Large Dynamic Range Dynamically Biased Log-Domain Filters

Nagendra Krishnapura

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2000

Abstract

Large Dynamic Range Dynamically Biased Log-Domain Filters

Nagendra Krishnapura

This dissertation investigates the enhancement of the dynamic range per unit power consumption of analog filters using dynamic biasing. A technique for realizing dynamically biased log-domain filters while maintaining input-output linearity is presented. This method is much simpler than previously known techniques for realizing large dynamic range filters using syllabic or instantaneous companding. The consequent advantages of the proposed technique are pointed out.

In order to demonstrate the capabilities of the proposed dynamically-biased log-domain filters, a third-order Butterworth filter with a cutoff frequency of 1 MHz is designed in a 0.25 μm BiCMOS technology. Circuit techniques to ensure proper operation of the filter over a wide range of input currents are presented. With suitable dynamic biasing, the fabricated filter can maintain a THD < 40 dB and $S/N > 53.7$ dB for differential input current amplitudes ranging from 3 μA to 2.5 mA (a range of 58.4 dB). In terms of the range of signals that can be handled, the performance is equivalent to that of a conventional filter with a maximum signal to noise ratio of 112 dB. The filter draws 575 μW from a 2.5 V supply in the quiescent condition and 26.1 mW with the maximum input amplitude of 2.5 mA. The maximum power consumption normalized to the order, the dynamic range, and the bandwidth is 5.9×10^{-20} J, which represents more than an order of magnitude of improvement over existing filters.

The design of a current mode peak detector that can provide the dynamic bias to the filter based on the input signal strength is presented. Satisfactory operation of the peak detector over a range of current amplitudes from 1.4 μA to 2.8 mA is verified experimentally. The envelope detector in a 0.25 μm BiCMOS technology occupies 0.12 mm^2 and consumes

162.5 μW in the quiescent condition. The attack time for a 1:2 increase in the input amplitude is less than 1.2 μs and the decay time for a 2:1 decrease in the input amplitude is less than 40 μs over the entire range of input amplitudes.

The feasibility of log-domain filtering in standard CMOS processes is verified by an experimental prototype of a 22 kHz second-order filter using lateral bipolar transistors and pMOS accumulation capacitors. This filter occupies 0.085 mm^2 in a 0.25 μmCMOS technology, consumes 4.1 μW from a 1.5 V supply and has a measured dynamic range of 56.1 dB.

The behavior of noise in companding systems is different from that in classical linear systems due to their inherent internal nonlinearity. Methods for analysis and simulation of noise in instantaneous companding processors are presented. Experimental results corroborating the theory are given.

Contents

List of Figures

List of Tables

Acknowledgments

I am grateful to the guidance provided by Prof. Yannis Tsividis, my thesis advisor, throughout the course of this research. His unwavering faith in this subject kept me from being discouraged. I would like to thank him for infecting me with some of his meticulous nature and optimistic attitude towards research.

The courses taught by Professor John Khoury and discussions with Professors Jesper Steensgaard, Ken Suyama and László Tóth made my stay at Columbia University an educational experience. I would like to express my gratitude to them.

I am thankful to the teachers from my undergraduate days, Professors Bhaskar Ramamurthi, Radhakrishna Rao and Anthony Reddy of the Indian Institute of Technology, Madras for preparing me with the basic knowledge required to pursue further studies in this area.

I am grateful to Dr. Douglas Frey of Silicon Laboratories for early access to some of his work.

I would like to thank my doctoral committee members—Drs. Douglas Frey, Krishnaswamy Nagaraj (Texas Instruments), Professors Jesper Steensgaard, László Tóth, and Yannis Tsividis for their time reading my dissertation and their helpful comments. I also appreciate their accommodation of the very tight schedule in which this dissertation was written.

I am grateful to Aleksander Dec, Shanthi Pavan, Maurice Tarsia and Gregory Ionis, past and present students at the Columbia Integrated Systems Laboratory for all the technical and non-technical discussions. Additionally, thanks are due to: Alex, for the excellent company and shared enthusiasm for experimental work in the lab; Shanthi, for using me (and others) as a sounding board for a barrage of creative ideas and sharpening my insight; Maurice, for helping me out with the tools and fabrication procedures for my design during my stay at Bell Laboratories; Greg, for helping me tame my computer on those occasions when it seemed to acquire a mind of its own.

I would like to extend my thanks to Glenn Cowan, Ruben Herrera, Dandan Li, Georgos Palaskas, Sanjeev Ranganathan, Hai Tao, Jiangtao Yi and Ilya Yusim for making Columbia Integrated Systems Laboratory a fun place to work.

I would like to thank Peter Kinget, formerly of the Design Principles Research Department at Bell Laboratories, for making my stay there in the summer of 1998 a pleasant and productive one.

I am grateful to Dr. Mihai Banu of the Silicon Circuits Research Department at Bell Laboratories for facilitating my research and the fabrication of my chips during the summer of 1999.

I would like to thank Lourdes de La Paz, Marlene Mansfield, Jim Mitchell, Judith Nicholson, Elsa Sanchez and Jody Schneider of the office of Department of Electrical Engineering for helping me with the official procedures of the University. Thanks are also due to John Kazana, the undergraduate lab manager in the electrical engineering department for the loan of measurement equipment.

I dedicate this thesis to my parents whose encouragement of my pursuit of higher studies and love and understanding enabled my research. I thank my brother Manju and sister Sowmya for all the fun times back home.

Chapter 1

Introduction

1.1 Motivation

Integrated continuous time active filters are used in various applications like channel selection in radios, anti-aliasing before sampling, and hearing aids. One of the figures of merit of a filter is the dynamic range; this is the ratio of the largest to the smallest signal that can be applied at the input of the filter while maintaining certain specified performance. The dynamic range required in the filter varies with the application and is decided by the variation in strength of the desired signal as well as the strength of unwanted signals that are to be rejected by the filter.

It is well known that the power dissipation and the capacitor area of an integrated active filter increases in proportion to its dynamic range [1]. This situation is incompatible with the needs of integrated systems, especially battery operated ones. In addition to this fundamental dependence of power dissipation on dynamic range, the design of integrated active filters is further complicated by the reduction of supply voltage of integrated circuits imposed by the scaling down of technologies to attain higher speed and lower power consumption in digital circuits. The reduction in power consumption with decreasing supply voltage does not apply to analog circuits. In fact, considerable innovation is required with a reduced supply voltage even to *avoid increasing* the power consumption for a given signal

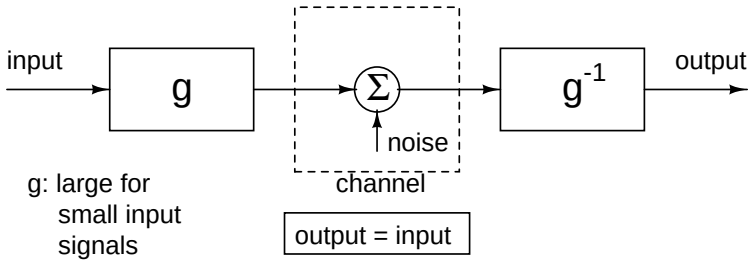

Figure 1.1: Companding with memoryless noisy channels.

to noise ratio (S/N). These aspects pose a great hurdle to the active filter designer.

A technique which has attracted attention recently as a possible route to filters with higher dynamic range per unit power consumption is *companding* [2, 3]. Traditionally companding has been applied to memoryless systems with a dynamic range limited channel (e.g. in telephony). The key idea is to ensure that the signal in the channel stays sufficiently above noise. To ensure this, preamplification is applied. However, it is necessary to avoid overloading the channel as well and for this reason, large signals are preamplified by much smaller amounts than small signals. Thus the entire dynamic range of input signals is amplified by appropriate amounts depending on their strength so that they are near the top of the channel's dynamic range. To restore the output of the channel to the original input levels, the opposite, i.e. small gain for small signals and large gain for large signals is applied. Fig. 1.1 depicts this situation.

The gain can be made to depend on the signal in one of the two following ways.

1. The input "amplifier" includes a nonlinearity whose slope (equivalently, the small signal gain) decreases as the input increases as shown in Fig. 1.2(a). It can be seen that the input is "compressed". The output amplifier has the opposite behavior, as shown in Fig. 1.2(b). This case, where the output of the amplifier is a nonlinear function of the instantaneous value of the input is termed "instantaneous companding".

2. Alternately, the input and output amplifiers can have characteristics of the form $y = gx$ with a variable gain g. The gain of these amplifiers are made to depend on the input signal. The characteristic of such amplifiers are depicted in Fig. 1.3 where the gain g is shown as a parameter.

   If the gain is made to depend on the instantaneous value of the input signal, this case reduces to instantaneous companding[1] described above.

   A distinct situation occurs when the gain is made to depend on an average measure of the input signal strength (e.g. the envelope or the root-mean-square value). This case is termed "syllabic companding".

Although either Fig. 1.2 or Fig. 1.3 can be used to describe the input and output blocks of instantaneously-companding filters, it is customary to use the former. The latter is typically used only in the description of syllabic companding filters.

Companding in telephony (A-law or μ-law [4]) is an example of instantaneous companding. Dolby noise reduction system used in tape recorders is an example of syllabic companding.

Merely substituting a filter in place of the "channel" shown in Fig. 1.1 with either type of input and output amplifiers described above results in a system that is not linear and time-invariant between its input and output. This general problem of applying companding to filters while maintaining input-output linearity and time-invariance has been solved earlier [5, 3, 6, 7]. Several practical implementations have been published as well. While some of them have significantly improved dynamic range per unit power consumption compared to traditional active filters, it is thought that companding can do

[1] Any nonlinearity $y = \alpha(x) = (\alpha(x)/x)x$ can be thought of as an amplifier with a gain $\alpha(x)/x$ which depends on the signal x.

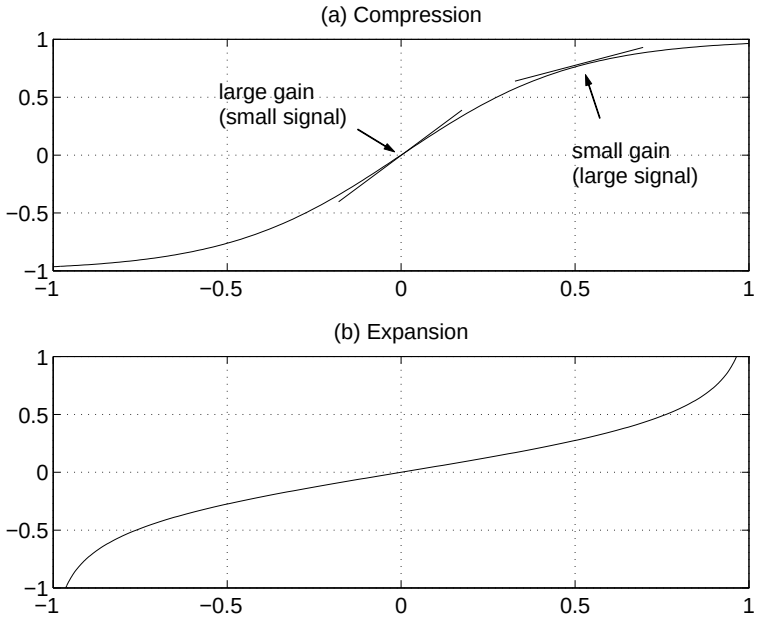

Figure 1.2: Companding using a nonlinearity: (a) Compression at the input, (b) Expansion at the output.

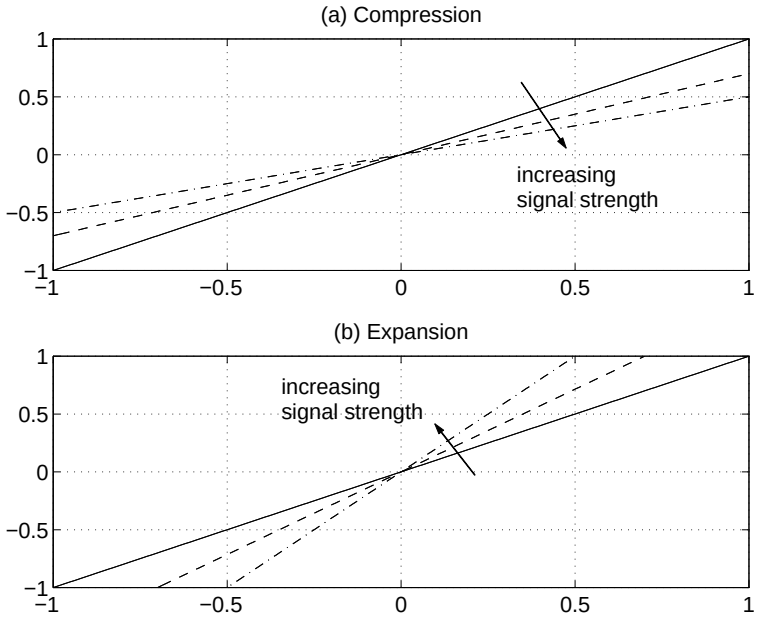

Figure 1.3: Companding using parameterized gains (a) Compression at the input, (b) Expansion at the output.

much better. It is in fact hoped that companding filters can be realized with a lower power consumption per dynamic range than *passive RC/RLC* filters which are assumed to be operating at the fundamental lower limit [1] of power consumption for a given dynamic range.

This thesis will demonstrate the design and implementation of a filter with high dynamic range per unit power consumption using one of several possible companding techniques.

1.2 History and state of the art of companding filters

It was mentioned earlier (Fig. 1.2) that instantaneously-companding filters use nonlinearities at the input and at the output. They are in fact a special case of externally linear, internally nonlinear (ELIN, [3]) filters in which the input nonlinearity is a compression and the output nonlinearity, an expansion. Research into ELIN filters and companding filters started out separately and intersected as they progressed. Although no clear distinction can be drawn between the two, papers that emphasize the synthesis of input-output linear relations using nonlinear blocks can be placed in the former category and papers that emphasize enhancement of dynamic range, in the latter. Both these types can be found in the examples listed below. But the focus of this dissertation is on enhancement of dynamic range using companding.

The idea of using syllabic companding to "improve" a filter was discussed in 1990 in [2]. It was recognized in that reference that the filter presented was not a linear system between the input and the output. The technique improved the selectivity and dynamic range of filters. The system behaved like an input-output linear system only for input signals with slowly varying envelopes.

In order to eliminate this restriction, it was pointed out in [5] that one needs to appro-

priately adjust the state variables of the filter. Techniques for doing this were discussed in [5] for discrete gain changes and in [8] for continuous gain changes. These techniques are applicable to both instantaneous and syllabic companding filters. [9] presents an example of a continuous time syllabic companding based on the formulation in [8].

The earliest form of ELIN filters, dubbed "log-domain" filters due to their use of logarithmic nonlinearity of diodes date back to 1978 [10]. The motivation was not companding, but wide tunability of filter parameters.

[6] presented a compact realization of first-order log-domain filters using translinear loops [11, 12] and through the use of class-AB circuits for high dynamic range, connected them to the concept of companding filters introduced in [2]. To date, log-domain filters have been the most thoroughly investigated species of companding filters.

Log-domain filters received a systematic treatment in [7] in which they were shown to be synthesizable using exponential mappings of state variables in the state equations of linear filter prototypes. Since then, several papers dealing with their analysis and synthesis have been published [13, 14, 15]. A state space formulation for class-AB log-domain filters, which are a class of filters capable of large dynamic range was presented in [16].

[17] presented a log-domain filter with syllabic companding. This was however still based on the formulation of [8]. [18] presented a technique for syllabic companding using dynamic biasing that is unique to log-domain filters and is much simpler to implement than [17]. The potential increase in the dynamic range of syllabically-companding filters was illustrated in [24].

The works mentioned above have dealt with the theoretical aspects of companding/ELIN filters. Notable experimental results can be found in [19, 20, 21, 22, 23]. [19] presented a class-AB log-domain filter in BiCMOS technology which outperformed most

published filters in terms of dynamic range per unit power consumption by a large factor. Log-domain filters at very high frequencies of hundreds of MHz to a GHz are explored in [21, 22]. [23] deals with programmable log-domain filters.

The above are a few examples of the published works in the area of companding filters. In this author's opinion, while theoretical aspects of companding filters have received a fair amount of attention, not enough experimental results are available as yet to conclusively prove the benefits of companding filters and also to verify if companding filters can indeed outperform passive RC/RLC filters in terms of dynamic range per unit power consumption. It is hoped that this dissertation will fill some of that gap.

1.3 Overview of the thesis

The next chapter is devoted to previously published techniques for the implementation of companding filters. After outlining the fundamental limits to dynamic range of traditional linear filters, an overview of previously published examples of instantaneously-companding and syllabically-companding filters is given. This is followed by a discussion of dynamic biasing, which is a recently introduced idea [18, 24] for the realization of high dynamic range filters.

Techniques greatly simplifying the practical implementation of dynamically-biased filters are discussed in Chapter 3. These are compared to existing methods of companding.

The generation of the gain control signal for syllabic companding in general and dynamic biasing in particular is considered in Chapter 4.

Chapter 5 deals with possible implementations of companding filters in pure CMOS technology.

The nonlinear relation between the internal state variables and input (and output) in

companding filters complicates the estimation of noise and renders traditional methods inadequate. Suitable techniques for estimating noise in presence of such internal nonlinearities are discussed in Chapter 6. Analytical methods for calculating noise in instantaneous companding filters are given.

Chapter 7 deals with the design of various prototypes to test the feasibility of ideas presented in Chapters 3 through 5. The design of a 3^{rd}-order dynamically-biased log-domain filter which can operate over a wide range of currents is given. The implementation of the peak detector which generates the dynamic bias is discussed. An experimental prototype for evaluating the feasibility of log-domain filters in standard CMOS processes is presented. In Chapter 8 the measurement techniques and implementation results of the prototypes are presented. In light of the experimental results, possible changes to circuit realizations that may improve the performance and enable more accurate measurements are discussed. The thesis concludes in Chapter 9 with a discussion of achieved results and suggestions for future work.

Chapter 2

Review of Companding Filters

2.1 Power dissipation and dynamic range of filters

2.1.1 Power dissipation and signal to noise ratio in simple circuits

It has been shown elsewhere [1, 25] that the power dissipated in a circuit is directly proportional to the desired signal to noise ratio. Although the exact value of the constant of proportionality is dependent on the circuit details, the relations can be worked out [1, 25] for some simple cases and are given below.

For a first-order passive RC filter (Fig. 2.1) with a pole frequency f_p, the power dissipated in the driving source P_{diss} at a frequency f_p to maintain a given signal to noise ratio S/N (expressed as a ratio of *mean square* quantities) is given by

$$P_{diss} = 2\pi kT f_p \, S/N \tag{2.1}$$

where k is the Boltzmann constant and T is the absolute temperature. This expression and similar expressions for other simple circuits are derived in Appendix A. (2.1) is derived assuming highly ideal conditions and the power dissipation in real circuits may be orders of magnitude higher.

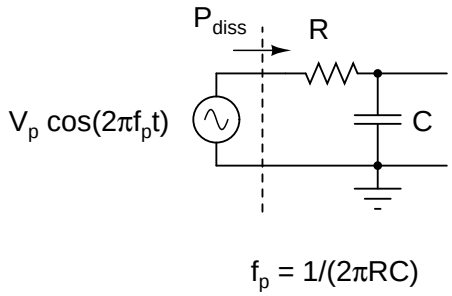

Figure 2.1: Sinusoidal source driving a first-order RC filter.

2.1.2 Power dissipation and dynamic range of conventional active filters

The linear relationship between the input and the output of a filter implies that the output signal strength is proportional to the input signal strength[1]. Also, in a conventional filter, the output noise is a constant. In an active filter, the input signal is limited to a certain upper limit which depends on the supply voltage and the topology used. Typically, the output distortion increases gradually with increasing input signal. In this discussion, for simplicity, we assume a hard limit for the signal power denoted by S_{max}.

This situation is depicted in Fig. 2.2. S denotes the output signal power and N denotes noise power measured in dB relative to some arbitrary reference. The signal to noise ratio S/N (measured in dB) at the filter's output is shown in Fig. 2.3. S/N increases by 1 dB for every dB increase in the input signal until the hard limit S_{max} is reached.

