Widely Programmable High-Frequency Continuous-Time Filters in Digital CMOS Technology

Shanthi Pavan, Yannis P. Tsividis, *Fellow, IEEE*, and Krishnaswamy Nagaraj

***Abstract*—We present design considerations for programmable high-frequency continuous-time filters implemented in standard digital CMOS processes. To reduce area, accumulation MOS capacitors are used as integrating elements. The filter design problem is examined from the viewpoint of programmability. To allow frequency scalability without deterioration of noise performance and of the frequency response shape, we employ a technique called "constant-capacitance scaling," which assures that even parasitic capacitances remain invariant when transconductors are switched in and out of the filter. This technique is applied to the design of a programmable fourth-order Butterworth continuous-time filter with a bandwidth programmable from 60 to 350 MHz implemented in a 0.25-μm digital CMOS process. The filter has a dynamic range of 54 dB, dissipates 70 mW from a 3.3-V supply, and occupies an area of 0.15 mm^2.**

***Index Terms*—Accumulation, CMOS, continuous-time filter, digitally programmable, fixed transconductance, scaling, tuning.**

I. Introduction

HIGH-SPEED communication and storage systems need continuous-time filters with bandwidths tunable over a wide range, while keeping the relative shape of the frequency response fixed irrespective of the set bandwidth and maintaining an adequate dynamic range [1]. For high levels of integration and to reduce cost, it is desirable to implement these filters in deep-submicrometer digital CMOS technologies. However, in such processes, high-density poly-poly capacitors are not available. In this paper, we use MOS accumulation structures in place of the low-capacitance density metal–metal capacitors typically used in digital CMOS technologies. These accumulation capacitors will be described in Section II. The design of programmable high-frequency filters presents a number of problems quite different from those encountered in the design of fixed bandwidth filters. These problems are discussed in Section III. To ameliorate these problems, we introduce a technique called "constant-capacitance scaling" in Section IV. This technique enables the design of filters without deterioration of noise performance and frequency response accuracy as they are programmed over a wide range. An efficient CMOS implementation of a constant capacitance scaled transconductor is presented in Section V. The design of a fourth-order, 60–350-MHz Butterworth filter in a 0.25-μm n-well technology is discussed in Section VI. A tuning system, required to keep the bandwidth of the filter constant in the face of process and temperature variations, is discussed in Section VII. Implementation results of the filter prototype are presented in Section VIII, and conclusions are given in Section IX.

Manuscript received June 10, 1999; revised November 18, 1999.

S. Pavan and K. Nagaraj are with Texas Instruments, Warren, NJ 07059 USA (e-mail: shanthi@ti.com).

Y. P. Tsividis is with the Department of Electrical Engineering, Columbia University, New York, NY 10027 USA.

Publisher Item Identifier S 0018-9200(00)02863-8.

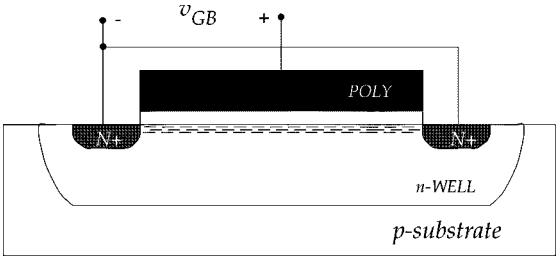

Fig. 1. Cross section of an n-polysilicon n-well capacitor.

II. Capacitors in Digital CMOS Technologies

High-density double-poly capacitors are not available in digital CMOS technologies. Metal–metal structures can be used, but they have the disadvantage of having very low specific capacitance, unless they are realized using special techniques [2]. An alternative is to use the capacitance formed by MOSFET gates [3], [4]. Polysilicon to n-well structures are naturally available in CMOS technologies. The advantages of using the gate-oxide capacitance are well-controlled values and large specific capacitance.

A two-terminal MOS structure can be operated in accumulation or inversion. It can be shown that for typical process parameters, and a given gate-body bias, accumulation capacitors tend to be more linear [5]. The cross section of an n-polysilicon, n-well capacitor is shown in Fig. 1. A physics-based model for accumulation capacitors, which predicts distortion accurately, can be found elsewhere [6]. There it is shown that for a gate-body bias voltage of 1 V and a peak voltage swing of 200 mV, the third-harmonic component of the capacitor voltage for a sinusoidal current drive is about 80 dB below the fundamental; even-order harmonics largely cancel out when two capacitors are used in a balanced fashion (in which case 800-mV peak-to-peak differential signals become possible). For several applications, therefore, nonlinearity of the capacitors is not of concern provided they are biased by about a volt into accumulation.

III. Issues in Programmable Filter Design

Apart from the usual issues associated with high-frequency CMOS filter design [7], the issue of *programmability* brings to the forefront the very important problem of *maintaining* adequate dynamic range and accurate frequency response across the tuning range. Problems in this respect become very acute when the pole frequencies are high (several hundreds of megahertz) *and* when they have to be tuned over a wide range. Although these problems exist in all filter architectures (Gm-C, MOSFET-C, Gm-OTA-C), here we only examine the specific case of Gm-C filters.

A brief discussion of the issues and choices in programmable continuous-time filter design will now be given. The simplest possible circuit that can illustrate a point will be used in each case.

A. Constant C Versus Constant G_m

We consider a first-order, low-pass G_m-C filter, which must be programmed over a wide range of frequencies, the highest of which is $f_{c,\max}$. We assume that an optimized design for the toughest case (the highest cutoff frequency) has been achieved, as in Fig. 2(a). The relation of this frequency to component values, total integrated mean-square output noise, and power dissipation is given next to the figure; it is assumed that the noise current power spectral density is $4kT\eta G_m$, where η is the excess noise factor of the transconductor [8], [9]. Also, as is often the case in optimized transconductor design, the power dissipation is assumed to be proportional to G_m, with a constant of proportionality b, depending on design details. For a given noise spec $\overline{v_{\rm on}^2}$, the value of C is set, which from the top relation in the figure sets G_m and thus P.

