant for Analog Devices and is presently a mixed-signal analog design engineer for Silicon Laboratories. He has nine patents and has presented and published a number of papers in his areas of interest.

Dr. Frey is a member of the Audio Engineering Society and a registered Professional Engineer in Pennsylvania.



Yannis P. Tsividis (S'71–M'74–SM'75–F'86) received the B.S. degree from the University of Minnesota, Minneapolis, and the M.S. and Ph.D. degrees from the University of California, Berkeley, in 1972, 1973, and 1976, respectively.

He is the Bachelor Professor of Electrical Engineering at Columbia University, New York. He has worked for Motorola Semiconductor and AT&T Bell Laboratories and has taught at the University of California at Berkeley, Massachusetts Institute of Technology, and the National Technical University

of Athens, Greece.

Dr. Tsividis is the recipient of the 1984 IEEE Baker Best Paper Award, the 1986 European Solid-State Circuits Conference Best Paper Award, and the 1998 IEEE Circuits and Systems Society Guillemin–Cauer Best Paper Award. He is co-recipient of the 1987 IEEE Circuits and Systems Society Darlington Best Paper Award. He has received the Great Teacher Award at Columbia University.



Giorgos Efthivoulidis was born in Greece in 1970. He received the Diploma and Ph.D. degrees in electrical and computer engineering from the National Technical University of Athens, Athens, Greece, in 1993 and 1998, respectively.

He joined Intracom S.A., a telecommunications company in Greece, for one year. He is currently affiliated with the Massachusetts Institute of Technology, Cambridge, conducting research in analog and hybrid systems.

Nagendra Krishnapura received the B.Tech. degree in electronics and communications engineering from the Indian Institute of Technology, Madras, in 1996. He is currently working toward the Ph.D. degree in electrical engineering at Columbia University, New York.

During the summer of 1997, he worked at Texas Instruments, NJ and during the summers of 1998 and 1999, at Lucent Technologies, Bell Laboratories, NJ. His research interests include analog and RF circuits.