A typical application demands that a certain minimum signal to noise ratio $S/N_{\ min}$ be maintained at the filter's output[2]. This level is indicated on Fig. 2.3. The range of signals over which the filter can be used under this constraint is the dynamic range DR, and is indicated on the figure[3]. It can be seen that the peak signal to noise ratio of the filter

[1]Signal strength could be specified as peak, root mean square (rms) or mean square values of voltage or current, or as power dissipated by the signal in a reference resistor. The decibel (dB) representation of any of these units referred to a desired standard could also be used (e.g. dBm).

[2]For simplicity we assume that this is the only criterion imposed on the filter although in some systems, $S/N_{\ min}$ itself could be a function of signal level.

[3]Conventionally the maximum signal to noise ratio $S/N_{\ max}$ is termed the dynamic range. This amounts to calculating the dynamic range assuming $S/N_{\ min} = 0$ dB. This convention will not be used here. To avoid

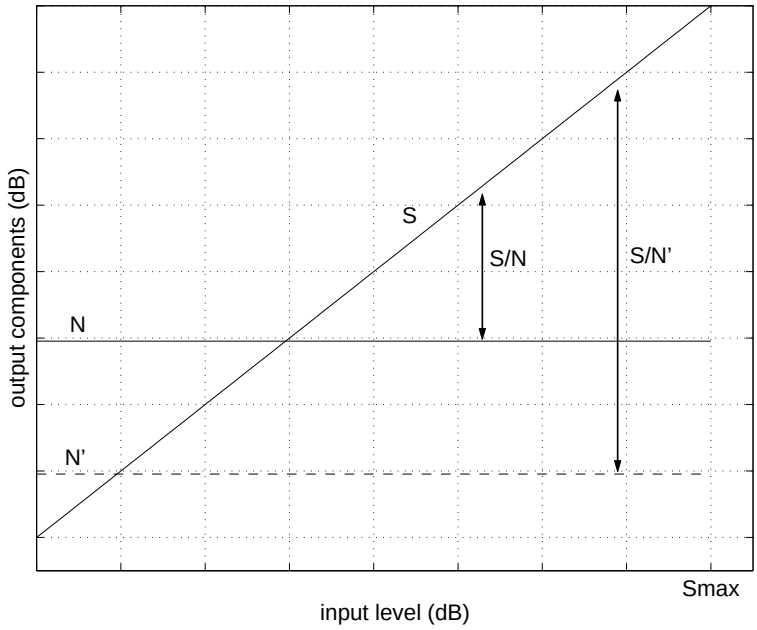

Figure 2.2: Output signal and noise in a typical active filter.

S/N_{peak} which occurs when the input is S_{max}, is far greater than the specified S/N_{min}. This excess signal to noise ratio is an unavoidable characteristic of conventional linear filters. From the unit slope of the S/N curve in Fig. 2.3, it can be inferred that the excess signal to noise ratio is equal to the dynamic range (DR) itself.

If an increase in dynamic range is desired, the noise floor of the filter has to be lowered. The new noise floor N' of the filter is shown using a dashed line in Fig. 2.2. The corresponding signal to noise ratio S/N' is plotted in Fig. 2.3. The dynamic range increases (by the amount the noise is lowered) to DR'. The peak signal to noise ratio also increases by the same amount from S/N_{peak} to S/N'_{peak}. Thus the peak signal to noise ratio is inextricably linked to the dynamic range in conventional filters.

This is an undesirable situation because the power dissipated in a circuit is directly proportional to the signal to noise ratio as discussed in section 2.1.1. To obtain the desired dynamic range, excessive signal to noise ratio and consequently excessive power

confusion while discussing companding filters, in this work, the term dynamic range is reserved to denote the possible range of signal variations subject to a given constraint.

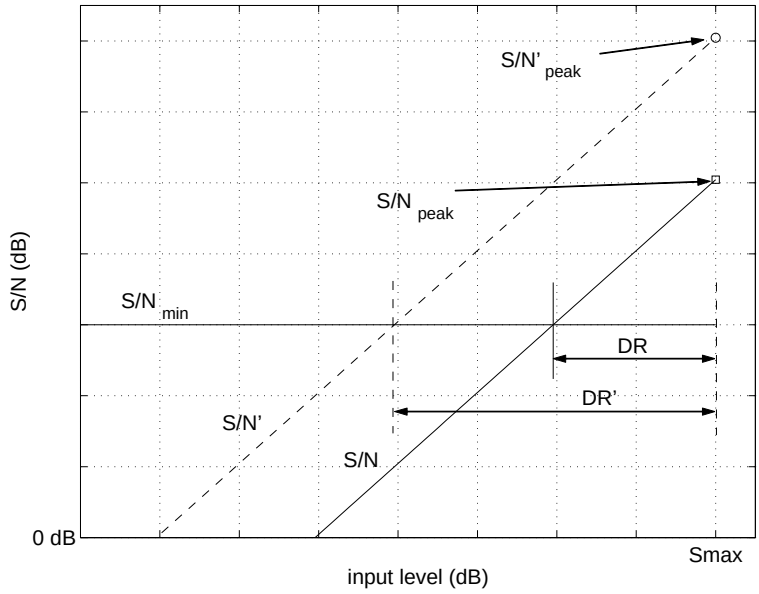

Figure 2.3: Signal to noise ratio and dynamic range of a typical active filter.

consumption are inevitable. An increase in dynamic range necessitates an equal increase in the peak signal to noise ratio although the signal to noise ratio demanded by the application remains unchanged and may be much lower than that.

Thus there appears to be a potential for significant power savings in those circuits intended to be used over a large dynamic range of input signals while maintaining a modest signal to noise ratio. Given that the excess signal to noise ratio is equal to the dynamic range and dynamic range of 10^6 (60 dB) is not unheard of, power reduction by a factor of a *million* is possible if this excess signal to noise ratio is eliminated![1].

2.1.3 Power dissipation and dynamic range of companding filters

Consider the two filters Filter 1 and Filter 2 shown in Fig. 2.4(a). Filter 1 is a linear time-invariant active filter with a transfer function $H(s)$. Let S_{max} and N respectively be the maximum possible signal level and the noise at the output of the filter. These quantities are marked on a log scale in Fig. 2.4(b). Assuming for simplicity that the specified S/N_{min} is 0 dB, the dynamic range of this filter DR is the ratio S_{max}/N.

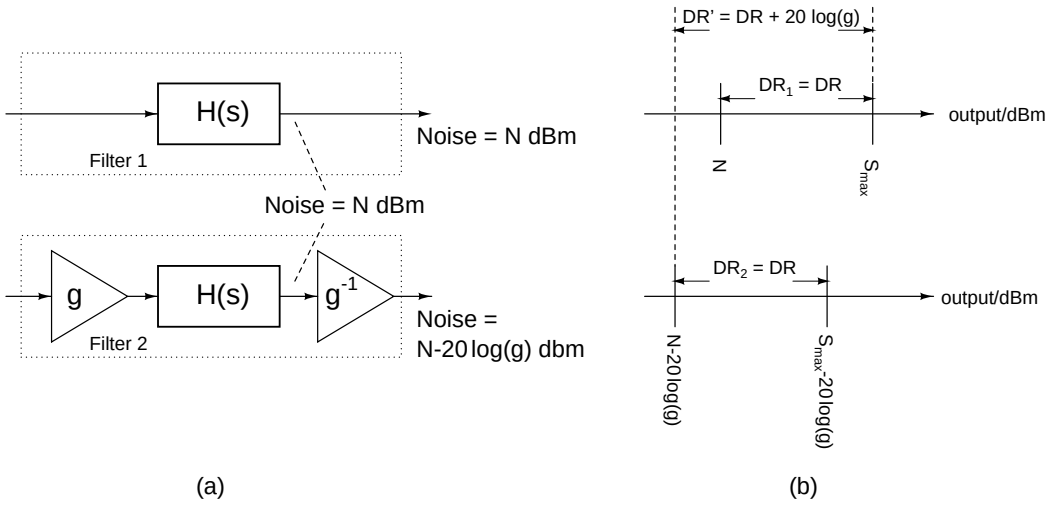

Figure 2.4: Filters with skewed operating ranges.

Filter 2 is the same filter embedded between two amplifiers of gain g and g^{-1} where g is greater than unity. The transfer functions of Filter 1 and Filter 2 are identical. Assuming noiseless amplifiers, both the maximum output signal and the output noise of Filter 2 are reduced by a factor g when compared to Filter 1. This is depicted in the lower part of Fig. 2.4(b). Hence, the dynamic ranges of the two filters are identical.

It can be seen that Filter 2 outperforms Filter 1 for small signals due to its lower noise level. However, the maximum signal that can be fed to Filter 2 is lower as well and Filter 1 would be preferable in a large signal situation.

By switching between these two configurations based on the input signal strength, optimal conditions can be had for both large and small input signals. However it must be ensured that the original linear time-invariant nature of the filter is not destroyed in presence of such switching. *If* this can be accomplished, we would have a filter whose dynamic range is increased to $DR' = DR + 20\log(g)$ as shown in Fig. 2.4(b). Thus, the dynamic range (expressed in terms of mean square quantities) has increased by g^2. Increasing the dynamic range of the original filter (Filter 1) by a factor of g^2 using conventional

means (i.e. lowering its noise floor) would also increase its power dissipation by a factor g^2 (section 2.1). Therefore, if the switching between the two situations while maintaining the linear time-invariant nature of the system could be implemented without a g^2 times larger power dissipation, we would have a filtering technique that has a superior dynamic range per unit power consumption when compared to conventional active filters.

The description above considered two discrete values for the gain used at the input and the output. In general, any number of discrete values can be used for g. g can be even be made to vary continuously with the input signal strength. Since the filter embedded between the amplifiers is the dominant source of noise, best performance is obtained if the signal fed to this filter stays as much above its noise as possible.

It was stated in section 1.1 that it is indeed possible to maintain input-output LTI [5, 8, 3] behavior in the presence of signal dependent gains at the input and the output of the filter. The following sections describe existing methods for realizing input-output linear filters that use the two types of companding mentioned in Chapter 1. The issues involved in the implementation of these techniques are briefly touched upon.

2.2 Companding techniques

2.2.1 Instantaneously-companding filters

Externally linear, internally nonlinear filters

A linear integrator is shown in Fig. 2.5(a). u, x and y denote the input, the output and the state variable respectively. The state variable description of this system is given by

$$\frac{dx}{dt} = ku \tag{2.2}$$

$$y = x \tag{2.3}$$

In [26], it is shown that the input-output relationship given above can be obtained with another system with a transformed state variable v [27]. Let the transformation be described

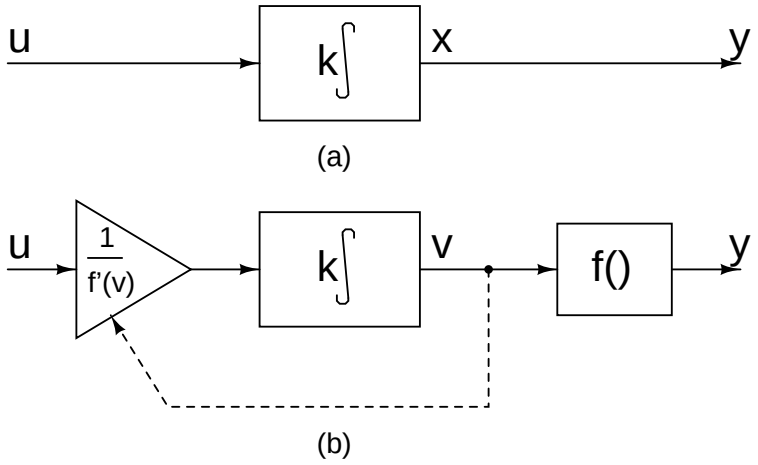

Figure 2.5: (a) Internally linear integrator, (b) ELIN integrator.

by

$$x = f(v) \tag{2.4}$$

where f is a continuous monotonic nonlinearity. Substituting (2.4) into (2.2) and (2.3), using the relation $df(v)/dt = f'(v)dv/dt$ and rearranging the terms, we obtain the state equations of the system in terms of the transformed state variable v.

$$\frac{dv}{dt} = k\frac{u}{f'(v)} \tag{2.5}$$

$$y = f(v) \tag{2.6}$$

Fig. 2.5(b) shows the realization of this transformed system. The input u is divided by $f'(v)$ before being fed to the integrator. The integrator is followed by the nonlinearity $f()$. The input-output behavior of this system is identical to that of the integrator in Fig. 2.5(a). Because of the presence of the nonlinearity $f()$ in the system, this is an externally linear, internally nonlinear (ELIN) integrator [26, 3].

The ELIN integrator in Fig. 2.5(b) turns into a companding integrator if an expanding nonlinearity $f()$ is used at the output. An expanding nonlinearity $f()$ has a slope $f'(v)$ that increases with increasing v. The input "amplifier" would thus have a gain $1/f'(v)$ that

decreases with increasing v.

Higher order filters are constructed using integrators. The ELIN versions of these filters can be derived either by substituting the integrators in the filters by ELIN integrators shown in Fig. 2.5(b) or by transforming the state variables of the system using nonlinear mapping [27]. These two methods are illustrated below for a first-order filter.

A first-order filter using feedback around the linear integrator in Fig. 2.5(a) is shown in Fig. 2.6(a). The transfer function of this filter is

$$H(s) = \frac{Y(s)}{U(s)} = \frac{k}{s+a} \tag{2.7}$$

To obtain an ELIN first-order filter with the transfer function above, the internally linear integrator must be substituted by the ELIN integrator in Fig. 2.5(b). Fig. 2.6(b) shows the resulting filter. Fig. 2.6(c) shows a further transformation whereby the feedback path is placed around the integrator used inside the companding filter. The equivalence between Fig. 2.6(b) and Fig. 2.6(c) can be easily seen.

The state equations of the first-order filter in Fig. 2.6(a) are

$$\frac{dx}{dt} = -ax + ku \tag{2.8}$$

$$y = x \tag{2.9}$$

Using $x = f(v)$ as before, these state equations can be expressed in terms of the new state variable v

$$f'(v)\frac{dv}{dt} = -af(v) + ku \tag{2.10}$$

$$y = f(v) \tag{2.11}$$

Dividing the first equation above by $f'(v)$, the state equations can be rewritten as

$$\frac{dv}{dt} = -a\frac{f(v)}{f'(v)} + \frac{k}{f'(v)}u \tag{2.12}$$

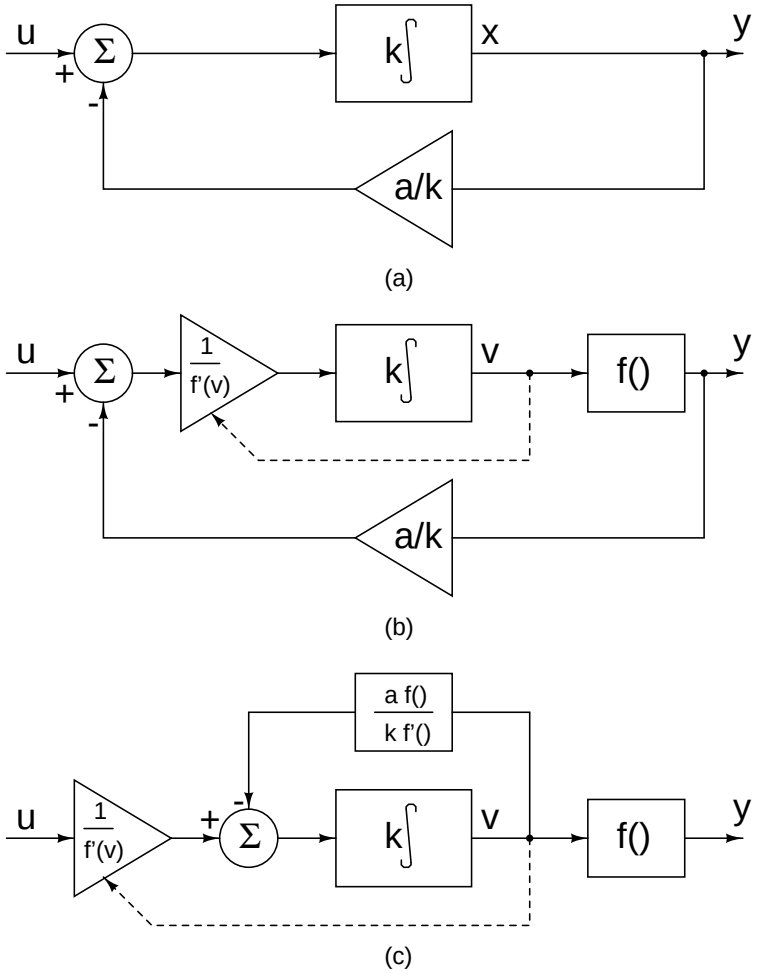

Figure 2.6: (a) Internally linear first-order filter, (b) ELIN first-order filter, (c) ELIN first-order filter with transformed feedback path.

$$y = f(v) \tag{2.13}$$

It can be seen that the block diagram in Fig. 2.6(c) realizes these equations.

When $f()$ is an exponential

Log-domain filters are obtained when the nonlinearity $f()$ is an exponential. The exponential nonlinearity could be realized using a bipolar transistor whose characteristic is given by

$$f(v) = I_s \exp(\frac{v}{V_t}) \tag{2.14}$$

where I_s, v and $f()$ are respectively the saturation current, the base-emitter voltage, and the collector current of the bipolar junction transistor and V_t is the thermal voltage kT/q. Using this equation for $f()$, the first-order filter in Fig. 2.6(c) can be redrawn as shown in Fig. 2.7(a). Since the derivative of an exponential is also an exponential, the feedback term reduces to subtraction of a constant. Note that the input u and the output y in this block diagram of a log-domain filter are currents.

In practice the integrator is realized using a capacitor C which has the input-output relation

$$v = \frac{1}{C}\int i_c(t)dt \tag{2.15}$$

where v is the voltage across the capacitor ("output" of the integrator) and i_c is the current through the capacitor ("input" to the integrator). Modifying the block diagram in Fig. 2.7(a) to use the capacitive integrator described by (2.15) results in Fig. 2.7(b).

Translinear loops

The output of the input amplifier in Fig. 2.7(b) is $u(kCV_t)/y$. This is a product of two currents divided by another current and can be realized [6] using translinear loops [11, 12]. Translinear loops are loops of base-emitter junctions of transistors. Fig. 2.8 shows an example of a translinear loop with four transistors Q_{1-4}. The following relations hold true for this circuit:

$$V_{BE1} + V_{BE2} = V_{BE3} + V_{BE4} \tag{2.16}$$

$$V_t \ln\frac{i_1}{A_1 I_{s0}} + V_t \ln\frac{i_2}{A_2 I_{s0}} = V_t \ln\frac{i_3}{A_3 I_{s0}} + V_t \ln\frac{i_4}{A_4 I_{s0}} \tag{2.17}$$

$$\frac{i_1 i_2}{A_1 A_2} = \frac{i_3 i_4}{A_3 A_4} \tag{2.18}$$

where V_{BE1-4}, i_{1-4} and A_{1-4} are respectively the base-emitter voltages, emitter currents and normalized areas of transistors Q_{1-4} and I_{s0} is the saturation current of a transistor

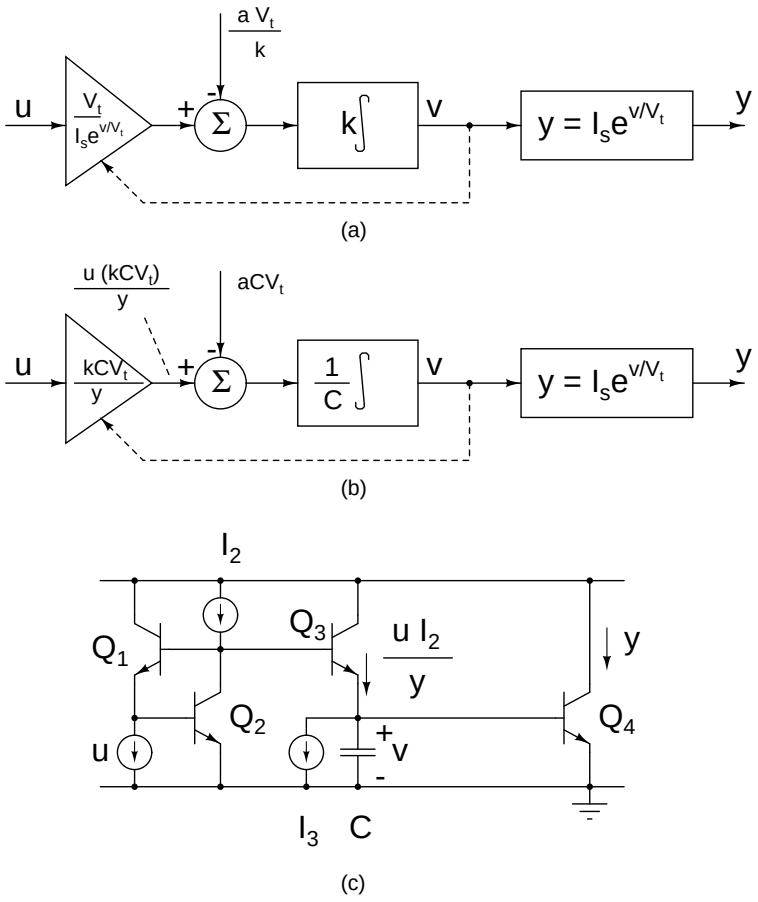

Figure 2.7: Realization of a log-domain first-order filter.