There are two fundamental ways of making the cutoff frequency programmable, as shown in Fig. 2(b) and (c); these are termed "constant G_m" and "constant C," respectively. The corresponding cutoff frequency, mean-squared noise, and power dissipation (denoted by primes) are given below each figure, in terms of those at $f_c = f_{c,\max}$. From these results, it follows that as f_c is programmed, the noise, total capacitance area, and power dissipation will vary as shown in Fig. 3. It is clear that the constant-G_m design can result in unnecessarily large capacitance in order to achieve low cutoff frequencies. In low-noise applications, where chip area can be capacitance-limited, this can be a serious problem. The fact that very low noise is achieved by this design for low values of f_c is irrelevant, as far as the spec (indicated on the top plot) is concerned; the total capacitance cannot be decreased, as then the entire noise plot will rise, and then the spec will not be satisfied at high frequencies. In contrast, the constant-C design needs no capacitance increase, and the noise spec is still satisfied everywhere. Finally, since the low f_c values are achieved by decreasing G_m, a lower power dissipation can be achieved at low f_c values, depending on the transconductor design (in the design reported in this paper, we give up the latter advantage in order to achieve parasitics immunity, as will be seen). Due to the above observations, the rest of this paper focuses on the constant-C approach.

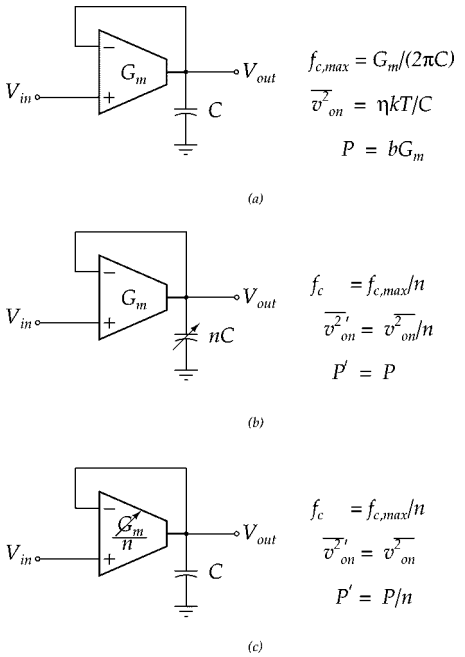

Fig. 2. (a) A first-order Gm-C filter designed to achieve the maximum cutoff frequency in a programmable range. (b) Lowering the cutoff frequency of (a) by increasing the capacitance ("constant G_m" approach). (c) Lowering the cutoff frequency of (a) by decreasing the transconductance ("constant C" approach).

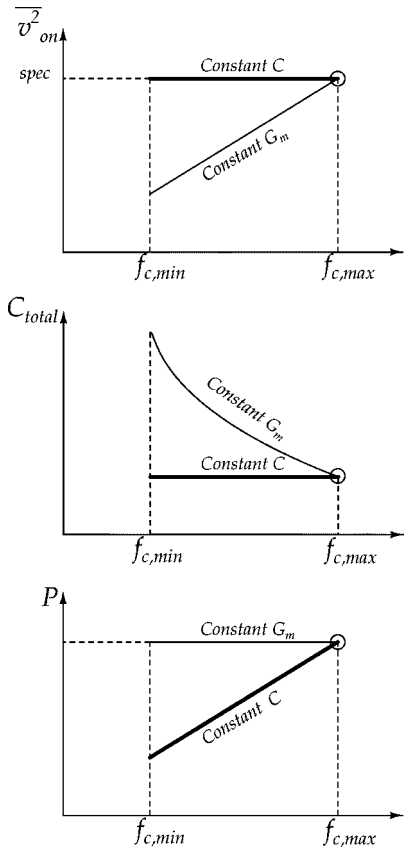

Fig. 3. Variation of mean-squared noise, total capacitance, and power dissipation versus cutoff frequency for programmable filters using the "constant G_m" approach (thin line) and the "constant C" approach (thick line). The circle indicates a design optimized for the maximum f_c setting.

B. Variable G_m Versus Multiple Unit Cells

A number of CMOS linearized transconductors have been reported in the literature. Many of these cannot be optimally designed if their transconductance is to be variable. For example, the transconductors of [10]–[13] exhibit varying excess phase shift and varying excess noise factor as their transconductance is varied; each of these parameters attains its worst case value at opposite extremes of the programmable frequency range [17]. This makes design difficult and results in suboptimum performance.

An alternative, which does not suffer from the above problems, is to compose G_m out of optimized unit transconductor cells and switch these in parallel to achieve the desired programmable transconductance values. (If continuous tuning is needed in between the discrete steps, one can control the G_m values over a small range using well-known continuous tuning techniques. Since this range will be very small, the variability of performance over this range will be small too. This case is not considered below.)

C. Effect of Parasitic Capacitances on Programmability

When switchable transconductor cells are used, a significant problem arises. To illustrate this problem, we consider the case of a second-order filter, as shown in Fig. 4(a). The pole frequency and quality factor of this filter are as follows [14]:

$$f_o = \frac{1}{2\pi}\sqrt{\frac{G_{m1}G_{m2}}{C_aC_b}} \tag{1}$$

and

$$Q = \sqrt{\frac{G_{m1}}{G_{m2}}\frac{C_a}{C_b}}. \tag{2}$$

To make the pole frequency programmable, G_{m1} and G_{m2} are commonly varied by the same factor, and thus (G_{m1}/G_{m2}) remains fixed. However, if switchable unit cells are used to implement the programmable G_{m1} and G_{m2}, then each time such cells are switched in and out, the total value of the parasitic capacitance at each node changes, in a manner that is very difficult to control. This is illustrated in Fig. 4(b), where the arrows indicated elements that are variable (intentionally or unintentionally). C_{i1} and C_{o1} are the input and output parasitic capacitances of transconductors G_{m1}, and C_{i2} and C_{o2} are the corresponding quantities for transconductors G_{m2}. The pole frequency and quality factor of this filter now become

$$f_o = \frac{1}{2\pi}\sqrt{\frac{G_{m1}G_{m2}}{(C_a+2C_{o1}+2C_{i2}+C_{o2})(C_b+C_{i1}+C_{o2})}} \tag{3}$$

and

$$Q = \sqrt{\frac{G_{m1}}{G_{m2}}\frac{(C_a+C_{o1}+2C_{i2}+C_{o2})}{(C_b+C_{i1}+C_{o2})}}. \tag{4}$$

The variation of parasitics causes two problems: the pole frequency will not change in proportion to the stepped G_{m1} and G_{m2}; and the changing parasitics will change Q, resulting in response shape errors. The latter effect is especially severe for higher order filters and/or ones that are supposed to have accurate group delay characteristics. It is possible that a design that meets the specs at the highest pole frequency will not meet them at lower pole frequencies, due to the variability of parasitics as transconductances are switched in and out of the filter. It is thus seen that it is not enough to achieve a working design at the highest pole frequency of interest and then hope that there will not be any problems at lower pole frequency settings.