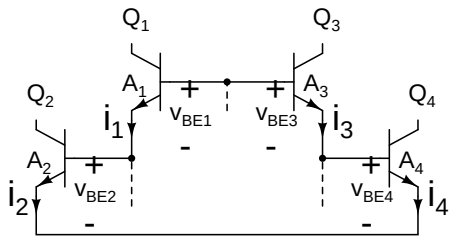

Figure 2.8: Translinear loop with four transistors.

with unit area. Assuming equal areas for all transistors, we have

$$i_1 i_2 = i_3 i_4 \tag{2.19}$$

Such equations relating products of currents are characteristic of translinear loops. The circuit in Fig. 2.8 is by itself incomplete; the actual currents flowing in the transistors depends on other elements connected to the circuit. The relation between the products of currents given above can be used to realize interesting nonlinear functions [11, 12].

More transistors can be included in the translinear loop and relations similar to (2.18) derived. Usually, the number of transistors is even and connected such that (2.16) has equal number of terms on either side. If this condition is not satisfied, the equation relating the products of currents (like (2.18)) contains the saturation current I_{s0}, making the relation dependent on process and temperature, usually an undesirable situation.

In the discussions of the circuits that follow, unless mentioned otherwise, transistors are assumed to have equal areas and infinite current gains β. The effects of deviations from this assumption are discussed separately. Therefore, the simplified relation (2.19) can be used with collector or emitter currents.

Log domain filters

Fig. 2.7(c) shows the first-order log-domain filter [6] realized using bipolar transistors. Q_4 implements the output nonlinearity. The capacitor C implements the integrator. I_3 is the

constant term aCV_t subtracted at the input of the integrator in Fig. 2.7(b). Q_1, Q_2, Q_3 and I_2 are connected as shown in Fig. 2.7(c) to realize the input "amplifier". The transistors Q_{1-4} form a translinear loop identical to the one in Fig. 2.8. (2.19), which relates products of emitter currents in a translinear loop, can be used to determine the emitter current of Q_3 in Fig. 2.7(c). Q_3's emitter current is found to be uI_2/y and has the same form as the output of the input amplifier in Fig. 2.7(b). It can thus be seen that the circuit in Fig. 2.7(c) emulates the block diagram of Fig. 2.7(b). By comparing Fig. 2.7(b) with Fig. 2.7(c), the correspondence between the currents I_2 and I_3 in the circuit and the constants a and k in the block diagram can be seen. The transfer function of the filter in Fig. 2.7(c) is given by

$$\frac{Y(s)}{U(s)} = \frac{I_2/CV_t}{s + I_3/CV_t} \tag{2.20}$$

$$= \frac{k}{s+a} \tag{2.21}$$

The pole of the low pass filter is determined by the current I_3 and capacitor C and the dc gain is determined by the ratio I_2/I_3.

The feedback used in Fig. 2.6(a) reduces to subtraction of a constant current in the log-domain version (Fig. 2.7(c)). This fact can be intuitively understood from Fig. 2.9. Fig. 2.9(a) shows a first-order RC filter (lossy integrator) with a zero input. As is well known, the output y decays exponentially from its initial value y_0. The first-order log-domain filter in Fig. 2.7(c) reduces to Fig. 2.9(b) when the input u is zero. The constant current I_3 causes the voltage v across the capacitor C to decrease linearly. This linear decrease, through the exponential nonlinearity of the bipolar transistor causes the output y to decay exponentially.

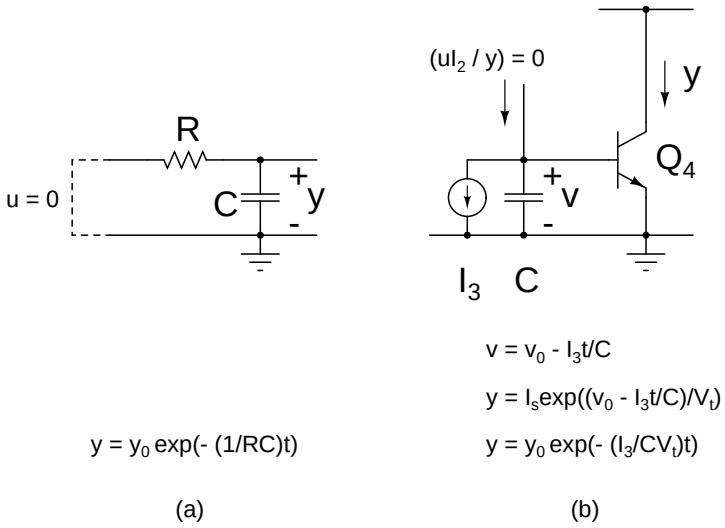

Figure 2.9: Damping in a log-domain filter.

Higher-order log-domain filters

Higher order log-domain filters can be synthesized either by using exponential mappings of the state variables in the state equations of the linear filter prototype [7] or by replacing the integrators in the linear filter prototype using log-domain integrators ([14] for LC ladder prototypes). State variables are a set of independent variables in the circuit that can be used to completely describe the operation of the circuit given the initial conditions [28]. In the state space descriptions of LC/active RC filters, inductor currents and/or capacitor voltages are usually used as state variables. Note that the scaled versions of inductor currents and/or capacitor voltages in the circuit can also be used as state variables. In the following discussion, the state variables in the prototype are referred to as "linear-domain" state variables. In a log-domain filter derived from a linear prototype, the voltages across the capacitors—the "log-domain" state variables—are related logarithmically to the linear-domain state variables, and the collector currents are scaled versions of the linear-domain state variables.

Fig. 2.10 shows a doubly terminated third-order RLC filter. x_2 is the inductor current

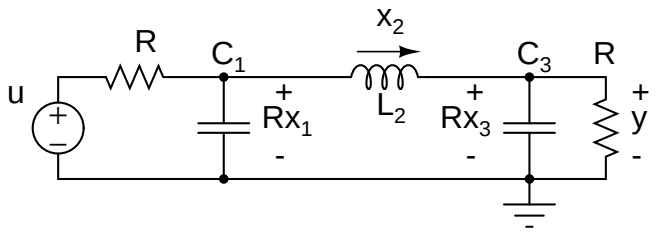

$v_1 = V_t \ln(x_1 / I_0)$	$x_1 = I_0 \exp(v_1 / V_t)$
$v_2 = V_t \ln(x_2 / I_0)$	$x_2 = I_0 \exp(v_2 / V_t)$
$v_3 = V_t \ln(x_3 / I_0)$	$x_3 = I_0 \exp(v_3 / V_t)$

Figure 2.10: Ladder filter prototype and logarithmic (exponential) mappings.

and Rx_1 and Rx_3 are the capacitor voltages where R is the termination resistance of the RLC filter. x_1, x_2, and x_3, which have dimensions of currents, are chosen to be the state variables of the prototype filter. The equations (Kirchoff's voltage and current laws) governing the ladder filter can be rewritten in terms of new state variables v_1, v_2 and v_3 that are logarithmically related to the original state variables x_1, x_2 and x_3. The mappings are shown in the figure. I_0 is a normalizing current (For the example in Fig. 2.7, it was the saturation current I_s). The equations so obtained can be realized using translinear loops. In the log-domain version of the filter v_1, v_2 and v_3 would be the voltages across the integrating capacitors and x_1, x_2 and x_3 (or their scaled versions) would be the collector currents of the bipolar transistors.

Quiescent currents in log-domain filters

Since the state variables of the prototype filter are scaled versions of the collector currents of the transistors of the log-domain filter, they need to have positive quiescent values and stay positive for all inputs for proper operation of the filter. In the LC ladder filter in Fig. 2.10 a positive quiescent value of the state variables can be achieved by having a positive quiescent input u. This however is not the case with all filters. The issue of main-

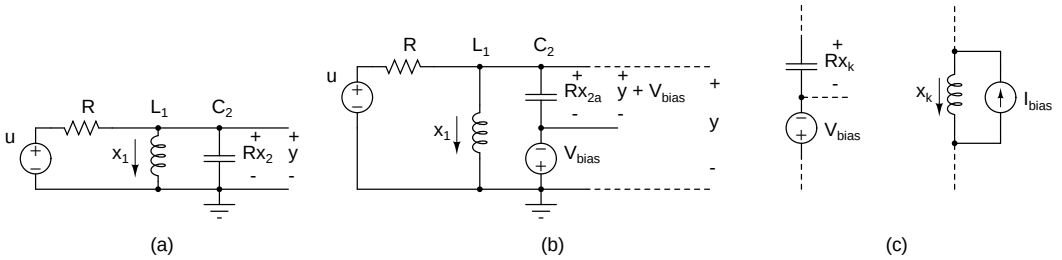

Figure 2.11: (a) RLC Bandpass filter, (b) Modification to ensure positive quiescent state variable x_2, (c) General modification to ensure positive quiescent state variables x_k.

taining positive quiescent collector currents in log-domain filters has been treated partially in [7, 27, 14, 29]. A systematic procedure to synthesize log-domain filters with positive currents in all their transistors is given in [30]. Without going into details, the principle behind these methods can be illustrated using Fig. 2.11.

Fig. 2.11(a) shows an RLC bandpass filter whose state variables are x_1 and x_2. It is obvious that the quiescent value of the capacitor voltage Rx_2 is identically zero regardless of the dc value of u due to the dc short circuit across the capacitor through the inductor. x_2 would swing both positive and negative on application of a sinusoidal signal u. This situation is problematic if the RLC band-pass filter has to be transformed into a log-domain filter (or if it is intended to be built with electrolytic capacitors that explode if a voltage of incorrect polarity is applied to them!). The solution, shown in Fig. 2.11(b), is to provide a dc bias to the capacitor by connecting a voltage source in series. With this arrangement, the quiescent value of the new state variable x_{2a} is V_{bias}/R. The relation between u and y (measured across the combination of capacitor C and the bias source) is unaffected by this addition. Thus, the log-domain filter synthesized using exponential mappings of state variables in the circuit equations of the filter in Fig. 2.11(b) would have positive quiescent current in all its transistors. The voltage Rx_{2a} across the capacitor would be $y+V_{bias}$ where y is the intended output and V_{bias} is the bias voltage. Therefore x_{2a}, which appears as the

collector current of a transistor in the log-domain version of the circuit, includes the bias. The bias should be subtracted from x_{2a} to obtain the desired output. The state variables in any RLC filter can be made to have the desired values by adding voltage sources in series with the capacitors or current sources in parallel with the inductors. Fig. 2.11(c) shows the general scheme. The filter so obtained can be transformed into a correctly operational log-domain filter through exponential mappings of state variables.

ELIN vs. companding log-domain filters

The paragraphs above discussed the synthesis of ELIN filters with an exponential nonlinearity at the output. It was mentioned earlier that ELIN filters are companding filters when an expanding nonlinearity is used at the output. Although the exponential nonlinearity at the output appears to be "expanding", log-domain filters of the type shown in Fig. 2.7(c) are not companding filters. This can be inferred from the examination of the nonlinearity [31]. It is noted that noise analysis [32, 33] shows that output noise does not vary appreciably with the signal.

In Fig. 2.7(c), the input u is the sum of a dc bias current and a signal current. The negative swings of the signal current should not be so large as to reduce the total input u to very small values which would cause the transistor's characteristic to deviate from the exponential. Assuming symmetric input signals this means that the peak value of the signal input is constrained to be somewhat less than the bias. With this constraint, the nonlinearity of the exponential is not exercised as much and the change in output noise with signal is not significant.

A type of log-domain filter which does incorporate instantaneous companding is the class-AB filter shown in Fig. 2.12 [6]. The filter has differential inputs u_p, u_n and differential outputs y_p, y_n. The two halves of the filter, which are identical to the first-order filter in

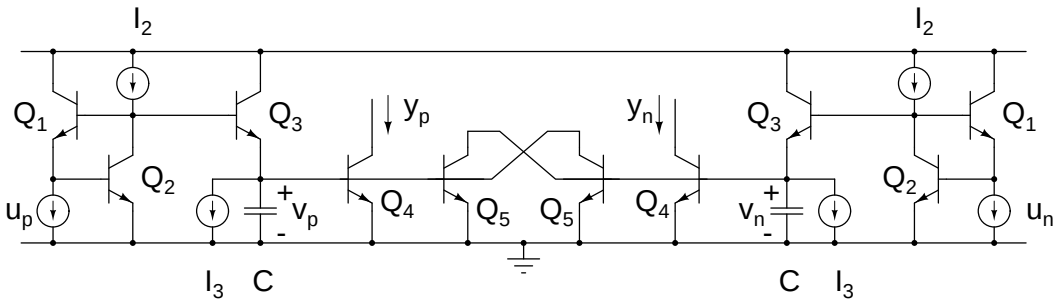

Figure 2.12: Class-AB instantaneous companding log-domain filter.

Fig. 2.7(c) are coupled using transistors Q_5. The resulting filter behaves as a first-order filter between the differential inputs $u_p - u_n$ and outputs $y_p - y_n$. In this case, the peak value of the input signal is no longer constrained by the bias current used at the input. This is because the input signal $u_p - u_n$ can be symmetric even if u_p and u_n individually are not. u_p and u_n could be half wave rectified or geometrically split [19] versions of the input, with their difference being the desired input signal. For large positive inputs, u_p is very nearly equal to the input and u_n is nearly zero. The opposite situation exists for a large negative input. Since only one half of the circuit is active for each polarity of the input signal, this circuit is known as a class-AB log-domain filter.

An integrated class-AB log-domain filter with a large dynamic range was presented in [19]. But the circuit used to split the input signal into two asymmetric signals was external to the chip. [22] presented a class-AB log-domain filter for high frequencies that included an on-chip signal splitter. A formal theoretical treatment of class-AB filters can be found in [16].

Practical realization of log-domain filters

The main problems in the practical realization of log-domain filters arise from the non-exponential behavior of the bipolar transistor. The deviation from the ideal characteristic results in distortion at the output of the log-domain filter. The Early voltage V_A and the

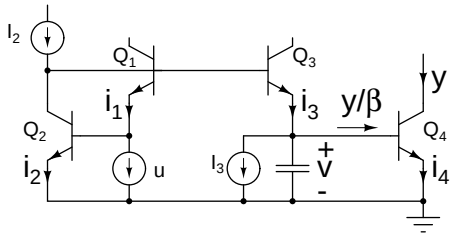

Figure 2.13: First order log-domain filter in Fig. 2.7(c).

ohmic resistances in series with the base and emitter contribute to deviations from the exponential. The problem of low Early voltage can be countered by cascoding the transistors so that a constant voltage is maintained across their collector and emitter. The simplest technique to combat series resistances is to use transistors of sufficiently large area so that the resistances are minimized. But too large transistors compromise the frequency response of transistors due to large parasitic capacitances. A technique to maintain ELIN behavior in presence of series ohmic resistances is presented in [34]. But due to its complexity, it may introduce frequency limitations similar to that of a large transistor.

Another major nonideality is the finite current gain of the bipolar transistors. The first order log-domain filter in Fig. 2.7 is redrawn in Fig. 2.13. The connections to the collectors are not shown. It is assumed that they are biased at voltages that are large enough to keep the transistors out of saturation. It can be seen that the current flowing through the capacitor is reduced from its ideal value by an amount equal to the base current of Q_4. Assuming that Q_4 has a current gain β and that Q_{1-3} have infinite current gains, the differential equation governing this circuit can be written as follows.

$$\frac{dy}{dt} = -\frac{I_3}{CV_t}y + \frac{I_2}{CV_t}u + \frac{y^2}{\beta CV_t} \tag{2.22}$$

The last term on the right hand side introduces a nonlinearity into the equation and disappears when $\beta = \infty$.

Base current compensation circuits that sense the output current and inject an appro-

priate current into that base are not very effective when the current is time varying. This is because the small base current causes the compensation circuits to be slow. The best known solution to this problem is to use an alternative topology in which the parasitic base current drawn from the capacitor is independent of the signal [19].

Log-domain filter topology immune to base current errors

The log-domain filter topology proposed in [19] is described in detail below. The derivation of the topology is shown as a logical sequence of steps, although that is probably not how it was originally conceived. Since at each step there is more than one choice, some seemingly arbitrary assumptions are made so that the topology in [19] is arrived at. While deriving this topology, it is assumed that the base currents of the transistors are zero. The effect of base currents can be examined after the derivation.

The translinear loop of transistors Q_{1-4} is the heart of the circuit in Fig. 2.13. As described earlier in this section, this translinear loop is used to derive the current that is to be integrated in the capacitor C. The emitter (and collector) currents of Q_{1-4} are denoted by i_{1-4} respectively. The translinear loop alone is shown in Fig. 2.14(a-i). Another translinear loop using four transistors is shown in Fig. 2.14(a-ii). The transistors and their currents in the two translinear loops are named identically to reflect an underlying correspondence. Disregarding for the moment the method by which these currents are established in the transistors, it can be seen that (2.19) holds for both of the translinear loops in Fig. 2.14(a).

$$i_1 i_2 = i_3 i_4 \tag{2.23}$$

This equation was used to derive the current flowing into the integrating capacitor C in Fig. 2.7(c). Since both of the translinear loops in Fig. 2.14(a) obey this equation, conceivably, both of them can be used to implement the first order filter in Fig. 2.13.

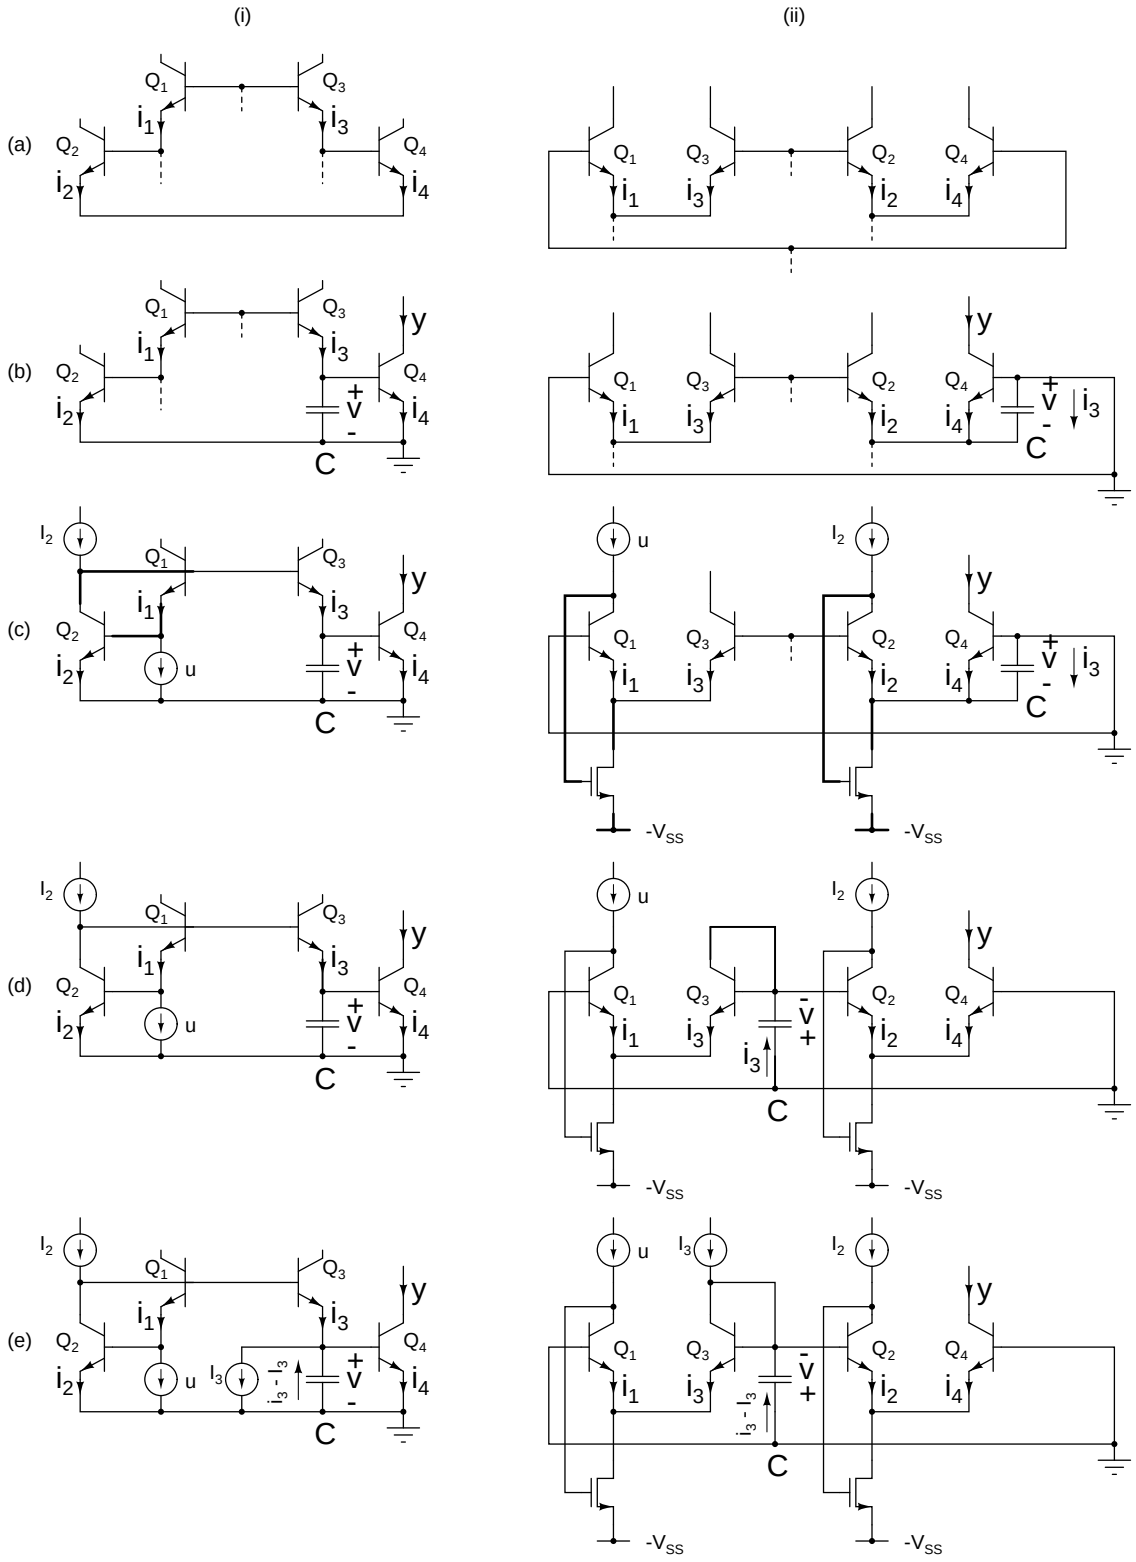

Figure 2.14: Deriving an alternative topology of a first-order log-domain filter.

Fig. 2.14(b) shows the same translinear loops with one of the nodes grounded in each case. Q_4 is assumed to be the transistor realizing the exponential nonlinearity. A capacitor C is connected across the base and emitter of Q_4. The collector current of Q_4 is taken as the output y. The voltage across the capacitor is denoted by v and is the state variable of the circuit. For Fig. 2.14(b-ii) to be identical to Fig. 2.14(b-i), a current equal to i_3 must be driven through the capacitor. The method used to accomplish this is discussed later.

Fig. 2.14(c-i) shows the circuit with the input current u and a constant current I_2 fed into transistors Q_1 and Q_2 respectively. To force I_2 into the collector of Q_2, a negative feedback loop—shown with thick lines—is formed around Q_2. u is pulled out of the emitter of Q_1.

Fig. 2.15 shows the feedback arrangement that is used to force I_2 into the collector of Q_2 in Fig. 2.14(c-ii)[19]. A transistor M_2 provides negative feedback around Q_2. If the collector current of Q_2 exceeds I_2, the collector voltage falls and vice versa. The collector voltage is sensed by the gate of M_2 whose drain current varies accordingly. When equilibrium is reached, Q_2's collector current equals I_2. M_2's drain current equals the sum of i_2 and i_{2x} which respectively are the emitter current of Q_2 and any other current that may be pumped into the emitter of Q_2. The collector voltage of Q_2 is determined by V_{SS} and the the parameters of M_1. These are assumed to be such that $Q_{1,2}$ are operating in the active region.

A similar feedback loop is used to force the input u into the collector of Q_1 as shown in Fig. 2.14(c-ii).

In Fig. 2.14(c-ii), both of the nodes to which the capacitor is connected are low-impedance nodes. Any current injected into the top node flows into ground and that injected into the bottom node flows into M_2. Therefore, the intended current i_3 cannot

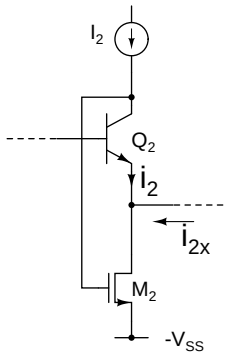

Figure 2.15: Feedback around Q_2 in Fig. 2.14 to force a current into its collector.

be independently forced into the capacitor in the present incarnation of the circuit. This situation can be remedied by connecting the capacitor as shown in Fig. 2.14(d-ii). The voltage across the capacitor is $v = V_{BE4} - V_{BE2}$ instead of V_{BE4}. Q_2 carries a constant current I_2 and therefore, V_{BE2} $(= V_t \ln(I_2/I_s))$ is a constant voltage. Since the capacitor is an open circuit at dc, a dc shift in the capacitor voltage can be introduced without changing the operation of the circuit. i_3 can now be drawn from the capacitor by connecting the collector of Q_3 as shown in Fig. 2.14(d-ii). The damping current I_3 can be connected as shown in Fig. 2.14(e) to complete the derivation.

The circuit in Fig. 2.14(e-ii) is complete and performs the same function as the one in Fig. 2.14(e-i) which is the first-order log-domain filter in Fig. 2.7. The transfer function between y and u is given by

$$\frac{Y(s)}{U(s)} = \frac{I_2}{I_3}\frac{1}{1 + sCV_t/I_3} \tag{2.24}$$

Fig. 2.16 shows the first order log-domain filter in Fig. 2.14(e-ii) which is modified such that it can be operated with a positive power supply. The sources of M_1 and M_2 are returned to ground and the bases of $Q_{1,4}$ are biased with a positive voltage V_{base} (equal to V_{SS} in Fig. 2.14(e-ii)).

When the current gains of the transistors are finite, the only base current that affects

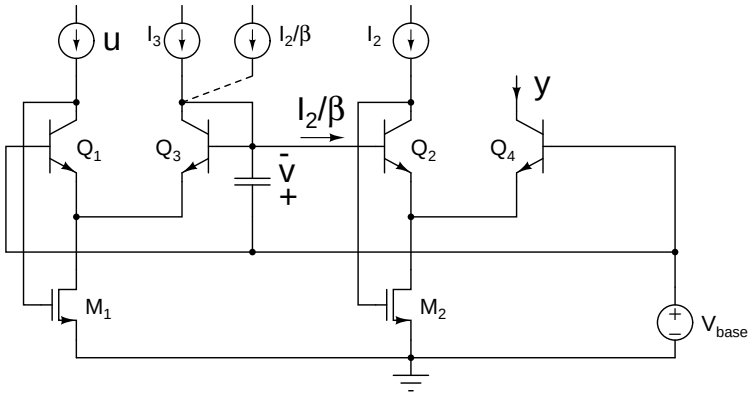

Figure 2.16: First-order filter in Fig. 2.14(e) operating from a positive power supply.

the operation of the circuit is that of Q_2. But Q_2's base current is a constant I_2/β which is subtracted from I_3. Thus no nonlinearity is introduced into the circuit and the only effect of the finite β (assuming β is constant w.r.t. the collector current) is a reduction of the bandwidth of the filter. If this reduction is not tolerable, it can be compensated either by injecting a current I_2/β as shown in the figure (used in [20]) or using automatic tuning techniques [35].

2.2.2 Syllabically-companding filters

A linear integrator, identical to the one in Fig. 2.5(a) is shown in Fig. 2.17(a). The state variable description of this system is given by (2.2) and (2.3). Let a new state variable v be related to the original state variable x by a multiplying factor g (g can be time varying in general) as follows

$$v = gx \tag{2.25}$$

The state equations of a new system with v as the state variable instead of x can be obtained by substituting $x = g^{-1}v$ in (2.2) and (2.3) [8, 3]

$$g^{-1}\frac{dv}{dt} - \frac{1}{g^2}\frac{dg}{dt}v = ku \tag{2.26}$$

$$y = g^{-1}v \tag{2.27}$$

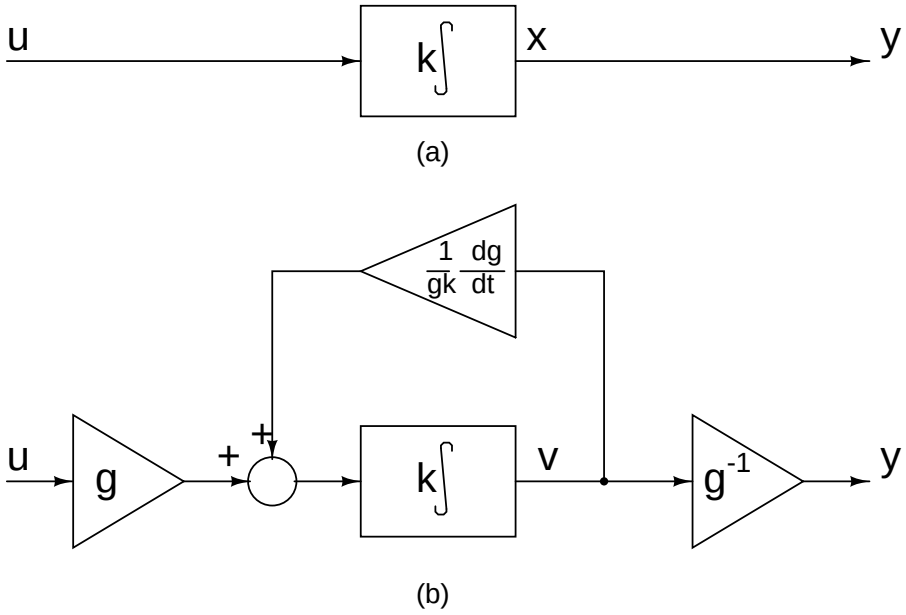

Figure 2.17: (a) Integrator, (b) Integrator with time varying gains at the input and the output.

Multiplying the first of the equations by g, we obtain

$$\frac{dv}{dt} = ku + \frac{1}{g}\frac{dg}{dt}v \tag{2.28}$$

$$y = g^{-1}v \tag{2.29}$$

Fig. 2.17(b) shows the realization of these equations in the form of a block diagram. Compared to Fig. 2.17(a), this has an extra feedback path which is in effect only when g is changing with time. This feedback provides the state variable compensation[4] mentioned in [5, 8].

More complicated syllabic companding filters can be synthesized by similarly transforming the state equations of the filters. A second order filter using this principle was presented in [9].

The integrator in Fig. 2.17(b) uses time varying gains g and g^{-1}, but has an LTI behavior between its input and output. To turn this into a companding integrator, g must be changed in accordance with the input signal strength so as to be small for large inputs and

[4]Modification of the state variable in order to maintain input-output linearity and time-invariance.

vice versa. In case of syllabic companding, the variation of g is much "slower" than the variations of the input signal u.

A syllabic companding filter needs a circuit that measures the average strength of the input to generate a signal that controls the value of g. This could be a rectifier followed by a filter, an envelope detector, or an rms detector. The biggest difficulty in realizing syllabic companding filters is in the state variable compensation circuit whose gain must be proportional to the time derivative of g. Direct differentiation of the gain control signal g is usually avoided to prevent excessive high frequency noise. But since the gain control signal is generated by a filtering process to average the input signal strength, direct differentiation is unnecessary. Fig. 2.18(a) shows a first-order filter that was used in [9] used to generate a gain control current I_g from the rectified input signal I_R. The signal at the input of the integrator in the filter is naturally the time derivative of I_g. I_g and its time derivative are fed to a divider to extract the desired control signal for the feedback gain. While this is simpler than generating the time derivative explicitly, a divider is still required.

A log-domain filter to which syllabic companding is applied using variable gain amplifiers at the input and the output is presented in [17]. A log-domain circuit to generate the state variable compensation circuit is used in this work. Fig. 2.18(b) shows a first-order log-domain filter that is used to generate the gain control signal I_g from the rectified input I_R. The current flowing through the capacitor is given by

$$i_c = C\frac{dv}{dt} \tag{2.30}$$

$$= C\frac{d}{dt}\left(V_t \log \frac{I_g}{I_s}\right) \tag{2.31}$$

$$= \frac{CV_t}{I_g}\frac{dI_g}{dt} \tag{2.32}$$

Which is precisely of the form required for the gain of the feedback path in Fig. 2.17(b). This

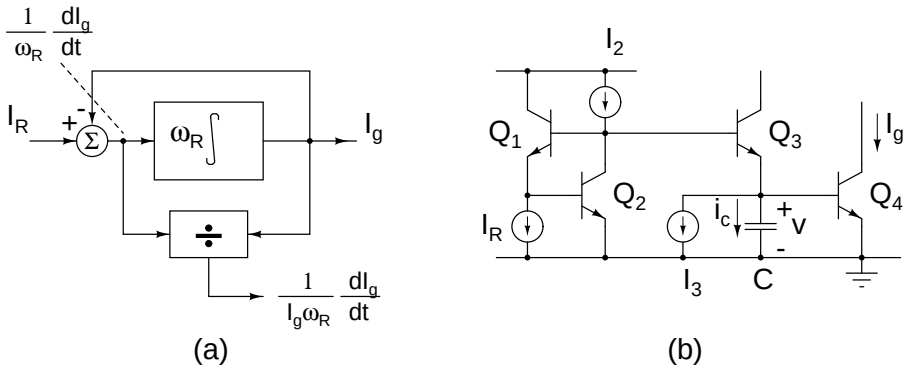

Figure 2.18: Generating the time derivative of the gain control current I_g.

current can be extracted from the collector current i_3 of the transistor Q_3 after subtracting a bias I_3.

Fig. 2.19 shows the block diagram of a syllabic companding filter in a general form [3]. A filter is enclosed between variable gain amplifiers whose gains are controllable. The input signal strength[5] is used to generate an appropriate control signal g which is used to control the gains of the amplifiers at the input and the output. To maintain LTI behavior, a state variable compensation circuit is used. In Fig. 2.17(b), the feedback path performed this function. In general, a state variable compensation circuit consists of amplifiers whose gains are proportional to the time derivative of g, divided by g.

Practical realization of syllabically-companding filters

As mentioned in sections 1.1 and sec 2.1.3, companding attempts to enhance the dynamic range of filters by ensuring that the internal signal (between the input and output amplifiers; e.g. gu in Fig. 2.19) is as large as possible even for small input signals. Ideally this means that g must be varied in inverse proportion to the input signal strength. This in turn implies that if a large dynamic range of input signals needs to be handled, g must vary over a very large range as well. Therefore the input amplifier is a multiplier which must be linear and be able to handle high dynamic range signals at both its inputs. The

[5] In general, signal strength anywhere in the filtering chain could be used to generate the control signal g.

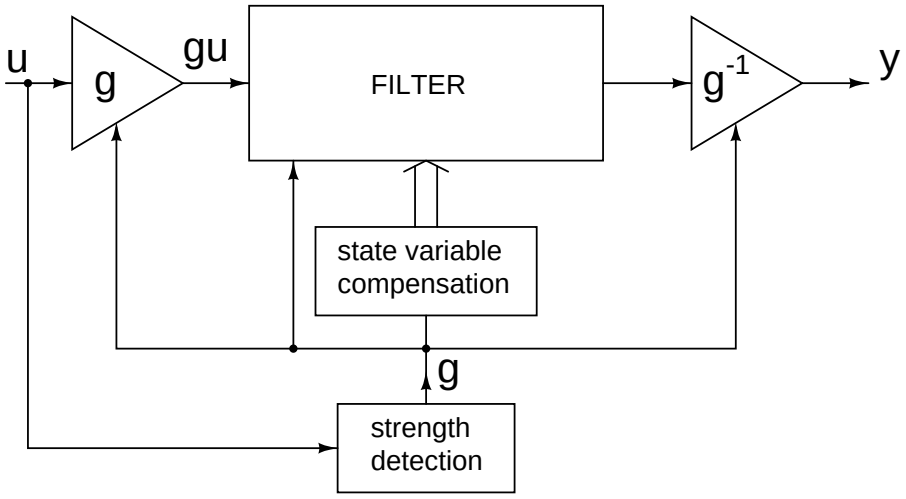

Figure 2.19: Block diagram of a general syllabic companding filter.

design of such an amplifier is by no means trivial. This fact combined with the complexities involved in the design of the state variable compensation circuit make it very difficult to realize a high dynamic range syllabic companding filter based directly on the prototype of Fig. 2.17(b). Work in this direction is under progress by others in the field.

2.2.3 Syllabically-companding log-domain filters using dynamic biasing

Fig. 2.20 shows the first-order log-domain filter in Fig. 2.7(c) with minor modifications. The dc bias and signal currents at the input and the output are shown separately. The currents i_1 and i_4 respectively in the input transistor Q_1 and the output transistor Q_4 are related by the transfer function in (2.20). This transfer function, with some terms rearranged can be expressed as

$$\frac{I_4(s)}{I_1(s)} = \frac{I_2}{I_3}\frac{1}{1+sCV_t/I_3} \tag{2.33}$$

The dc gain is I_2/I_3 and the pole is located at $a = I_3/CV_t$. The current i_1 in the transistor Q_1 is the sum of the signal input u and a bias I_{bias}. The current i_4 in the transistor Q_4 under quiescent conditions ($u = 0$) is $(I_2/I_3)I_{bias}$. The signal output y is obtained by subtracting this quiescent output current from i_4. The transfer function between the signals u and y is

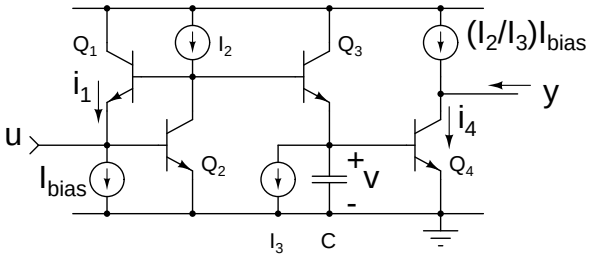

Figure 2.20: First-order filter in Fig. 2.7 with bias currents shown separately at the input and the output.

the same as (2.33) and is not affected by the bias current I_{bias}.

$$\frac{Y(s)}{U(s)} = \frac{I_2}{I_3}\frac{1}{1+sCV_t/I_3} \tag{2.34}$$

The output noise of the filter, along with the largest signal that can be handled, determines the dynamic range of this class-A filter. Several papers: e.g. [32, 36, 33] have been published outlining the methods for calculation of noise in such a filter. Chapter 6 describes in detail the steps involved in calculating the noise. It is also known [32, 33] that in class-A operation, the output noise does not change appreciably from the quiescent value when an input signal is applied. Therefore it is sufficient to consider the quiescent output noise in discussing the signal to noise ratio or dynamic range of class-A log-domain filters. The mean square integrated output noise of the filter in Fig. 2.20 in the quiescent state can be shown to be[6]

$$\overline{i^2_{n,out}} = \frac{qa}{2}A_{dc}(A_{dc}+1)\left(I_{bias} + \frac{{I_{bias}}^2}{I_3}\right) \tag{2.35}$$

Where q is the electron charge and A_{dc} and a are the dc gain and the pole frequency (in rad/s) of the filter respectively. Apart from the dc gain A_{dc} and the bandwidth a of the filter, the output noise depends on the bias current I_{bias}. The output noise increases with

[6]This expression includes the shot noise from the transistors Q_{1-4}. The power spectral density of noise from Q_4, which is not bandlimited by the circuit is multiplied by the noise bandwidth (in Hz) of the circuit $0.25a$ to obtain the expression for $\overline{i^2_{n,out}}$.