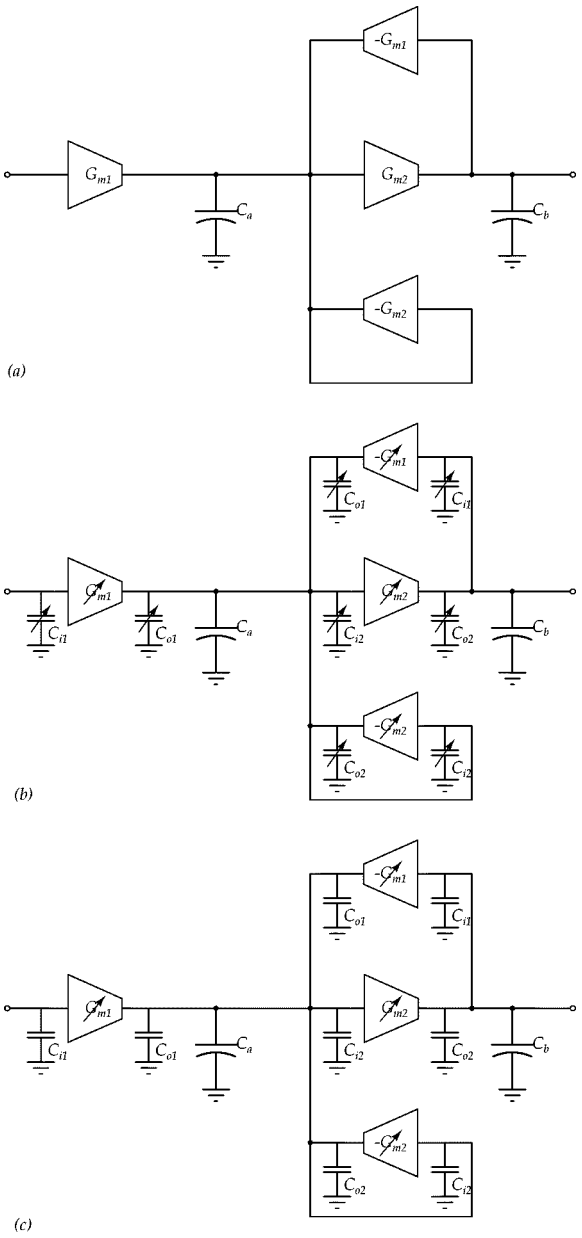

Fig. 4. (a) A second-order Gm-C filter. (b) The filter in (a), made programmable by varying transconductances; the parasitic input and output capacitances also vary, as indicated by arrows. (c) The filter in (a), assuming all parasitics remain fixed while transconductances are varied.

Consider now the situation in Fig. 4(c). Assume that all parasitic capacitances can be kept fixed and only the G_ms are variable. Once the fixed value of parasitics has been incorporated into the design, it is seen from (1) that the pole frequency will now change in proportion to the G_m's and Q will be fixed, as seen from (4). This paper discusses techniques that make possible the reliable implementation of this principle.

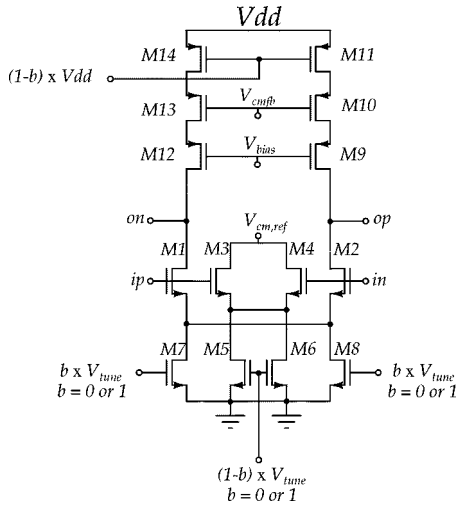

Fig. 5. A unit transconductance element.

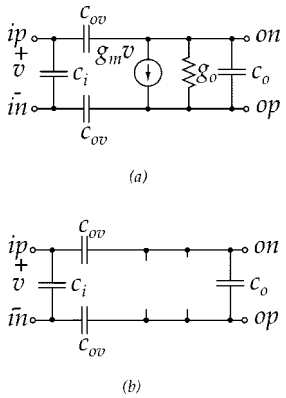

Fig. 6. The small-signal equivalent circuit of the transconductor of Fig. 5 under differential excitation, for the cases (a) $b = 1$ and (b) $b = 0$ (where b denotes the state of the corresponding control signals in Fig. 5). Note that the transconductance and output conductance can be switched off without changing any of the parasitic capacitances.

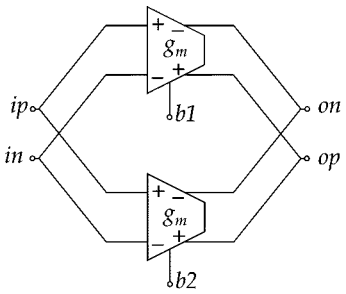

(a)

b1	b2	G_m	G_o	C_i	C_{ov}	C_o	DC Gain
0	1	g_m	g_o	$2c_i$	$2c_{ov}$	$2c_o$	g_m/g_o
1	1	$2g_m$	$2g_o$	$2c_i$	$2c_{ov}$	$2c_o$	$2g_m/2g_o = g_m/g_o$

(b)

Fig. 7. (a) Parallel connection of two unit transconductors. (b) Properties of the composite transconductor in terms of the properties of the individual transconductors.

IV. True Constant-Capacitance Scaling

The principle just discussed can be generalized as follows. Consider any arbitrary linear network of capacitors, conductances, and transconductances, and let the voltage transfer function between any two ports be denoted as $H_{21}(f)$. If *every* conductance and transconductance in the network is multiplied by an arbitrary positive number α while keeping *every* capacitance the same (*including* the parasitic ones), then it is easy to show that the voltage transfer function between the ports is changed to $H_{21}(f/\alpha)$. We call the above modification "constant-capacitance scaling." It corresponds to a stretching of the frequency axis. The properties of scaled linear and nonlinear networks are derived elsewhere [16]. Here, we only summarize the main results.

It can be shown that the integrated output thermal noise of a constant-capacitance scaled network is independent of the scaling factor α [16]. Further, if the capacitors and trans(conductors) are weakly nonlinear, it can be shown [16], [17] that the worst case total harmonic distortion of a constant-capacitance scaled network is also independent of α. Thus the dynamic range is also independent of the scaling factor (assuming that thermal noise is dominant). Since, as already mentioned, the frequency response shape is also independent of the scaling factor, no overdesign is needed. If the original filter is designed optimally in terms of thermal noise and distortion, then any constant-C scaled version of it will also be optimal, irrespective of frequency setting [16]. This property is key to making possible the performance of the chip described in Sections V–VIII.