I_{bias}.

Reducing I_{bias} results in a reduced output noise but leaves the transfer function of the filter unchanged. The bias current I_{bias} need only be large enough to accommodate the input signal u. Therefore, the signal to noise ratio for small inputs can be improved considerably by decreasing I_{bias}. The effect of varying the bias I_{bias} on the dynamic range of the filter is similar to that of companding. Also reducing I_{bias} reduces the power consumption of the filter. However, if the bias I_{bias} is time varying, the output contains the filtered version of this time varying bias in addition to the desired output. Some manner of compensation is required in order to obtain only the filtered version of the input signal u at the output.

The solution offered in [18] to this problem is loosely based on the state variable compensation technique discussed in section 2.2.2. It involves injecting a state variable compensating current into the capacitor of the log-domain filter. Fig. 2.21(a) shows the circuit in Fig. 2.20 with the compensating current i_x included and a slightly modified output bias mI_{bias} subtracted from i_4. m is a parameter that is larger than the dc gain I_2/I_3 and can be chosen for convenience in the final circuit design.

The equations governing the new circuit in Fig. 2.21(a) contain i_x, I_{bias}, and the time derivative of I_{bias}. Imposing the condition that the new equation relating u and y be the same as that with a constant bias results in a particular form for i_x. The equation for i_x can then be used to derive the circuit that generates it. In line with the principle enunciated in Fig. 2.18, it is also assumed that I_{bias} itself is the output of a low pass filter so that the time derivative of I_{bias} does not have to be generated explicitly.

The log-domain filter in Fig. 2.20 with a constant bias is governed by the following

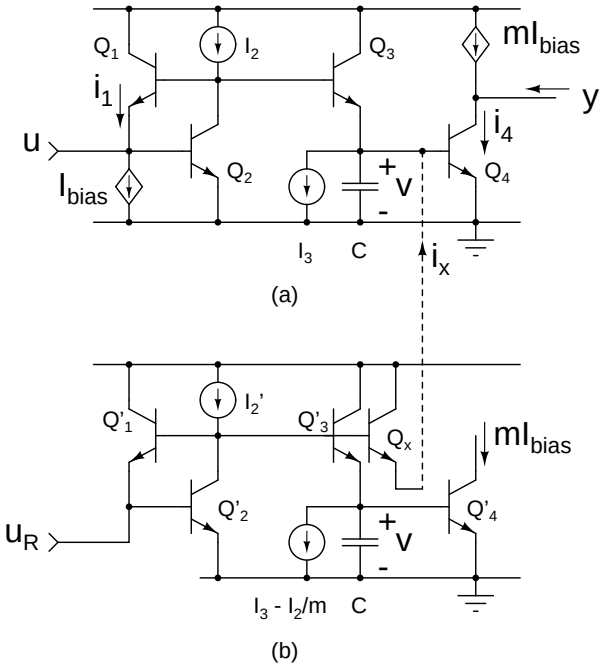

Figure 2.21: (a) First order log-domain filter with a current i_x injected into the capacitor, (b) Circuit for generating i_x and the filter used to average I_R.

equation:

$$CV_t\frac{dy}{dt} = -I_3 y + I_2 u \tag{2.36}$$

In the translinear loop Q_{1-4} in Fig. 2.21(a) the product of currents in the transistors Q_{1-2} is the same as the product of currents in the transistors Q_{3-4}. Thus, the equations governing this circuit in presence of a time varying I_{bias} are

$$\begin{aligned}(u + I_{bias})I_2 &= \left(C\frac{dv}{dt} + I_3 - i_x\right) i_4 \\ &= \left(\frac{CV_t}{i_4}\frac{di_4}{dt} + I_3 - i_x\right) i_4 \end{aligned} \tag{2.37}$$

$$i_4 = y + mI_{bias} \tag{2.38}$$

where the exponential relation between v and i_4 is used to express the relation entirely in terms of currents. Using (2.38), (2.37) can be rewritten as

$$CV_t\frac{dy}{dt} = -I_3 y + I_2 u + \left[i_x(y + mI_{bias}) - (mI_3 - I_2)I_{bias} - mCV_t\frac{dI_{bias}}{dt}\right] \tag{2.39}$$

If the term enclosed in square brackets on the right hand side is zero, (2.39) is identical to (2.36). For this to happen, i_x should be of the following form:

$$\begin{aligned} i_x &= \frac{(mI_3 - I_2)I_{bias} + mCV_t\frac{dI_{bias}}{dt}}{(y + mI_{bias})} \\ &= \frac{mI_{bias}(I_3 - \frac{I_2}{m} + \frac{CV_t}{I_{bias}}\frac{dI_{bias}}{dt})}{y + mI_{bias}} \end{aligned} \tag{2.40}$$

This equation containing products and ratios of currents can be implemented using a translinear loop. The principle in Fig. 2.18(b) can be used to generate the term containing the time derivative of I_{bias}. The resulting circuit is shown in Fig. 2.21(b). The transistors Q_{1-4} form the log domain filter that generates the bias I_{bias} from u_R which could, for instance, be the rectified version of the input u. Q_x generates i_x derived above. I_2' is adjusted to obtain the correct scaling factor between the peak value of u and the bias I_{bias}.

Simulation results

The dynamically biased filter in Fig. 2.21 was simulated to verify its operation. $I_2 = I_3 = 1.61\,\mu$A, $C = 100\,$pF and $m = 1.11$ were used in Fig. 2.21 to obtain a cutoff frequency of 100 kHz for the main filter 10 kHz for the auxiliary filter that generates I_{bias}. In these simulations, the envelope was assumed to be extracted using an ideal envelope detector. Consequently, u_R was not the rectified input, but the envelope itself. To evaluate the advantages of dynamic biasing, a conventional log-domain filter with a constant bias (Fig. 2.20) which is large enough to accommodate the largest input was also simulated. Identical inputs were fed to the two filters. These simulations are discussed in greater detail in [30].

Fig. 2.22(a) shows the input u with a changing envelope and a corresponding u_R (slightly larger than the input envelope) fed to the auxiliary filter. The *total* current i_4 in the output transistors of the filters in Fig. 2.20 and Fig. 2.21 are shown in Fig. 2.22(b) and Fig. 2.22(c) respectively. i_4 is small for small signals in the latter case reducing the transconductance of

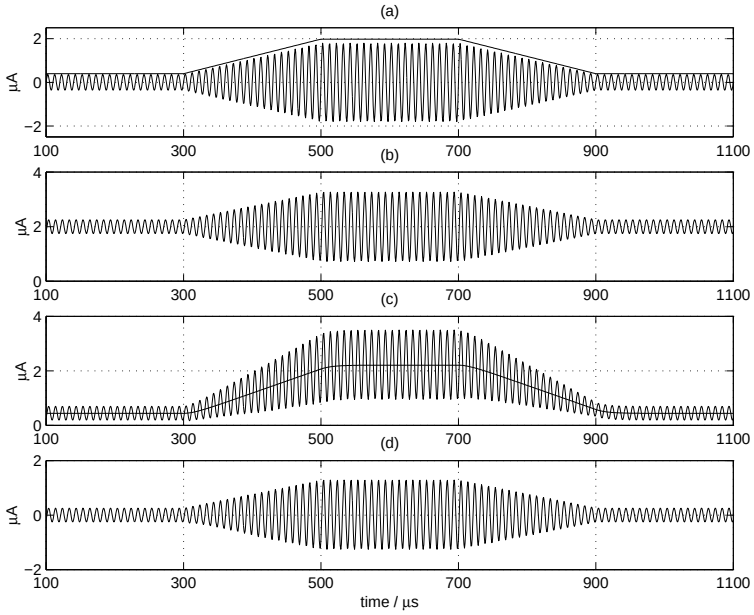

Figure 2.22: Simulation of the filter in Fig. 2.21: (a) Input u with a changing envelope and the corresponding u_R, (b) Output transistor current in the filter with a constant bias, (c) Output transistor current and the output bias current mI_{bias} in the filter with a dynamic bias, (d) Output signal y of the filter for cases (b) and (c).

Q_4 in Fig. 2.21. This reduces the gain for noise from internal points to the output and also reduces the noise from the output transistor Q_4. The current mI_{bias} that is subtracted from i_4 in Fig. 2.21 is also shown in Fig. 2.22(c). The output signal y obtained by subtracting the bias in the two filters is shown in Fig. 2.22(d). As expected, they are identical.

The output noise power spectral density, obtained using transient noise simulation[7] is shown in Fig. 2.23 with the input amplitude as a parameter. Two versions of the filter in Fig. 2.21, with and without noise are simulated with a sinusoidal input u at 100 kHz and a bias I_{bias} that is 10% larger than the peak value of u. The difference in their outputs taken to determine the output noise. The noisy version of the filter has uncorrelated noise sources that are connected between the collector and emitter of each transistor. The variation of the power spectral density of the output noise with the input signal (due to a proportional change in the bias current) is easily seen. Due to class-A operation, the results from the

[7]Chapter 6 deals with this technique in greater detail.

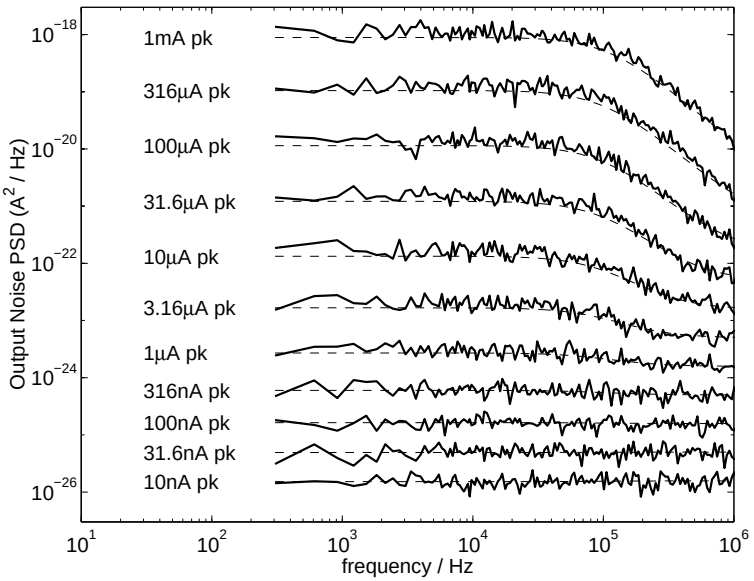

Figure 2.23: Noise PSD at the output of the filter in Fig. 2.21: solid line—from transient analysis, dashed line—from ".AC" analysis.

large signal transient analysis and a small signal ".AC" analysis at the quiescent point of the filter are very close to each other.

The rms integrated noise (from 0 to 2.5 MHz) is plotted versus the input current in Fig. 2.24. The output noise increases with the input signal. Two distinct slopes (on a logarithmic scale) of increase can be identified in this figure. At very small input currents, the rms integrated output noise increases as the square root of the input current. At very large input currents, the rms integrated output noise increases in proportion to the signal. This behavior is consistent with (2.35) whose right hand side has terms containing I_{bias} and ${I_{bias}}^2$ ((2.35) is an expression for the *mean square* integrated output noise). At very small currents, the transconductance of the output transistor Q_4, and hence, the amount noise form the internal circuits that appears at the output, is very small. The output noise is dominated by the shot noise of the input and output transistors (Q_1 and Q_4) which increases as the square root of the output current. At very large currents, the noise from the input and output transistors is insignificant. Most of the noise is from the internal nodes

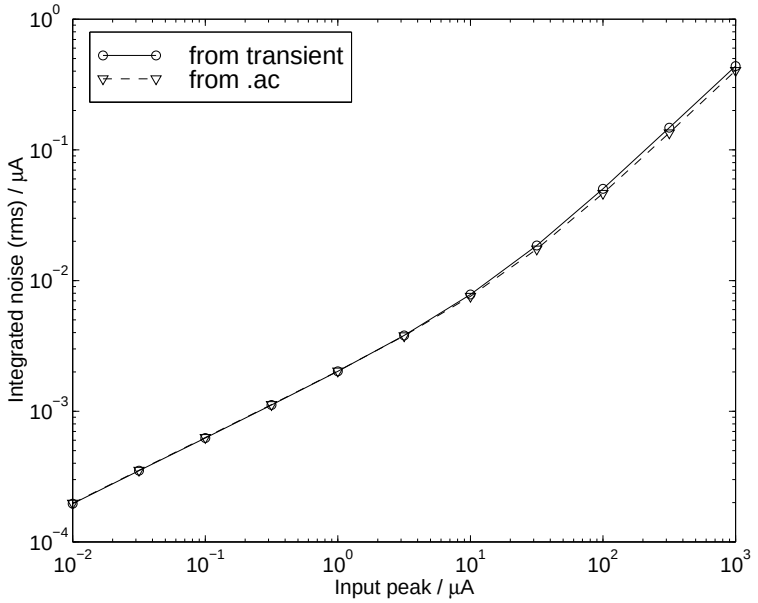

Figure 2.24: Integrated noise at the output of the filter in Fig. 2.21 with the input amplitude as a parameter: solid line—from transient analysis, dashed line—from ".AC" analysis.

converted into an output current by the transconductance of Q_4. Since the transconductance of the bipolar transistor Q_4 is proportional to its collector current, the output noise is proportional to the output current.

Fig. 2.25 shows the output signal to noise ratio of (i) the dynamically biased filter in Fig. 2.21 and (ii) a filter with a constant bias in Fig. 2.20. For the latter, the bias current is set to 10% larger than the highest value of the input peak. At small currents, when the output noise of the dynamically biased filter has a square root dependence on the input current, the signal to noise ratio increases by 0.5 dB for every dB increase in the input signal. At large currents, the output noise is proportional to the input current resulting in a constant signal to noise ratio. The signal to noise ratio of the filter with a constant bias is much smaller than that of the dynamically biased filter for the most part. This is because, with a constant (large) bias I_{bias}, the output noise has a constant large value. At the largest input signal, the dynamically biased filter has a slightly smaller signal to noise ratio due to extra

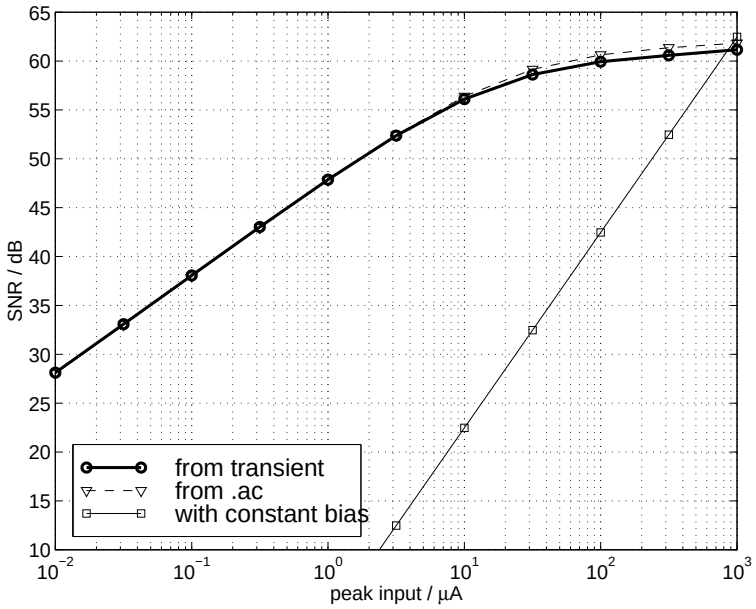

Figure 2.25: Signal to noise ratio at the output of the filter in Fig. 2.21.

noise added from the state variable compensation circuit. This plot shows the advantage of the companding systems very clearly. If a S/N_{min} of 40 dB is desired, the dynamically biased filter can accept signals that are five hundred times smaller than that are acceptable with the constant biased filter.

Chapter 3

A Simplified Technique for Dynamic Biasing and Syllabic Companding in Log-Domain Filters

3.1 Introduction

This chapter deals with the implementation of syllabic companding log-domain filters through dynamic biasing. It is shown here that the large signal linearity of log-domain filters can be exploited to eliminate the state variable compensation circuit [37]. Consequently, this approach is much simpler than the previously known techniques for syllabic companding and dynamic biasing described in Chapter 2. The new technique is described in the following section. The proposed technique is compared to other approaches to companding in section 3.3.

3.2 Proposed technique

The first order log-domain filter from Fig. 2.20 is shown in Fig. 3.1(a). As mentioned in the previous chapter, a linear time-invariant (LTI) relation exists between the large signal currents i_1 and i_4 in the input and the output transistors. When I_{bias} is constant with time, the same relation holds between the input signal u and the output signal y. The transfer

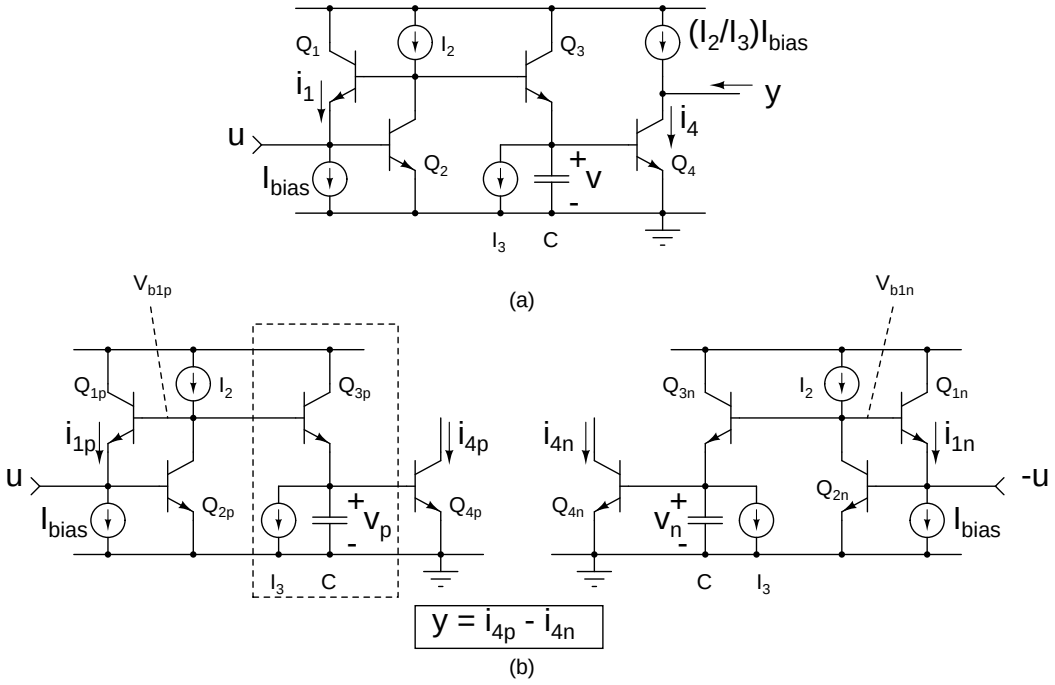

Figure 3.1: (a) First order log-domain filter, (b) Pseudo-differential operation with time varying bias.

function relating the two currents i_1 and i_4 and also u and y is ((2.33), (2.34))

$$\frac{Y(s)}{U(s)} = \frac{I_4(s)}{I_1(s)} = \frac{I_2}{I_3}\frac{1}{1+sCV_t/I_3} \tag{3.1}$$

In the time domain, the same relations can be expressed as

$$i_4(t) = i_1(t) * h(t) \tag{3.2}$$

$$y(t) = u(t) * h(t) \tag{3.3}$$

Where $h(t)$ is the impulse response of the filter between i_1 and i_4 or equivalently, the inverse Laplace transform of the transfer function in (3.1) and $*$ denotes convolution.

Assume that the filter shown in Fig. 3.1(a) is duplicated as shown in Fig. 3.1(b). One of the two filters receives a signal u added to a bias I_{bias}. The other filter is operated with the same bias I_{bias} but an opposite input $-u$ as shown in the figure. The output transistor

currents i_{4p} and i_{4n} in the two halves of the filter can be written as

$$i_{4p}(t) = (u(t) + I_{bias}(t)) * h(t) \tag{3.4}$$

$$i_{4n}(t) = (-u(t) + I_{bias}(t)) * h(t) \tag{3.5}$$

The difference output $y = i_{4p} - i_{4n}$ is given by

$$y(t) = 2u(t) * h(t) \tag{3.6}$$

Thus, in the difference output, the bias dependent term $I_{bias}(t) * h(t)$ disappears. The relation between the input u and the output y is LTI and is the same as that given by (3.3) for the original log-domain filter in Fig. 3.1(a) operating with a constant bias except for a factor of two which appears due to differential operation. In contrast with the technique described in the previous chapter (section 2.2.3, [18]), no extra circuit is required for state variable compensation here.

Deriving this result using the approach in [38] would require an internal description of a log domain filter, and would be more involved.

As described in section 2.2.1, the quiescent currents in the transistors of a log-domain filter are established either by adding a bias to the input or by using auxiliary bias sources. The relationship from the input and the auxiliary bias sources to the output is linear and time invariant. Therefore, the method described above for canceling the bias components at the output is applicable generally to all log-domain filters. Although the specific examples cited in this work relate only to log-domain filters, the proposed technique also applies to other filters in which the input bias (current or voltage) controls the internal biasing.

Two log-domain filters that are fed with identical bias currents and input signals of opposite polarity realize a pseudo-differential filter. The bias sources can now be varied in accordance with the signal strength to achieve dynamic biasing. The differential output

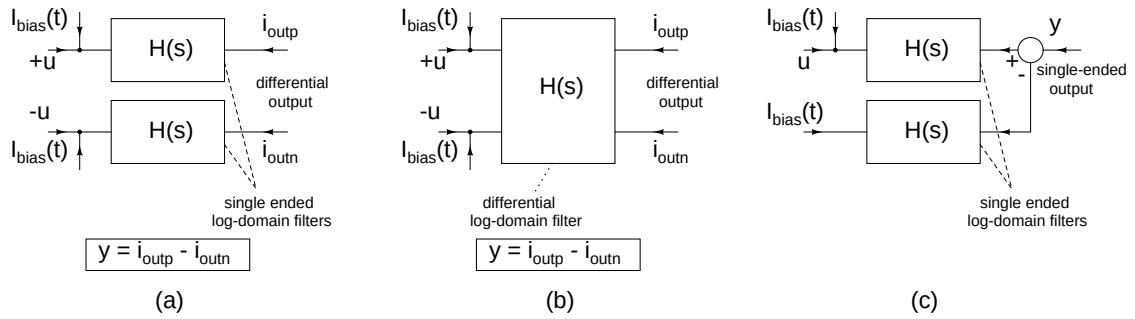

Figure 3.2: Distortionless dynamic biasing: (a) Pseudo-differential operation, (b) Using differential (class-AB) filters, (c) Single ended operation.

contains no component due to the time varying bias. This general scheme is shown in Fig. 3.2(a).

There are log-domain filter structures that are intended for differential operation—e.g. the class-AB circuits in [6, 19]. Instead of operating in class-AB mode, they could also be operated with a dynamic bias at the input as shown in Fig. 3.2(b). Again, the differential output contains no component due to the time varying bias.

An additional benefit of the pseudo-differential or differential operation as described in the previous paragraphs is the cancellation of even order nonlinearities and common interferences to the two halves of the circuit.

There may be cases in which single ended input and output are desirable. This could, for example, be in order to interface conveniently with preceding and following circuits. In that case, the configuration shown in Fig. 3.2(c) could be used. The sum of the signal and bias is fed to one of the filters and the time varying bias alone is fed to the other. In the difference of the outputs of the two filters, the bias component cancels out.

At this point, we would like to distinguish the proposed technique for companding from the one presented in [39]. Pseudo-differential circuits are used in that reference to implement *instantaneously-companding* filters. Although the system in Fig. 3.1(a) may be topologically similar to the circuits presented in [39], the goal here is syllabic compand-

ing. The circuits in [39] share some of the disadvantages of instantaneously companding circuits pointed out in section 3.3.

Dynamic biasing as syllabic companding

It was mentioned in section 2.2.3 dynamic biasing in log-domain filters is analogous to syllabic companding. This fact can be confirmed by observing the internal voltage swings of the filter. The voltage V_{b1p} at the base of Q_{1p} in Fig. 3.1(b), which is the sum of the base emitter voltages of Q_{1p} and Q_{2p}, is given by:

$$V_{b1p} = V_t \ln\left(\frac{u + I_{bias}}{I_s}\right) + V_t \ln\left(\frac{I_2}{I_s}\right) \tag{3.7}$$

where V_t is the thermal voltage and I_s is the saturation current of the transistors. For dynamic biasing, I_{bias} is made proportional to the envelope of u. Let I_{bias0} be the bias current for a given input $u = u_0$. If u is increased by a factor α to αu_0, the bias current increases by the same factor to αI_{bias0}. The voltage at the base of the transistor in these two cases, denoted respectively by $V_{b1p,a}$ and $V_{b1p,b}$, are

$$V_{b1p,a} = V_t \ln\left(\frac{u_0 + I_{bias0}}{I_s}\right) + V_t \ln\left(\frac{I_2}{I_s}\right) \tag{3.8}$$

$$\begin{aligned} V_{b1p,b} &= V_t \ln\left(\frac{\alpha u_0 + \alpha I_{bias0}}{I_s}\right) + V_t \ln\left(\frac{I_2}{I_s}\right) \\ &= V_t \ln\left(\frac{u_0 + I_{bias0}}{I_s}\right) + V_t \ln\left(\frac{I_2}{I_s}\right) + V_t \ln(\alpha) \\ &= V_{b1p,a} + V_t \ln(\alpha) \end{aligned} \tag{3.9}$$

Therefore, when the input amplitude increases by α and dynamic bias is derived from the input envelope, V_{b1p} experiences no change in its ac component and only a dc shift $V_t \ln(\alpha)$.

Since i_{4p} is linearly related to i_{1p}, i_{4p} is increased by the same factor α. This means that v_p, the base emitter voltage of Q_{4p} too sees only a dc shift of $V_t \ln(\alpha)$. It can similarly be reasoned that V_{b1n} and v_n in Fig. 3.1(b) experience only a dc shift when the input is increased.

Thus the ac component of the voltage signals V_{b1p} or v_p remains the same, provided the bias is increased in proportion to the signal. V_{b1p} and v_p are respectively the input and output voltages of the *voltage mode* filter enclosed by dashed lines. V_{b1p}, which is the voltage applied to this filter, has a constant amplitude regardless of the amplitude of the input u. This is precisely the syllabic companding action described in Chapter 2, where a filter is driven by a signal dependent gain stage, so as to always operate with a constant amplitude near the top of its dynamic range. The input transistor Q_{1p} and the output transistor Q_{4p} are analogous to the input and output amplifiers in Fig. 2.19 whose "gain" (current to voltage and voltage to current respectively) is modified by varying the bias current.

In general, in a higher order log-domain filter, all the internal voltages experience a dc shift provided all the bias and signal currents are increased in the same proportion.

Distortion in dynamically biased log-domain filters

Fig. 3.3 shows the general structure of a high-order dynamically-biased log-domain filter. Only one half of the filter is shown. For simplicity it is assumed that proper quiescent currents in all the transistors of the filter can be established using only the bias at the input rendering the auxiliary bias sources (see Fig. 2.11) described in section 2.2.1 unnecessary. But the following discussion holds even in presence of such auxiliary bias currents as long as they are changed in the same proportion as the input bias current. The filter consists of an input structure using Q_{1-2} that converts the sum of the input signal current and the bias current to a logarithmically compressed voltage denoted by v_{in}. v_{in} is fed to the core of the filter consisting of translinear loops and capacitors across certain nodes. The filter core's output voltage v_{out} is connected to an output transistor Q_4 realizing the exponential nonlinearity at the output.

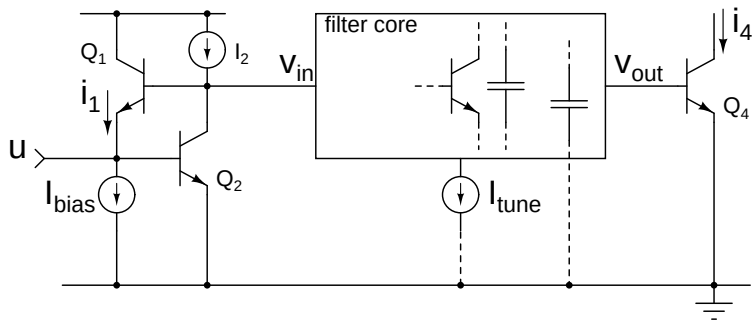

Figure 3.3: Generic structure of a high order log-domain filter

Some of the capacitors in the filter's core may be connected to ground. The only other connection to ground is through current sources used to control the time constants of the circuit. These are labeled I_{tune} in Fig. 3.3 to distinguish them from the "bias" currents used at the input. It is assumed that the current sources have infinite output resistances. Under these conditions, it is clear that the core of the filter is "floating" at dc[1]. This implies that a dc shift in v_{in} causes an equal dc shift in all the node voltages in the core of the filter and in the voltage v_{out}. The voltages across any two nodes in the circuit remains unchanged. As described earlier, in a dynamically biased filter, an increase in the input causes only a dc shift in the voltage v_{in} applied to the filter's core. Thus the filter's core, which is floating for dc, sees the same ac input regardless of the input amplitude, for a given input waveform. Therefore, the distortion contributed by the filter's core does not change with the input amplitude in a dynamically biased log-domain filter. If the input and the output stages do not add any distortion, a dynamically biased filter has a distortion performance that is independent of the signal strength for a given input waveform. The filter can thus be designed to have a constant "performance" over a wide dynamic range of inputs.

In practice, the input and the output stages do add distortion. But due to their simplicity, these stages are far easier to optimize than the large dynamic range linear amplifiers

[1]Other topologies of log-domain filters, for example, the one in Fig. 2.16, also possess this property.

required in the classical syllabic companding filter of Fig. 2.19. The finite output resistance of the current sources I_{tune} also causes some deviation from the ideal behavior described above.

3.2.1 Simulation results

HSPICE simulations were carried out to test the validity of the dynamic biasing scheme described above. Bipolar transistors with $\beta_F = 50$, $V_{AF} = 10\,\text{V}$ were used in the simulation. The pseudo differential filter in Fig. 3.1(b) was used with $I_2 = I_3 = 1\,\mu\text{A}$ and $C = 61.5$ pF to realize a low pass filter with a dc gain of unity and a -3 dB frequency of 100 kHz.

A sinewave at 100 kHz with a changing envelope as shown in Fig. 3.4(a) was fed to the filter. Having the input at the -3 dB frequency as opposed to deep in the passband ensures that substantial currents flow through the capacitor and that the circuit truly behaves like a dynamic nonlinear circuit as opposed to a static nonlinear circuit. Two cases were simulated. In the first case, a dynamic bias current I_{bias} whose value was 10% larger than the envelope of the input signal was used. This corresponds to dynamic biasing as described in Section 3.2. The envelope was assumed to be generated from an ideal envelope detector. In the second case, a constant bias which was 10% larger than the largest input envelope ($= 2\,\mu\text{A}$ in Fig. 3.4(a)) was used. This case corresponds to classical class-A operation.

The differential outputs in the two cases are overlaid in Fig. 3.4(b). As expected, they are exactly identical.

Fig. 3.4(c, d) show the base emitter voltage of Q_{4p} in the two cases. From Fig. 3.4(c) it is clear that the change in input amplitude causes only a dc shift in the internal voltage and that the internal peak-peak swing is being held constant. This confirms the syllabic companding action in the dynamically biased filter. With a constant bias, as can be seen

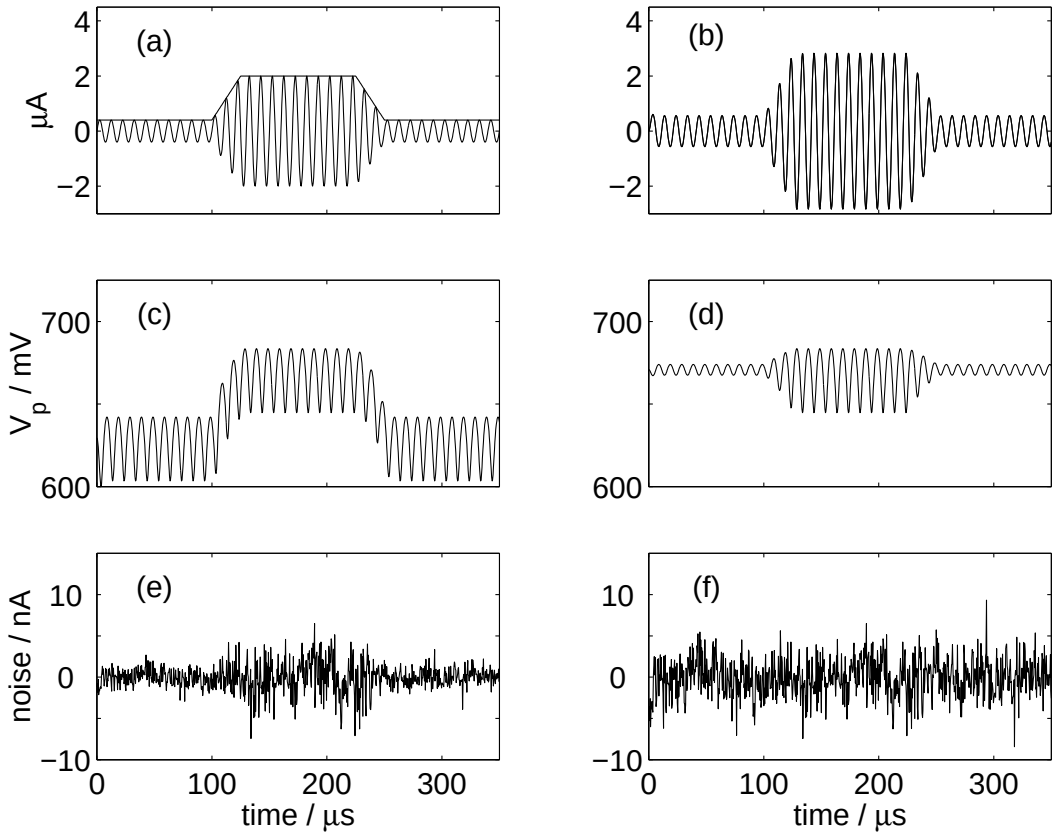

Figure 3.4: Results of simulation of the circuit in Fig. 3.1(b): (a) Input u and its envelope, (b) Differential output, (c) v_p with dynamic bias, (d) v_p with a constant bias, (e) Output noise with a dynamic bias, (f) Output noise with a constant bias.

in Fig. 3.4(d), the internal voltage swing varies in accordance with the input amplitude. If equal noise is added to these two waveforms, it is clear that the one in Fig. 3.4(d) would be corrupted much less than the one in Fig. 3.4(c).

The results of transient noise simulation[2] of the two cases using HSPICE is shown in Fig. 3.4(e, f). Two versions of the filter, with and without noise, are simulated and the difference in their outputs taken to determine the output noise. The noisy version of the filter has uncorrelated noise current sources that are connected between the collector and emitter of each transistor. The variance of each noise source is made proportional to the corresponding collector current. The reduction in output noise current for small signals in the dynamically biased filter (Fig. 3.4(e)) when compared to the constant biased filter (Fig. 3.4(f)) is evident.

These results prove the external linearity and syllabic companding nature of the proposed dynamically biased filter. For visual clarity, the envelope in these simulations is changed only by a small factor of four. The change in the envelope, and consequent savings in noise at small currents, can be much larger in real situations. It is limited mainly by the range of currents over which the input and output devices behave as exponential devices.

3.2.2 The dynamic biasing signal

Several factors that need to be considered while choosing the dynamic bias that is added to a given ac input signal are listed below.

1. In practice, the bipolar transistors in a log-domain filter need to be operated above a certain minimum collector current in order for them to have sufficient bandwidth [18, 7]. Therefore, the dynamic bias should not be decreased below a predetermined

[2] Chapter 6 deals with this technique in greater detail.

minimum level. Above this level, it can be varied in proportion to the input signal.

2. It must be ensured that the currents in *all* the transistors of the log-domain filter stay positive and above the minimum value required for sufficient bandwidth at all times. The general technique [30] for ensuring adequate quiescent currents in all the transistors of a log-domain filter was outlined in section 2.2.1 and Fig. 2.11.

3. The amount by which the bias I_{bias} is above the peak value of the input signal is determined by distortion considerations.

Issues related to the generation of the dynamic biasing signal are discussed in the next chapter.

3.3 Comparison to other companding filter realizations

3.3.1 Syllabic companding

The difficulties in implementing a syllabic companding filter using the general ideas of [3] (Fig. 2.19) were discussed in section 2.2.2. The state variable compensation circuit and the high dynamic range amplifiers at the input and the output can significantly increase the design complexity of the syllabic companding filter.

The dynamic biasing scheme proposed in [18, 30] (and described here in section 2.2.3) simplifies the implementation of syllabic companding log-domain filters. The amplifiers at the input and the output are eliminated. But the state variable compensation circuit is still required and necessitates extra design effort. Its complexity depends on the main filter being implemented. On the other hand, the technique proposed here needs no extra effort, since the required filter is merely duplicated.

When a circuit is duplicated and operated pseudo-differentially, its power consumption and signal to noise ratio (in terms of mean-squared quantities) are doubled. The latter

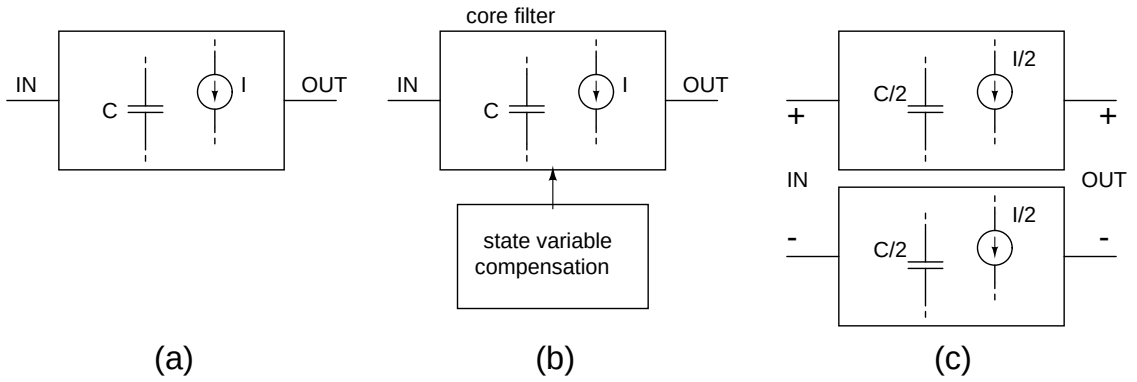

Figure 3.5: (a) Class-A single ended log-domain filter, (b) Dynamic biasing as in section 2.2.3, (c) Dynamic biasing as in section 3.2.

is because the *amplitude* of the signal is doubled from its single ended value (a 6 dB increase) whereas the *mean square* value of noise is doubled[3] (a 3 dB increase). Thus, for a given power consumption, the pseudo-differential circuit has the same signal to noise ratio as the single ended circuit. If the values of the bias currents and the capacitors in each half of the pseudo-differential circuit are half the corresponding values in the single ended circuit, the single ended and the pseudo-differential versions have identical power consumption and signal to noise ratio.

The approaches to distortionless dynamic biasing in section 2.2.3 and section 3.2 can now be compared. Fig. 3.5(a) shows a single ended class-A log-domain filter with certain bias currents and capacitor values. Fig. 3.5(b) shows a dynamically biased version of this filter with a state variable compensation circuit as described in section 2.2.3. The core filter has the same bias currents and capacitors as before. Fig. 3.5(c) shows a pseudo-differential version of the class-A filter where the capacitors and bias currents in each section have been halved. It is clear from the above discussion that the power consumed and the noise contributed by the core filter in Fig. 3.5(b) and by the pseudo differential filter in Fig. 3.5(c) are identical. But there will be additional noise and power consumption in Fig. 3.5(b) due

[3]The noise from the two halves of the filter are uncorrelated.

to the state variable compensation circuit. It can thus be seen that the dynamically biased filter described in section 3.2 (Fig. 3.5(c)) is superior to the one in section 2.2.3 (Fig. 3.5(b)) in terms of signal to noise ratio per unit power consumption.