V. Constant-Capacitance Programmable Transconductors

Fig. 5 shows the schematic of a unit transconductance element. It is built around a differential pair without special linearization. This was found to be the best choice for the combination of distortion and wide-band operation specifications we intended to meet. M1, M2, M3, and M4 are identical devices, with M1–M2 forming the main differential pair of the transconductor and M3–M4 forming a dummy differential pair. M7–M8 and M5–M6 form the current sources for the main and dummy differential pairs, respectively. When M7–M8 are supposed to act as current sources $(b = 1)$, their gates are connected to a voltage $V_{\rm tune}$, which is generated by a bandwidth tuning circuit (see Section VII). When these current sources are to be shut off $(b = 0)$, their gates are connected to ground and M5–M6 are connected to $V_{\rm tune}$. This ensures that a differential pair (either the main one, M1–M2, or the dummy one, M3–M4) is always connected at the input, and thus the input capacitance is maintained constant independent of b. A simple pass-transistor network is used to accomplish the switching of the gates of the current sources between $V_{\rm tune}$ and ground.

$V_{\rm cm,ref}$ is a dc voltage equal to the common-mode dc level of the transconductor outputs (Section VI-A). M11 and M14 operate in the triode region and are used to switch the load current sources on and off. M9 and M12 are cascode devices that enhance the output impedance of the PMOS current sources. $V_{\rm cmfb}$ is a voltage derived from a common-mode feedback circuit (see Section VI-A) that assures that the common-mode level of *on* and *op* is equal to $V_{\rm cm,ref}$. In our implementation, this value

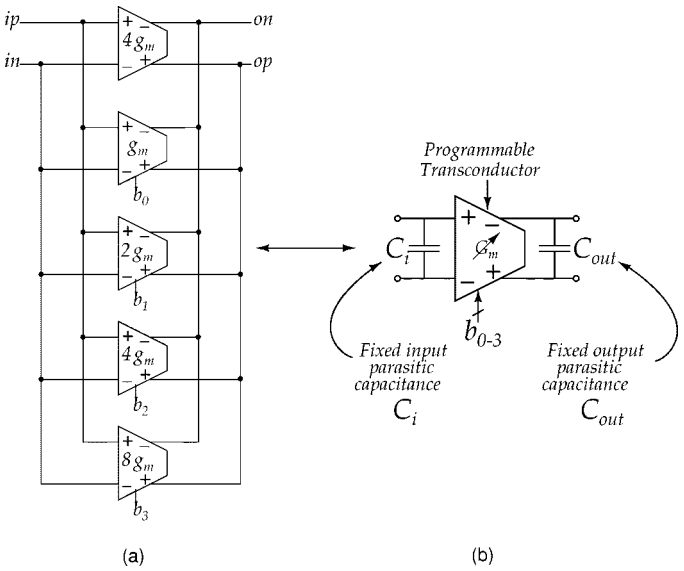

Fig. 8. A 4-bit programmable constant capacitance transconductor used in the chip implementation.

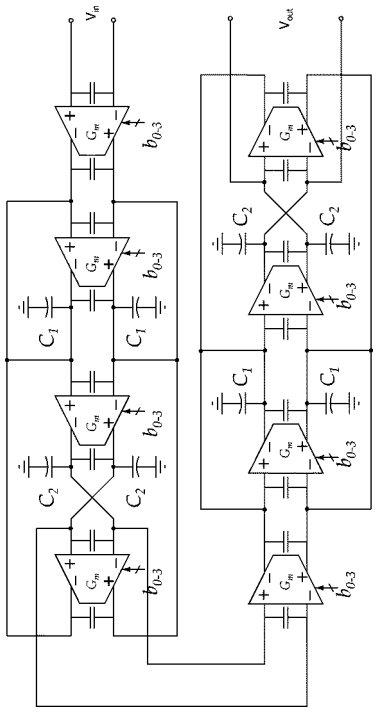

Fig. 9. A fourth-order Gm-C filter using the transconductors of Fig. 8. The unmarked capacitors are input and output parasitics, which remain invariant as the transconductor values are changed.

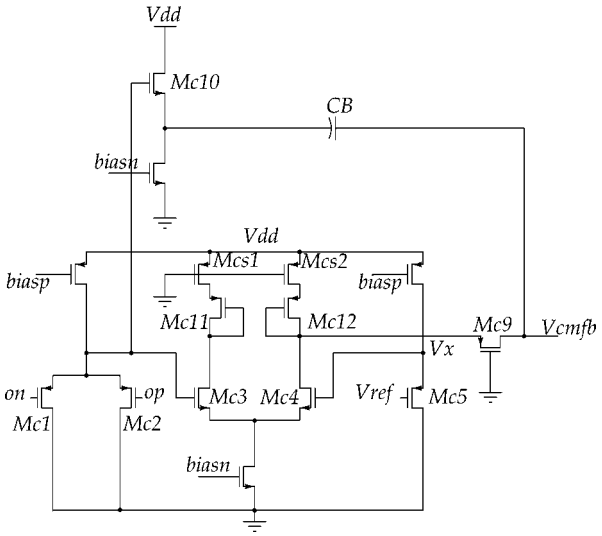

Fig. 10. Common-mode feedback circuit.

was chosen to be 1.2 V. The integrating capacitors are grounded accumulation MOS devices (Fig. 1). Since $V_{\mathrm{cm,ref}} = 1.2$ V, the MOS capacitors have a sufficient bias voltage across them for low-distortion operation (Section II).

The small-signal equivalent circuits for differential excitation for the cases $b = 1$ and $b = 0$ are shown in Fig. 6(a) and (b), respectively. Part (a) is for the case in which M1–M2 operate in strong inversion while M3–M4 are off. Part (b) is for the case in which M3–M4 operate in strong inversion while M1–M2 are off. In *both* cases, *all* nodal capacitances in the network remain identical. Notice that when the transconductor M1–M2 is turned off, *every* transconductance and conductance in the network is removed, while *every* parasitic capacitance remains the same. The transconductor M1–M2 does not contribute noise when it is turned off. The total power dissipation is constant, independent of the setting of b. This is the price one has to pay in order to keep all capacitances constant. (Power could have been saved when $b = 0$ by switching in a passive dummy device instead of M3–M4, but this would have compromised on the robustness of the technique.)

Now consider the parallel connection of two identical unit transconductors, as shown in Fig. 7. The table in the figure lists the various parameters of the composite transconductor (denoted by upper case letters) as a function of those of the individual transconductors (denoted by lower case letters), for the cases of one transconductor's being on (upper entry) and both transconductors' being on (lower entry). These results can be deduced by using the models in Fig. 6. Again, note that *every* node capacitance remains the same while *every* transconductance/conductance scales appropriately. The results can be extended to many unit transconductors connected in parallel. Fig. 8(a) shows the composite transconductors used in the chip implementation. When the digital word corresponding to $b_0b_1b_2b_3$ is set to 0000, only four unit elements are switched on. When the word is 1111, all 19 elements are on. The coarse programming range is therefore $19/4 = 4.75$. As transconductance is stepped over this range, *all* parasitic capacitances remain fixed and are equal to $19C_i$ for the input and $19C_o$ for the output capacitance. Fig. 8(b) shows the symbol for the composite transconductor of Fig. 8(a). For simplicity, the small input–output capacitance due to c_{ov} in Fig. 6 is not shown; note, however, that even this capacitance remains invariant as the transconductance value is changed. The scheme accomplishes, therefore, true constant-C scaling, in the sense discussed in Section IV. In applications where the programming resolution is not sufficient, one can either use more unit cells or tune the transconductances continuously in between steps (i.e., in a small range) using well-established techniques.

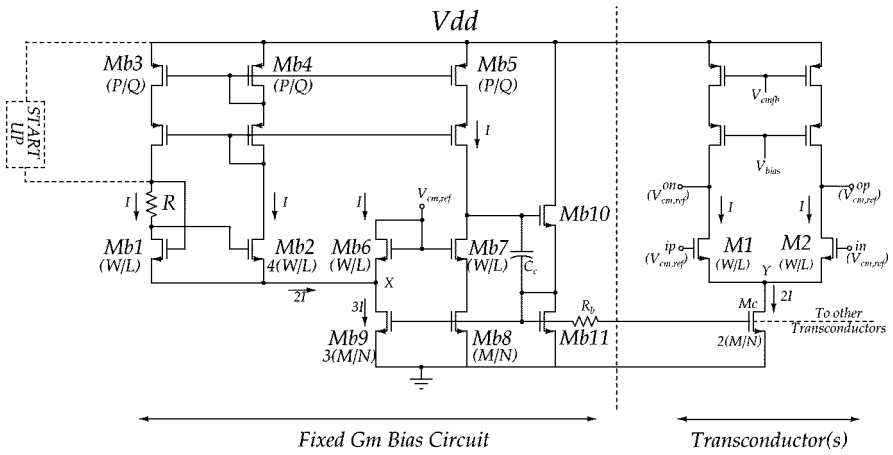

Fig. 11. Complete "fixed transconductance bias" circuit.

VI. Integrated Circuit Implementation

In this section, we describe the design of a fourth-order Butterworth filter with a maximum bandwidth of 350 MHz, designed in a 0.25-μm n-well digital CMOS process. The filter consists of two second-order filter sections and is shown in Fig. 9 [15]. All transconductors are made equal-valued for matching reasons. Each of the transconductors shown is implemented as in Fig. 8. The unmarked capacitors are input and output parasitics, which remain invariant as the transconductor values are programmed. Since the integrating capacitors and the parasitic input capacitance of the transconductors are both formed from gate oxide, a good degree of matching of capacitor ratios can be achieved.

A. Common-Mode Feedback Circuit

Fig. 10 shows the common-mode feedback (CMFB) circuit. The differential pair, consisting of Mc1–Mc2 (whose gates are connected to the outputs of the filter transconductor), forms the common-mode detector. An advantage of this kind of detector is its capacitive input impedance, which can be incorporated into the filter integrating capacitance. The detected common-mode voltage is compared with a common-mode reference ($V_{\rm ref}$), which goes through a PMOS source follower (to mimic the common-mode detector). Mc3, Mc4, Mc11, and Mc12 form the servo amplifier. Mc10 is a low-output impedance source follower. Mc9 is a p-channel device operating in the triode region. Along with CB, it acts as a crossover network between the servo amplifier and the source follower. This arrangement allows a high-dc loop gain at low frequencies while providing a feed-forward path at high frequencies. CB is a large capacitor realized using an accumulation MOS structure. Mcs1 and Mcs2 are replicas of the corresponding devices used as the top switches in the transconductor (M11 and M14 in Fig. 5, which operate in the triode region when on). The output Vcmfb of the circuit in Fig. 10 is connected to the gates of M13 and M10 in Fig. 5.

VII. Tuning System

To keep the bandwidth of the filter invariant with temperature for a particular frequency setting, the tail current of every differential pair is varied in a manner such that its transconductance tracks a precise external resistor. Rather than using an explicit dc loop to do this, we exploit the square law characteristics of the MOSFET. The bias circuit used to accomplish this is shown on the left in Fig. 11. The current generated in this circuit is used to derive the tail currents for all the differential pairs used as transconductors in the filter. A startup circuit should be included as shown (but was not implemented in our test chip).

One transconductor is shown on the right in Fig. 11. $V_{\rm cm,ref}$ is the common-mode reference voltage for the chip. In the analysis to follow, mobility reduction due to gate field and channel-length modulation are neglected; as will be seen, simulation and measurements confirm that these effects are not troublesome.

The current mirror formed by Mb3 and Mb4 forces identical currents through Mb1 and Mb2. Assuming that Mb1 and Mb2 are operating in strong inversion and saturation, we can write