3.3.2 Instantaneous companding

High dynamic range log-domain filters can also be realized using class-B or class-AB ([6, 19]) operation. These filters have differential inputs whose difference equals the desired input signal. The individual inputs can be half-wave rectified waveforms, or waveforms whose geometric mean is held constant. Recently [39] proposed the use of another type of input waveforms—waveforms with their harmonic mean held constant—with a pseudo-differential filter structure for instantaneous companding.

In this author's opinion, syllabic companding using the dynamic biasing technique described in this chapter has several advantages over instantaneous companding techniques for enhancement of the dynamic range of log-domain filters. These are listed below.

1. The preprocessing circuit: The envelope detector's accuracy is relatively unimportant as long its output is larger than the actual envelope, whereas the class-AB splitter has to accurately reproduce the input signal in its difference output to avoid added distortion.

2. The envelope detector is simpler to design than a class-AB splitter, partly due to the reduced accuracy requirement.

3. Mismatch between two halves of the filter leads to distortion because of internal nonlinearity in a class-AB filter and incomplete cancelation of bias components in a dynamically biased filter. But slow varying bias components may be more acceptable than intermodulation distortion in many cases.

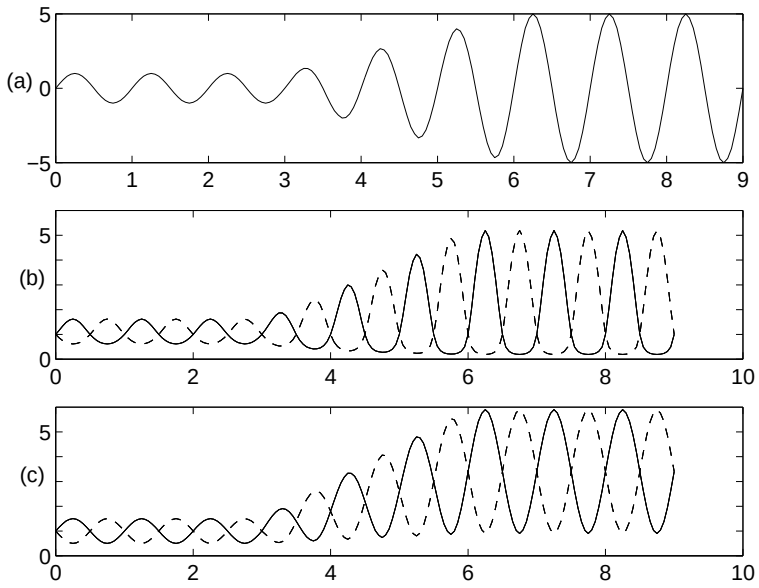

Figure 3.6: (a) Input with a changing envelope, (b) Geometrically split inputs for a class-AB filter, (c) Input with a dynamic bias added to it.

4. Fig. 3.6(a) shows a sinusoidal signal with a changing envelope. If this signal has to be fed to a class-AB filter, the signal pair shown in Fig. 3.6(b) must be generated. If it is to be fed to a dynamically biased filter, the signal pair in Fig. 3.6(c) must be used. For the smaller envelope, the inputs to the class-AB filter (Fig. 3.6(b)) is nearly identical to that to the dynamically biased filter (Fig. 3.6(c)). But for the larger envelope, the inputs to the class-AB filter reach near-zero values, whereas the inputs to the dynamically biased filter do not. The current swings over a much larger relative range in the former case. So even in absence of any mismatch, the class-AB configuration is more likely to distort than the dynamically biased class-A configuration.

5. The envelope detector's noise cancels at the output of the dynamically biased filter, whereas the noise in the class-AB splitter's outputs appears in opposite phases of the input signal for large signals and do not cancel. This may degrade the filter's SNR.

Chapter 4

Generation of the Control Signal for Dynamic Biasing and Syllabic Companding

Syllabic companding filters, of which the dynamically biased log-domain filters are a subset, need an average measure of the input signal strength to adopt themselves to the changing input signal. Generating a suitable control signal is non-trivial and requires careful consideration of the context in which the filter operates. This chapter briefly deals with the issues involved in the generation of the control signal for syllabic companding or dynamic biasing.

In many systems, an example of which is given in the following section, the envelope of the signal may be known in "advance" and no control generation circuitry is required. In cases where no such prior knowledge of the signal strength is available, a measure of the average strength must be derived from the signal itself. Such averaging introduces some inherent limitations in syllabic companding systems and they are discussed section 4.2. Section 4.3 deals with a few techniques for envelope detection that can be used to generate the control signal for dynamic biasing or syllabic companding.

4.1 Situations where envelope detection is unnecessary

There are many systems in which a complicated analog signal may be obtained by feeding a digital version of the signal to a D/A converter. The digital version of the envelope of the analog signal may already be present in the system or may be generated with the addition of a small amount of digital circuitry. An analog version of the envelope can be generated by using a D/A converter. As mentioned in the previous chapter, the dynamic bias added to the filter need not be very accurate provided it is larger that the signal peaks. Thus a fairly low resolution D/A converter can be used in its generation. Such "advance" knowledge of the envelope eliminates the need for envelope detection circuits and the problems associated with them.

Fig. 4.1(a) shows an amplitude shift keyed signal that is used in some forms of digital communication systems. The amplitude of a carrier is modulated by multi-level digital data[1]. The modulated signal can be generated digitally at a low frequency and converted to analog form using a D/A converter. In this case, a smoothing filter is required to remove the high frequency components of the sampled signal. If a dynamically biased filter is intended to be used for smoothing, the dynamic bias need not be generated from Fig. 4.1(a). Since the transmitted data is known, the envelope (Fig. 4.1(b)) is known beforehand. As mentioned earlier, a low resolution D/A converter can be used in this situation to generate the dynamic bias.

[1] Fig. 4.1(a) shows a single phase carrier. Two carriers in quadrature can be used with independent data modulating each of them—QAM.

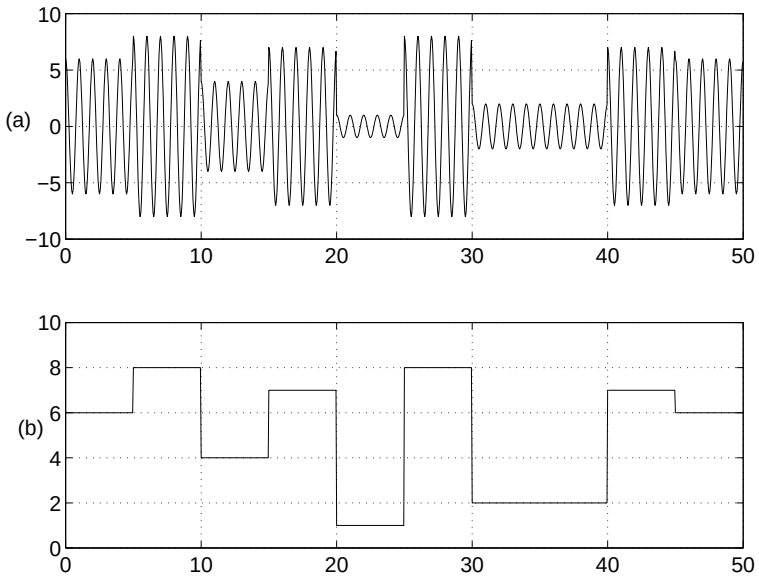

Figure 4.1: (a) Amplitude shift keyed waveform, (b) Envelope of (a).

4.2 Problems due to the averaging present in strength detection.

4.2.1 Syllabic companding systems

The syllabic companding system with memoryless channels shown in Fig. 1.1 is shown in Fig. 4.2 with the addition of an envelope detector and a circuit for generation of the gain control signal. The input and output of the input amplifier are denoted by u and $\hat{u}$ respectively. In general, as the signal strength v_a increases, the gain g of the input amplifier decreases. In this example, it is assumed that the gain g is proportional to the inverse square root[2] of the detected signal strength v_a. The input u is assumed to be a sinusoid with a changing envelope given by

$$u = a(t)\cos(\omega t) \tag{4.1}$$

In ideal conditions, $v_a = a$ so that the output $\hat{u}$ of the input amplifier is $\hat{u} = \sqrt{a}\cos(\omega t)$. It is assumed that the input amplifier saturates for $\hat{u}$ greater than 2.5. A noise n which has a

[2]This is an arbitrary choice used here for the purpose of illustration. Such a variation in gain compresses the dynamic range of the input signal (in dB) by factor of two and has been employed in [2, 40].

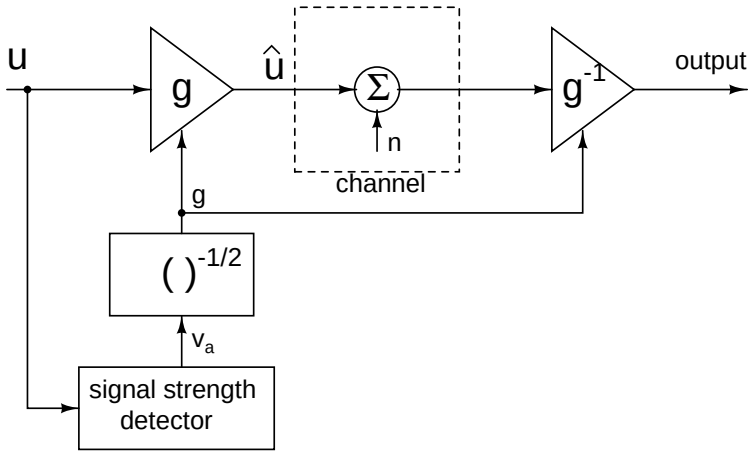

Figure 4.2: Companding with memoryless noisy channels.

constant variance over time is added to the signal $\hat{u}$ in the channel[3].

Fig. 4.3(a) shows a sinusoidal input u whose envelope is initially 0.5, changes to 4.0 at $t = 1$ s and returns to 0.5 at $t = 5$ s. Under ideal conditions, the detected envelope v_a and the gain control signal g are as shown in Fig. 4.3(b) and Fig. 4.3(c) respectively. Fig. 4.3(d) shows the ideal compressed signal $\hat{u}$. A factor of 8 variation in u is compressed to a factor of $2\sqrt{2}$ variation in $\hat{u}$. Fig. 4.3(e) shows the output signal of the system in Fig. 4.2 and it is identical to the input. Fig. 4.3(f) shows the noise n_o at the output of the system.

In practice, when the envelope of the input signal changes abruptly as shown in Fig. 4.3(a), it is impossible to obtain the signal that perfectly represents the envelope. Fig. 4.4 shows a more realistic situation. Fig. 4.4(a) shows the input signal u, which is the same as in Fig. 4.3(a). The signal v_a representing the envelope of the input and its inverse square root are shown in Figs. 4.4(b) and 4.4(c) respectively. The latter is used to control the gain of the amplifier in Fig. 4.2. During long periods of constant input envelope, the signal v_a in Fig. 4.4(b) faithfully represents the input envelope. But abrupt changes in the input envelope are not tracked. This behavior results if the average of the rectified input

[3]The parameters of the system are chosen such that the nonideal effects that we wish to highlight are clearly seen.

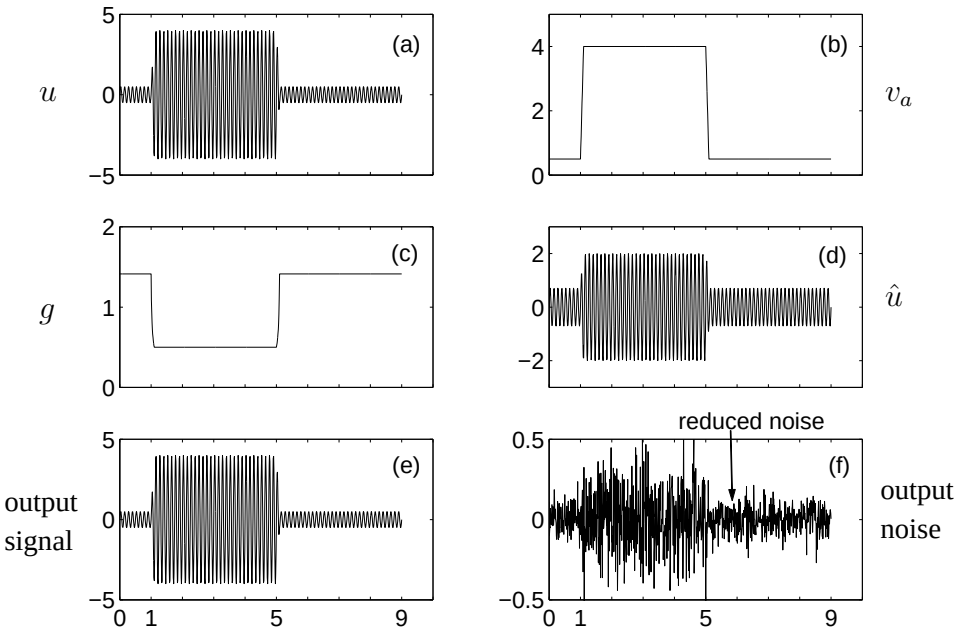

Figure 4.3: Syllabic companding under ideal conditions: (a) Input u with a changing envelope, (b) v_a, the detected envelope of u, (c) Gain g of the input amplifier, (d) Compressed signal $\hat{u}$, (e) Output signal of the system, (f) Output noise of the system.

signal (described in section 4.3.3) is used to estimate the input envelope. Fig. 4.4(d) shows the signal $\hat{u}$. There is an overshoot (i.e. insufficient compression) in $\hat{u}$ immediately following the increase in the input envelope at $t = 1\,\text{s}$. Due to the swing limit of the amplifier, the signal that overshoots is clipped at 2.5. This happens because the gain g is not adjusted immediately to the value that is appropriate for the increased input amplitude. The information in the signal is irretrievably lost in this period. Similarly, after the envelope decreases at $t = 5\,\text{s}$, the input is compressed "too much" before attaining the correct steady state. Fig. 4.4(e) shows the output signal of the system. Immediately following the envelope increase at $t = 1\,\text{s}$, the output signal is smaller than the ideal case shown in Fig. 4.4(e). This is due to the clipping observed in Fig. 4.4(d). The output noise is shown in Fig. 4.4(f). It is larger than in Fig. 4.3(f) just after the decrease in the input envelope at $t = 5\,\text{s}$. This is due to the slowly changing gain control signal.

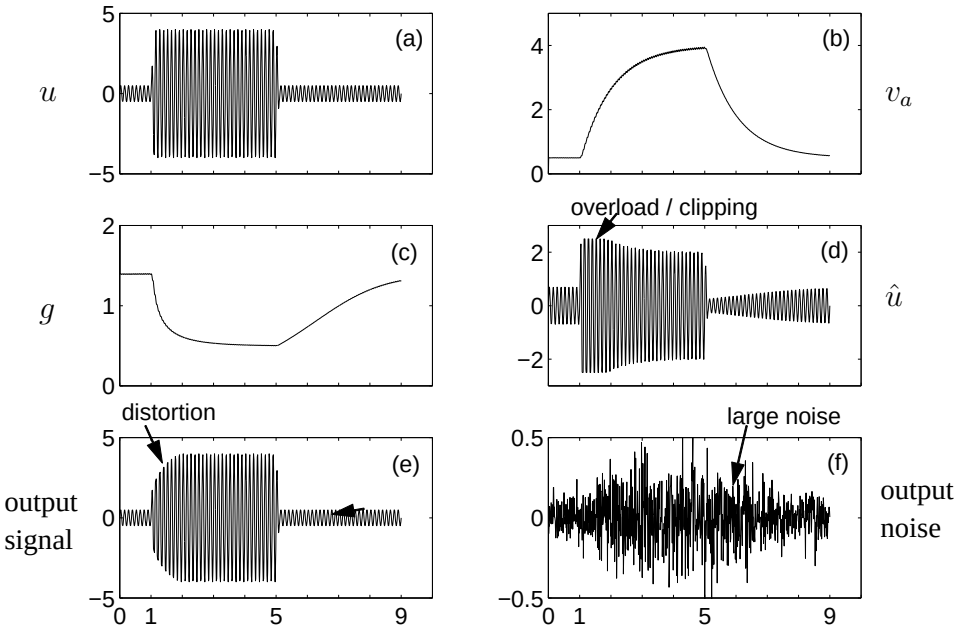

Figure 4.4: Syllabic companding under practical conditions: (a) Input u with a changing envelope, (b) v_a, the detected envelope of u, (c) Gain g of the input amplifier, (d) Compressed signal $\hat{u}$, (e) Output signal of the system, (f) Output noise of the system.

The two nonidealities described above, namely (i) Overload following an increase in the input envelope and (ii) Decreased signal to noise ratio following a decrease in the input envelope are characteristic of syllabic companding systems.

4.2.2 Syllabic companding filters

The implementation of syllabic companding *filters* (Fig. 4.5) poses an additional problem that is not encountered in syllabic companding systems with memoryless channels. When the input to the filter is a single sinusoid, the peak values of the signal at various points in the filter are related to the peak value of the input through the magnitude responses of the filter from the input to these points. This is no longer the case when a combination of frequencies is fed to the filter because different frequencies can experience different phase shifts in a filter (unlike the memoryless channel in Fig. 4.2). Fig. 4.6(a) shows the first five harmonics of a square wave at a frequency f_0. The sum of these waveforms is shown at

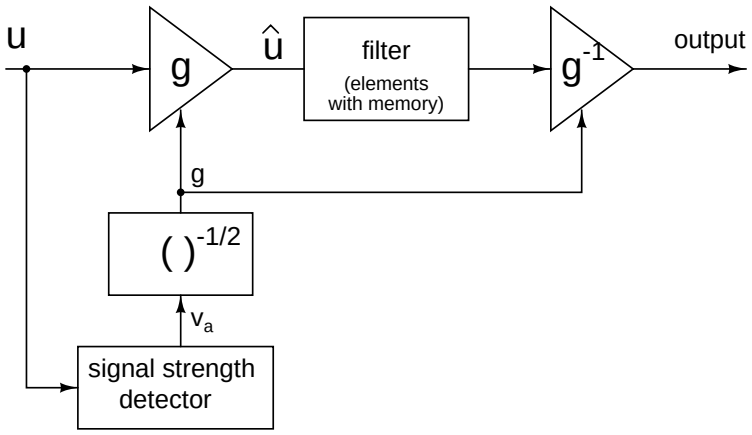

Figure 4.5: Companding filter with memory elements between the compressor and the expandor.

the bottom and resembles a square wave. Fig. 4.6(b) shows five harmonics which have the same amplitude as in Fig. 4.6(a). The third and the seventh harmonic are inverted. The sum of these is shown at the bottom [41]. It has a peak value nearly twice as large as that of the square wave in Fig. 4.6(a). While this is an extreme example, it illustrates that the peak values of the internal signals in the filter can be quite different from that of the input. Adjusting the value of the added dynamic bias based on the peak values of all the state variables of the filter can solve this problem. But this requires a peak detector for each of the state variables of the filter and is impractical. The dynamic bias added to the input must therefore be based on the transfer function of the filter from the input to the state variables and some assumptions regarding the nature of the input signal [30].