$$I_{Mb1} = k'\left(\frac{W}{L}\right)(V_{{\rm GS},b1} - V_T)^2 \tag{5}$$

$$I_{Mb2} = k'\left(\frac{4W}{L}\right)(V_{{\rm GS},b2} - V_T)^2 \tag{6}$$

$$I_{Mb1} = I_{Mb2} = I \tag{7}$$

$$V_{{\rm GS},b1} - V_{{\rm GS},b2} = IR \tag{8}$$

where k' is a process-dependent constant. From (5)–(7), we get

$$V_{{\rm GS},b1} - V_T = 2(V_{{\rm GS},b2} - V_T). \tag{9}$$

From (8) and (9), we can write

$$\frac{2I}{V_{{\rm GS},b1} - V_T} = \frac{1}{R}. \tag{10}$$

But $g_m = (dI_{\rm DS}/dV_{\rm GS}) = (2I_{\rm DS}/(V_{\rm GS} - V_T))$. Thus

$$g_{m,Mb1} = \frac{1}{R}. \tag{11}$$

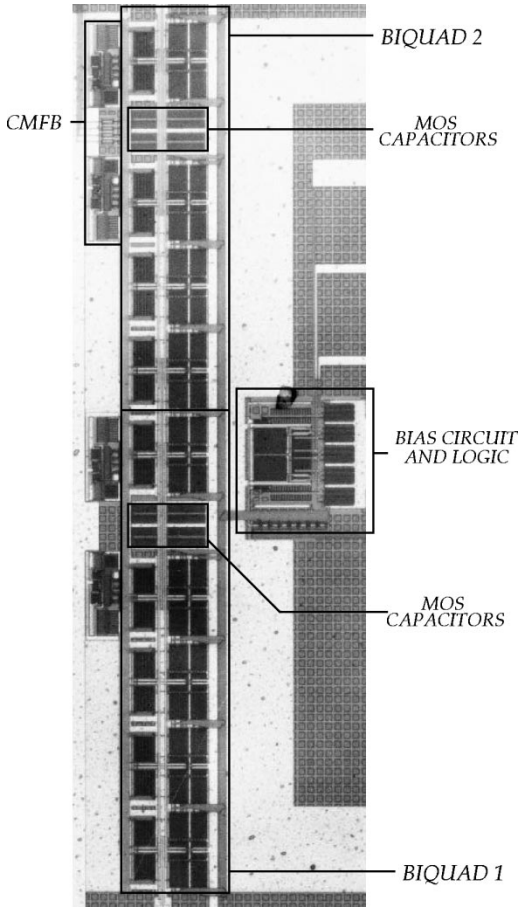

Fig. 12. Chip microphotograph.

Hence, the circuit stabilizes to a state where the current is such that the transconductance of Mb1 is maintained at $1/R$, irrespective of V_T, μ, or temperature [18]. Simulations show that the g_m of Mb1 is maintained within $\pm 1.5\%$ over a 100 °C range. Recall that our aim is to generate a tail current for the differential pairs in the filter, represented in Fig. 11 by M1 and M2, such that their transconductance remains constant with temperature.

Transistors Mb5 through Mb11 ensure that M1–M2 and Mb1 operate exactly in the same manner. To see that this is indeed so, consider the following. The current I flowing through Mb5 is equal to that in Mb3. The negative feedback loop formed by Mb10, Mb11, Mb8, and Mb7 forces the current through Mb8 to be I. Since Mb9 is three times as wide as Mb8, $3I$ flows through it. Of this, $2I$ is supplied by Mb1 and Mb2, leaving I to flow through Mb6. The gate of Mb6 is held at a voltage $V_{\mathrm{cm,ref}}$, the common-mode reference for the filter. Consider now the transistor M1. In the quiescent state, its gate is at a potential $V_{\mathrm{cm,ref}}$ (This condition is forced by the common-mode feedback circuitry for every transconductor output in the filter, and one such output is assumed to be driving the gates of M1 and M2). Since Mc is twice as wide as Mb8, $2I$ flows through it. Hence, the current through M1 is I. Therefore, Mb6 and M1 are equal in size, have the same gate voltage and the same drain voltage, and have the same current flowing through them. This means that the voltage at node Y in the transconductor is exactly equal to that of node X in the bias circuit, thus ensuring that the source-substrate bias is the same for Mb6 and M1. So, Mb1, M1 and M2 operate in exactly the same fashion (neglecting output conductances of the devices). Since the g_m of Mb1 is maintained at $1/R$ (to within $\pm 1.5\%$) over a 100 °C variation in temperature, so is the transconductance of M1–M2.

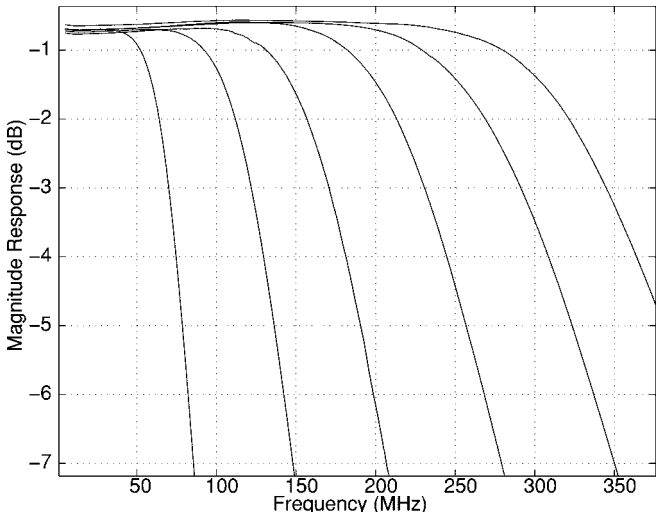

Fig. 13. Passband detail for various frequency settings.

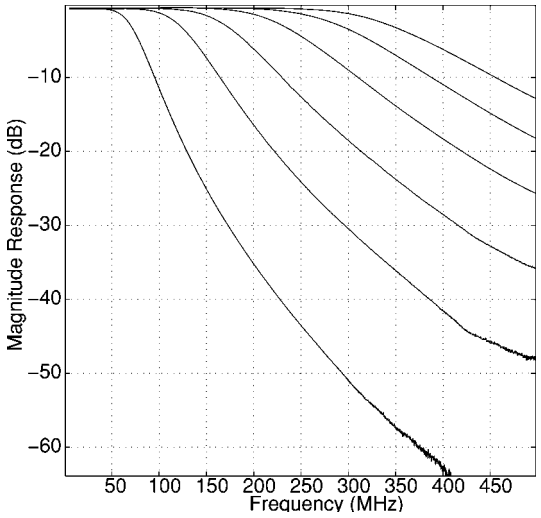

Fig. 14. Frequency response for various frequency settings.

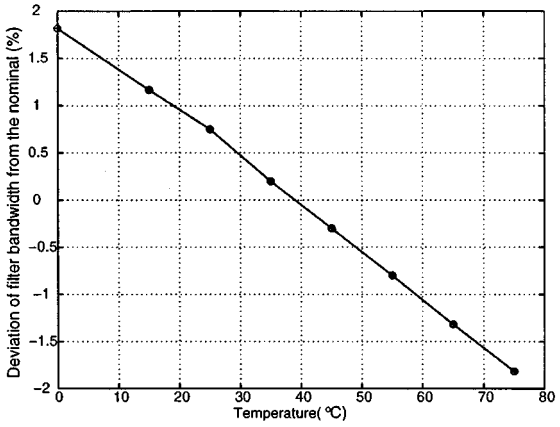

Fig. 15. Percentage variation of filter bandwidth over temperature.

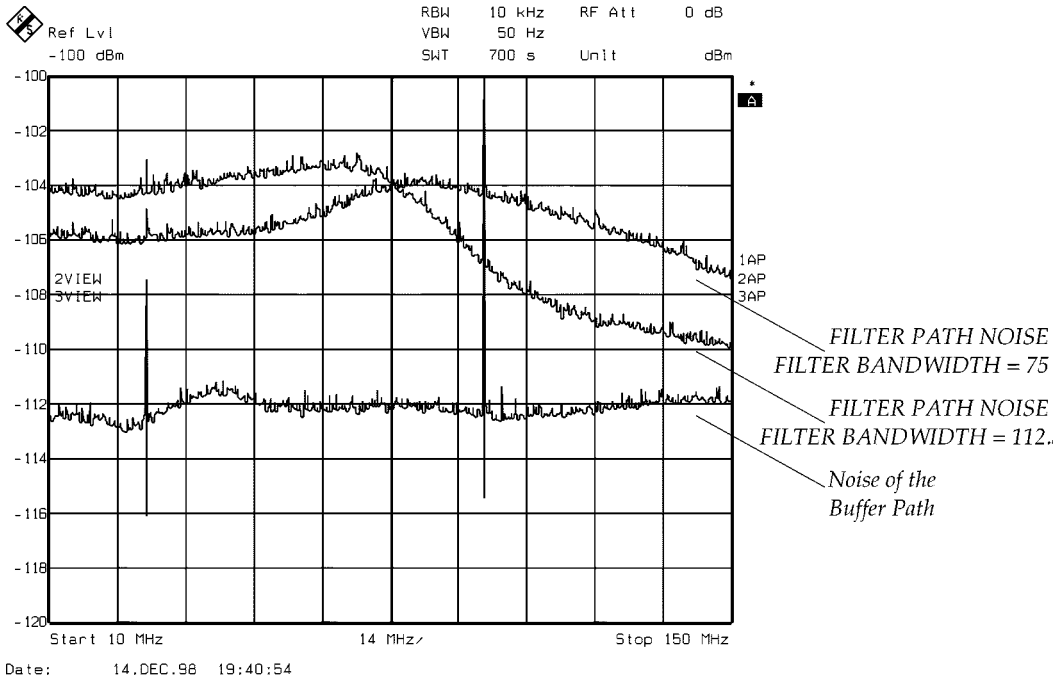

Fig. 16. Filter noise spectrum. Horizontal axis is from 50 to 150 MHz; vertical axis is 2 dB/division.

C_c is a large valued MOS accumulation capacitor, which feeds forward across Mb10 at high frequencies, improving the stability of the loop formed by Mb7, Mb8, Mb10, and Mb11. The resistor R_b isolates the bias voltage generated at the gate of Mb8 from the capacitive load the circuit sees and increases the effectiveness of C_c at high frequencies.

VIII. Implementation Results

In this section, we discuss the filter measurement strategy and test results of a prototype fabricated in a 0.25-μm, n-well CMOS technology. Fig. 12 shows the die photograph of the filter test chip. The active area is 0.15 mm^2. Since MOS capacitors are used, the integrating capacitors are a small fraction of the total filter area. The filter test setup was similar to the scheme used in [19]. Test buffers were included on-chip.

Fig. 13 shows the passband detail of the filter for various settings of the frequency control word. As can be expected, the response scales cleanly, thanks to the constant-C technique described above. Notice that the y-axis in the figure has a resolution of 1 dB per division. Fig. 14 shows the stopband performance of the filter for different frequency settings. (In interpreting this plot, notice that the frequency axis is linear.)

Fig. 15 shows the relative variation of the filter bandwidth as temperature is varied. From this plot, it is clear that the fixed-transconductance bias circuit performs well, even in the presence of mobility degradation effects.

The noise spectra of the filter and buffer paths from 10 to 150 MHz is shown in Fig. 16. The filter output noise spectrum is shown for two different bandwidth settings: the top trace is for the case when the filter bandwidth is set to 75 MHz, while the bottom trace is for a bandwidth of 112.5 MHz.

Table I summarizes important results.

TABLE I
Summary of Measured Characteristics (25 °C, Unless Noted Otherwise). † Remains Practically Fixed at All Bandwidths

Technology	0.25 μm n–well CMOS
Filter type	4th Order Butterworth
Supply voltage	3.3 V
Bandwidth programmability	60–350 MHz
Chip area	0.15 mm^2
Power†	70 mW
DC gain	−0.8 dB
Integrated output noise†	257 μV rms
f_{-3dB} variation for V_{dd} 3 − 3.6V	±2%
f_{-3dB} variation with temperature (0 − 75°C)	±1.80%
Test tone at $f_{-3dB}/3$:	
$V_{ipp,max}$ (THD ≤ −40 dB)†	380 mV
Dynamic range (THD ≤ −40 dB)†	54 dB

IX. Conclusion

Techniques for the realization of very high-frequency programmable CMOS continuous-time filters have been presented. An architecture has been described for programmable transconductors in which *all* parasitic capacitors remain fixed, independent of the programmed transconductance value. This makes possible proper scaling of the frequency response, without affecting noise and dynamic range and allows for a near-optimum design. Simple differential pairs and MOS accumulation capacitors have been used to keep area and noise low. These techniques

were demonstrated with a prototype fourth-order Butterworth low-pass filter tunable from 60 to 350 MHz while consuming 70 mW from a 3.3-V supply in a 0.25-μm, n-well digital CMOS technology.

ACKNOWLEDGMENT

The authors gratefully acknowledge the support of T. R. Viswanathan. They also thank Dr. B.-S. Song, and their colleagues at CISL—Columbia University Integrated Systems Laboratory, especially M. Tarsia, for useful discussions. The authors also wish to thank L. Conyers, K. Kistner, and S. Willard for their support.

REFERENCES

[1] J. Fields, P. Aziz, J. Bailey, F. Barber, J. Barner, M. Heimann, P. Kempsey, L. Mantz, A. Mastrocola, R. Peruzzi, T. Peterson, J. Raisinghani, R. Raushmeyer, M. Saniski, N. Sayiner, P. Setty, and S. Tedja, "A 200 Mb/s CMOS EPRML channel with integrated servo demodulator for magnetic hard disks," in *Proc. Int. Solid-State Circuits Conf.*, San Francisco, CA, Feb. 6–8, 1997, pp. 314–315.

[2] H. Samavati, A. Hajimiri, A. Shahani, G. Nasserbakht, and T. Lee, "Fractal capacitors," in *Proc. Int. Solid-State Circuits Conf.*, San Francisco, CA, Feb. 5–7, 1998, pp. 256–257.

[3] L. J. Pu and Y. Tsividis, "Transistor only frequency selective circuits," *IEEE J. Solid-State Circuits*, vol. 25, pp. 821–832, June 1990.

[4] M. Snelgrove and A. Shoval, "A CMOS biquad at VHF," in *Proc. Custom Integrated Circuits Conf.*, May 1991, pp. 9.1.1–9.1.6.

[5] A. Behr, M. Schneider, S. Filho, and C. Montoro, "Harmonic distortion caused by capacitors implemented with MOSFET gates," *IEEE J. Solid-State Circuits*, vol. 27, pp. 1470–1475, Oct. 1992.

[6] S. Pavan, Y. Tsividis, and K. Nagaraj, "Modeling of accumulation MOS capacitors for analog design in digital VLSI processes," *Proc. IEEE Int. Symp. Circuits and Systems*, vol. 6, pp. 202–205, June 1999.

[7] Y. Tsividis, "Integrated continuous-time filter design—An overview," *IEEE J. Solid-State Circuits*, vol. 29, pp. 166–176, Mar. 1994.

[8] G. Groenewold, "The design of high dynamic range continuous-time integratable bandpass filters," *IEEE Trans. Circuits Syst.*, vol. 38, pp. 836–852, Aug. 1991.

[9] G. Efthivoulidis, L. Toth, and Y. P. Tsividis, "Noise in Gm-C filters," *IEEE Trans. Circuits Syst. II*, vol. 45, pp. 295–302, Mar. 1998.

[10] Z. Czarnul, Y. Tsividis, and S. C. Fang, "MOS transconductors and integrators with high linearity," *Electron. Lett.*, vol. 22, pp. 245–246, 1986.

[11] V. Gopinathan, Y. Tsividis, K. Tan, and R. Hester, "Design considerations for high-frequency continuous-time filters and implementation of an antialiasing filter for digital video," *IEEE J. Solid-State Circuits*, vol. 25, pp. 1368–1378, Dec. 1990.

[12] J. Khoury, "Design of a 15-MHz CMOS continuous-time filter with on-chip tuning," *IEEE J. Solid-State Circuits*, vol. 26, pp. 1988–1997, Dec. 1991.

[13] I. Mehr and D. Welland, "A CMOS continuous-time Gm-C filter for PRML read channel applications at 150 Mb/s and beyond," *IEEE J. Solid-State Circuits*, vol. 32, pp. 499–513, Apr. 1997.

[14] R. Schaumann, "Continuous-time integrated filters—A tutorial," in *Integrated Continuous-Time Filters—Principles, Design and Applications*, Y. P. Tsividis and J. O. Voorman, Eds. Piscataway, NJ: IEEE Press, 1992.

[15] R. Schaumann, K. Laker, and M. Ghausi, *Modern Filter Design—Passive, Active-RC and Switched-Capacitor*. Englewood Cliffs, NJ: Prentice-Hall, 1992.

[16] S. Pavan and Y. Tsividis, "Time scaled electrical networks—Properties and applications in the design of programmable analog filters," *IEEE Trans. Circuits Syst. II*, vol. 47, pp. 161–165, Feb. 2000.

[17] S. Pavan, "High Frequency Continuous-Time Filters in Digital CMOS Processes," Ph.D. Dissertation, Columbia University, New York, NY, May 1999.

[18] R. Zele and D. Allstot, "Low power CMOS continuous-time filters," *IEEE J. Solid-State Circuits*, vol. 31, pp. 157–168, Feb. 1996.

[19] B. Nauta, "A CMOS transconductance-C filter technique for very high frequencies," *IEEE J. Solid-State Circuits*, vol. 27, pp. 142–153, Feb. 1992.

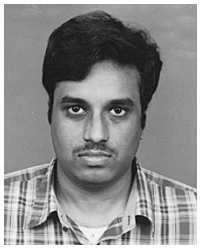

Shanthi Pavan was born in Narasaraopet, South India, in 1973. He received the B.Tech. degree in electronics and communication engineering from the Indian Instiiute of Technology, Madras, in 1995 and the M.S. and D.Sc. degrees from Columbia University, New York, in 1997 and 1999, respectively.

Since 1997, he has been with Texas Instruments R&D Center, Warren, NJ, where he has been involved in designing front-end circuits for hard-disk-drive read channels. His research interests include continuous-time filters, analog-to-digital converters, and MOS device modeling.

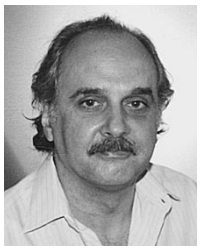

Yannis P. Tsividis (S'71–M'74–SM'75–F'86) received the B.S. degree from the University of Minnesota, Minneapolis, in 1972 and the M.S. and Ph.D. degrees from the University of California, Berkeley, in 1973 and 1976, respectively.

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 National Technical University of Athens, Greece.

Prof. Tsividis received 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 was a co-recipient of the 1987 IEEE Circuits and Systems Society Darlington Best Paper Award. He received the Great Teacher Award from Columbia University.

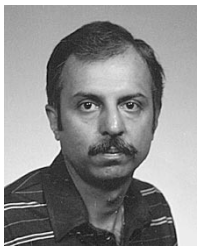

Krishnaswamy Nagaraj received the B.E. degree in electronics from Bangalore University, India, in 1972 and the M.E. and Ph.D. degrees in electrical communication engineering from the Indian Institute of Science, Bangalore, in 1974 and 1983, respectively.

From 1974 to 1984, he was with the R&D laboratories of the Indian Telephone Industries, Bangalore, where he was engaged in the design of telecommunication systems and integrated circuits. During 1985–1986, he was with the VLSI group at the University of Waterloo, Canada, engaged in research and teaching in analog and digital circuits. From 1986 to 1996, he was with AT&T Bell Laboratories, Murray Hill, NJ, and Allentown, PA, where he was engaged in the design of mixed-signal circuits and systems for mass storage, telecommunications, and wireless applications. He is presently a Distinguished Member of Technical Staff and Branch Manager at the Texas Instruments R&D Center, Warren, NJ, where he is leading the development of mixed-signal circuits and systems. He is an Adjunct Associate Professor of electrical engineering with the University of Pennsylvania, Philadelphia.

Dr. Nagaraj received the Distinguished Member of Technical Staff Award from Bell Laboratories in 1991.