4.2.3 Dynamically biased filters

The dynamically-biased log-domain filter (Fig. 3.1(b)) described in the previous chapters shares the above described shortcomings inherent to syllabic companding systems. In case of dynamic biasing, instead of a multiplicative gain g, a bias is added to the input signal. The bias is derived from the average signal strength and cannot follow abrupt changes in

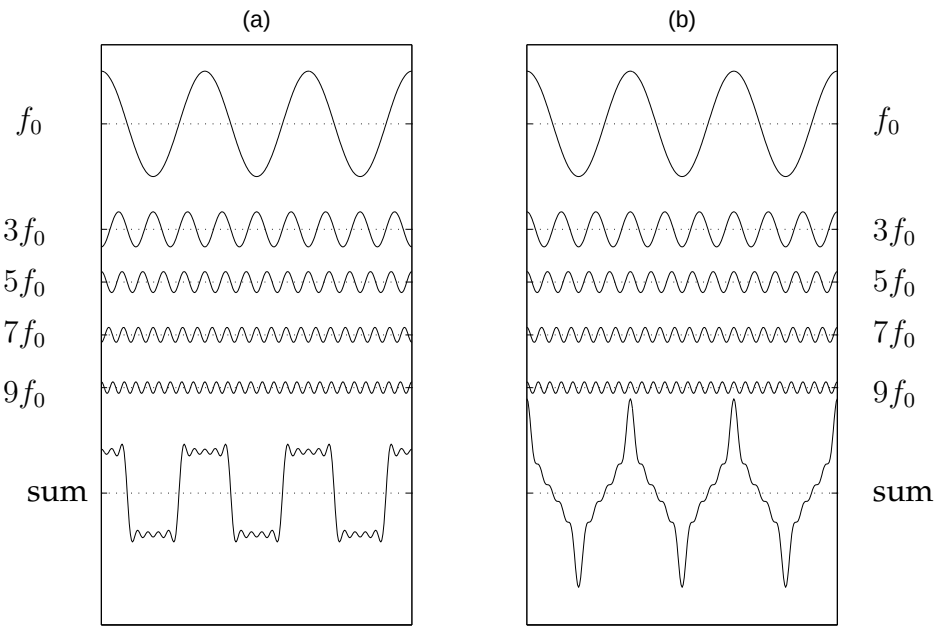

Figure 4.6: (a) First five harmonics of a square wave and their sum, (b) Same as (a), but with the third and seventh harmonic inverted.

the input envelope. This has the following consequences which are analogous to those depicted in Figs. 4.4(e, f).

1. Immediately after an abrupt increase in the input envelope, the added bias is insufficient to keep the overall input (i_{1p} in Fig. 3.1(b)) to the log-domain filter positive at all times. The negative peaks of the overall input are clipped by the input transistor (Q_1 in Fig. 3.1(b)), resulting in distortion.

2. Immediately after an abrupt decrease in the input envelope, the added bias is too large for the small input signal. This results in an excessive output noise (see (2.35)), compromising the signal to noise ratio of the signal.

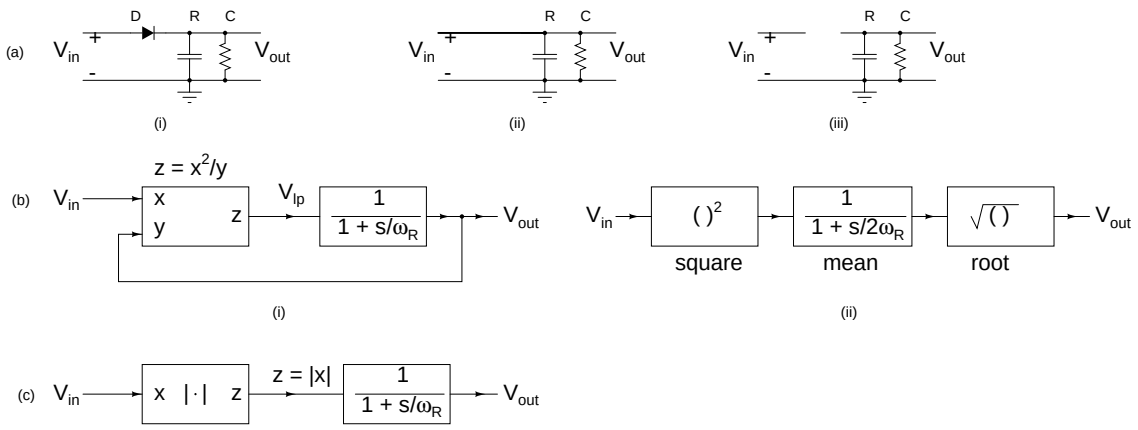

Figure 4.7: Techniques for envelope detection: (a) Classical diode-RC peak detector, (b) RMS detector, (c) Rectifier and averaging.

4.3 Estimating signal strength

4.3.1 Peak detector

Fig. 4.7(a-i) shows the classical peak detector using an ideal diode and D and an RC network. In this context, an ideal diode is one which has a zero voltage drop across it when a forward current is flowing through it and a zero current through it when a reverse voltage is applied across it. If the input voltage tends to be larger than the output voltage, the diode is turned on and the output voltage follows the input voltage (Fig. 4.7(a-ii)). When the input voltage V_{in} is less than the output voltage V_{out}, the diode is turned off (Fig. 4.7(a-iii)) and V_{out} decays exponentially due to the discharge of the capacitor C through the resistor R.

Fig. 4.8(a) shows the output of such a peak detector for a sinusoidal input of a constant amplitude. The output voltage (on the capacitor C) is recharged to the positive peaks through the diode and droops between successive peaks. The amount of droop is determined by the product of the time constant of the RC filter and the input frequency.

Fig. 4.9(a) shows the output of such a peak detector for a sinusoidal input with a changing envelope. The "attack"—the response to an increase in the input envelope—is

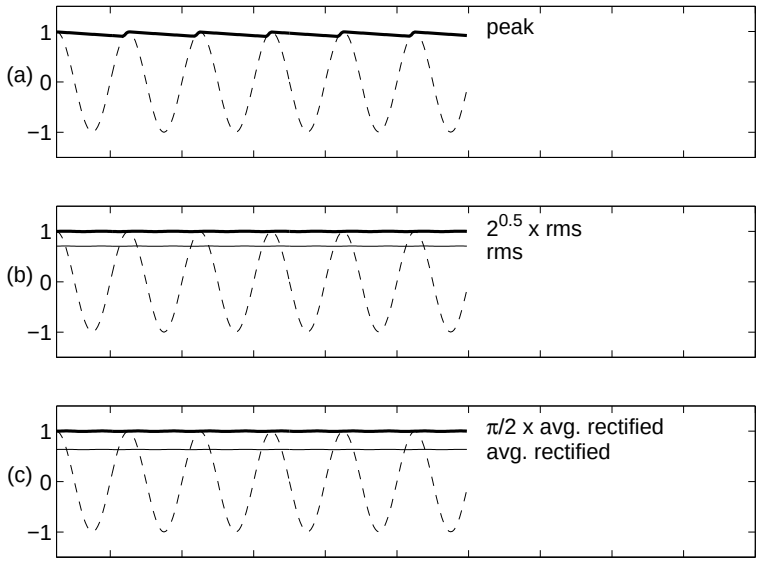

Figure 4.8: Steady state responses of (a) Classical diode-RC peak detector, (b) RMS detector, (c) Rectifier and averaging, to a sinusoidal input. The time constant used in the filters in Fig. 4.7(a-i), Fig. 4.7(b-i) and Fig. 4.7(c) is ten times as long as the period of the input.

fast since the capacitor C is charged rapidly through the diode D which is on (Fig. 4.7(a-ii)). The "decay" is slow—it is the exponentially decaying natural response of the RC combination (Fig. 4.7(a-iii)).

The peak detector responds to the peak value of the input regardless of the input waveform.

4.3.2 RMS detector

The envelope of the input signal can also be determined using an rms detector whose output is multiplied by the crest factor (The crest factor is the ratio of the peak value to the rms value of a signal).

Fig. 4.7(b-i) shows the well-known block diagram of an rms detector [42]. In steady state, the feedback network comprising the squarer-divider and the low-pass filter has an output V_{out} that is the root mean squared value of V_{in}.

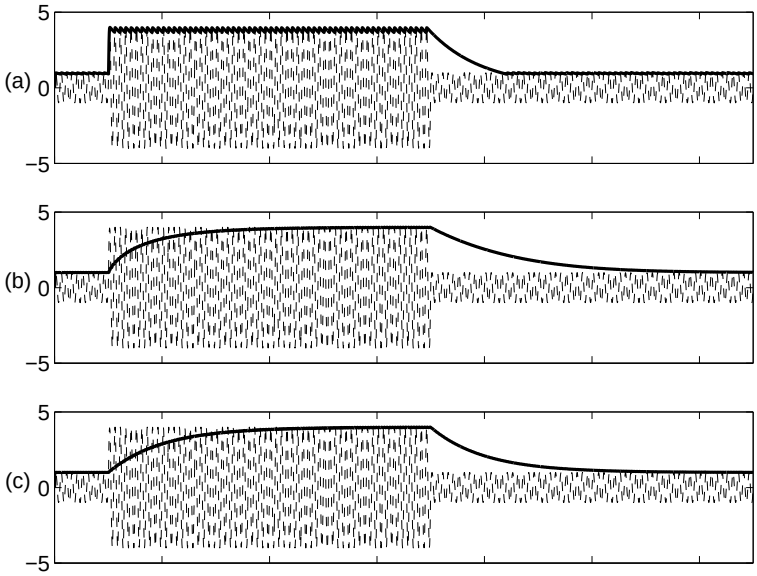

Figure 4.9: Responses of (a) Classical diode-RC peak detector, (b) RMS detector, (c) Rectifier and averaging, to a sinusoidal input with a changing amplitude. The time constant used in the filters in Fig. 4.7(a-i), Fig. 4.7(b-i) and Fig. 4.7(c) is ten times as long as the period of the input.

This can be understood as follows [43, 44]: Denoting the input to the low-pass filter in Fig. 4.7(b-i) by V_{lp} the equation governing the low pass filter can be written as follows

$$\frac{dV_{out}}{dt} = -\omega_R V_{out} + \omega_R V_{lp} \tag{4.2}$$

where ω_R is the cutoff frequency of the low pass filter. Substituting $V_{lp} = V_{in}^2/V_{out}$ (the squarer divider in Fig. 4.7(b-i)) in this equation we obtain

$$\frac{dV_{out}}{dt} = -\omega_R V_{out} + \omega_R \frac{V_{in}^2}{V_{out}} \tag{4.3}$$

$$\frac{dV_{out}^2}{dt} = -2\omega_R V_{out}^2 + 2\omega_R V_{in}^2 \tag{4.4}$$

where the second equation is obtained from the first by multiplying both sides with $2V_{out}$ and using the relation $2V_{out}dV_{out}/dt = dV_{out}^2/dt$. Fig. 4.7(b-ii) shows the equivalent system realizing (4.4). The output voltage V_{out} is the square root of the average (low pass filtered value) of the square—the root mean square (rms) value—of the input V_{in}.

Fig. 4.8(b) shows the output of the rms detector for a sinusoidal input of a constant amplitude. The output voltage is multiplied by[4] $\sqrt{2}$ to obtain the peak of the signal.

The output voltage has a ripple at twice the frequency of the sinusoid. This is the result of filtering V_{in}^2 (4.4). The amplitude of the ripple is determined by the time constant of the filter relative to the period of the input signal.

Fig. 4.9(b) shows the output of such a circuit for a sinusoidal input with a changing envelope. The attack and decay times are similar and are determined by the time constant of the filter used in the rms detector in Fig. 4.7(b).

Current mode rms detectors using translinear circuits (realizing the differential equation (4.4)) are presented in [43, 44].

4.3.3 Rectifier + averaging

Another measure of the input signal strength can be obtained using the combination of full wave rectification[5] and averaging shown in Fig. 4.7(c). Ideal diodes are assumed to be used. The full-wave rectified input is filtered and multiplied by an appropriate scale factor in order to obtain the peak value of the input.

Fig. 4.8(c) shows the output of Fig. 4.7(c) for a sinusoidal input of a constant amplitude. The output voltage is multiplied by[6] $\pi/2$ to obtain the peak of the signal.

The output voltage has a ripple at twice the frequency of the sinusoid. This is the result of filtering the full wave rectified input. The amplitude of the ripple is determined by the time constant of the filter relative to the period of the input signal.

Fig. 4.9(c) shows the output of such a circuit for a sinusoidal input with a changing envelope. As with the rms detector, the attack and decay times are similar and are

[4] This value is specific to a sinusoidal input. It depends on the shape of the signal.
[5] Half wave rectification can also be used.
[6] This value is specific to a sinusoidal input. It depends on the shape of the signal.

determined by the time constant of the filter used in the rms detector in Fig. 4.7(b).

This method does not respond to the peak value of the input signal. The averaged rectified output must be multiplied by a suitable factor to estimate the peak value.

4.3.4 Comparison

Figs 4.8 and 4.9 can be used to compare the three types of peak detectors described above.

The peak detector in Fig. 4.7(a) provides a fast attack which helps minimize the overshoot/clipping illustrated in Fig. 4.4(d) and the consequent distortion shown in Fig. 4.4(e). The decay is slower than the attack, but is faster than the decay of the rms and the average detectors[7]. The envelope detector responds to the peak value of the input. Therefore, its output can be used directly for dynamic biasing (after scaling up to obtain the desired overhead). Among the three detectors, the envelope detector has the largest ripple in its output for the same input period and the time constant of the filter used.

The rms detector has a slow attack and decay which means that the nonidealities depicted in Figs 4.4(e, f) will both be present. The rms detector has the advantage that it is not confused by the phases of the frequency components of the input signal: i.e. the rms detector responds identically to the two waveforms shown at the bottom of Fig. 4.6(a, b). But the output should be amplified by the expected maximum crest factor before being used for dynamic biasing.

The attack, decay and the output ripple of the combination of the rectifier and averaging shown in Fig. 4.7(c) are similar to that of the rms detector. Its virtue is that it is the easiest to design among the three detectors presented above.

None of the detectors has a fast decay which means that the signal to noise ratio is always compromised after an abrupt decrease in the input amplitude.

[7]When low-pass filters of equal time constants are used in the three detectors shown in Fig. 4.7. A time constant that is ten times the input period was used to obtain the simualtion results in Fig. 4.8 and Fig. 4.9.

Under ideal conditions, the ripple in the dynamic biasing signal does not appear at the output of the dynamically biased filter. But mismatch between the two halves of the filter can cause some of the ripple to leak to the output and manifest itself as distortion of the input signal.

The peak detector was chosen for the dynamically-biased filter in this work mainly because of its fast attack characteristics. The design of a current mode peak detector is described in Chapter 7.

More sophisticated envelope detectors can be used to tailor the attack and decay times as required. A technique using a combination of slow and fast average detectors is outlined in [41].

4.3.5 Choosing the time constant of the strength detectors

The time constant of the filter used in any of the techniques described above to measure the input signal strength has an influence on the attack and decay times as well as the ripple in the output. The choice of the time constant depends on the frequency components of the expected input signal (i.e. the signal u in Fig. 4.5).

Wideband input signals

Fig. 4.10(a) shows the representative power spectral density of a wideband input signal. The center frequency f_c is comparable to half the bandwidth f_b, or equivalently, the upper frequency limit f_u is much larger than the lower frequency limit f_l. This situation can be expected when the filter that is to be designed is a low pass filter[8]. For dynamic biasing the envelope of the signal that can have any frequency component in this band has to be detected. Since the *average* strength of a signal at a frequency f_l (the lowest possible

[8]Note that there is always a lower limit f_l below which the spectral components are of no interest to us since we are unwilling to wait infinitely long for a "dc" signal.

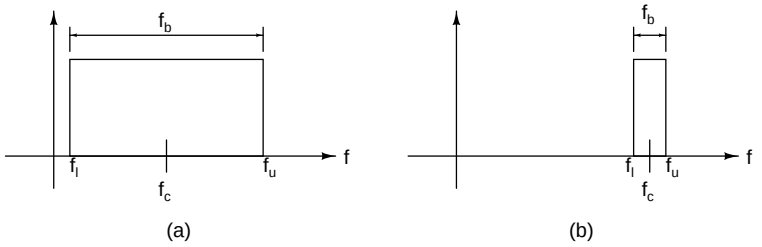

Figure 4.10: Power spectral density of (a) A narrowband signal, (b) A wideband signal.

frequency) has to be detected, the cutoff frequency (∼ reciprocal of the time constant) of the filter used in the envelope detector has to be smaller than f_l.

Narrowband input signals

Fig. 4.10(b) shows the representative power spectral density of a narrowband signal. The bandwidth f_b is much smaller than the center frequency f_c or equivalently, the upper frequency limit f_u and the lower frequency limit f_l are close to each other. The envelope variations of such a signal are constrained to frequencies less than $f_b/2$. In such a case, it is possible to use any one of the envelope detectors described above and track the variations in the envelope. The cutoff frequency (∼ reciprocal of the time constant) of the filter used in the envelope detector must be smaller than the center frequency f_c and larger than the bandwidth f_b. The output of the envelope detector tracks the envelope with a delay that is related to the time constant of the envelope detector. Since the time constant is much smaller than the period (∼ reciprocal of f_b) of the envelope, the output of the envelope detector practically follows the envelope of the input signal.

Such narrowband signals are present in AM radio receivers where the bandwidth of the modulating signal is much smaller than the center frequency. Diode-RC peak detectors are successfully used in the envelope detection of AM radio signals.

However, complications can arise if such a signal is passed through a bandpass filter (e.g. the filter in Fig. 4.5 could be a bandpass filter). Fig. 4.11(a) shows an ampli-

tude modulated waveform whose carrier frequency is 20 times the modulating frequency. This is passed through a second-order bandpass filter with a quality factor of 10 centered at the carrier frequency. Fig. 4.11(b) shows the output of the bandpass filter. The curves Fig. 4.11(c-i) and Fig. 4.11(c-ii) respectively show the envelopes of Fig. 4.11(a) and Fig. 4.11(b) as detected by an envelope detector whose time constant is smaller than the period of the modulating signal. Fig. 4.11(c-i) can be added as a dynamic bias to the signal in Fig. 4.11(a) to obtain a signal that is always positive. But the same is not necessarily true when Fig. 4.11(c-i) is added as a dynamic bias to the signal in Fig. 4.11(b) for the following reasons:

1. The envelope of Fig. 4.11(b) is delayed w.r.t. the envelope of Fig. 4.11(a).

2. The sidebands are attenuated by the bandpass filter resulting in a smaller modulation index in Fig. 4.11(b) than in Fig. 4.11(a). This causes the minima in the envelope of Fig. 4.11(b) to be larger than those in Fig. 4.11(a).

The effect of these phenomena can be seen in Fig. 4.11(c) where the curve (ii) is above curve (i) during certain intervals.

To remedy this situation, there are three possibilities:

1. A large safety margin must be added to Fig. 4.11(c-i) before using it as the dynamic bias so that it is sufficient for both Fig. 4.11(a) and Fig. 4.11(b). If the minima of Fig. 4.11(c-i) are very small the safety margin has to be very large in order for the dynamic bias to be larger than the phase shifted envelope in Fig. 4.11(b). This technique results in a waste of power during the signal peaks.

2. The dynamic bias added at the input and the output should be generated using separate envelope detectors at the input and the output. This methods needs extra hard-

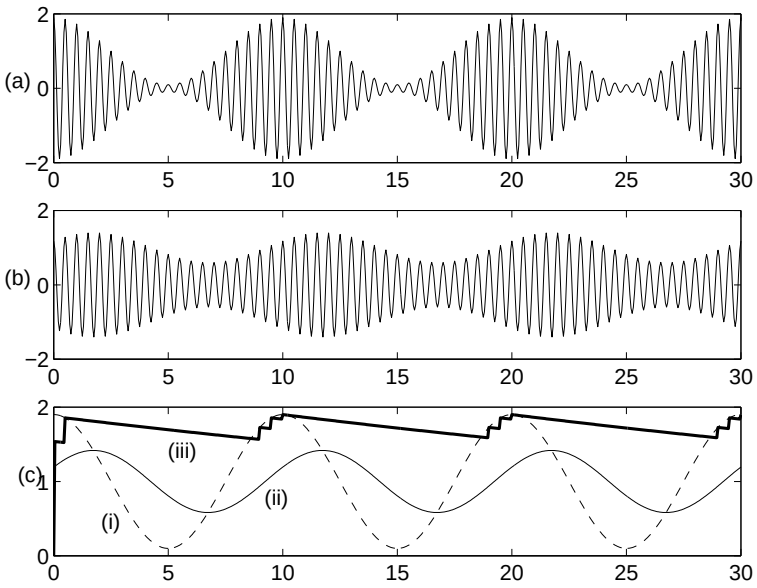

Figure 4.11: (a) Amplitude modulated signal, (b) Output when (a) is passed through a bandpass filter with $Q = 10$, (c-i) Envelope of (a), (c-ii) Envelope of (b), (c-iii) Envelope of (a) as detected by a "slow" peak detector.

ware.

3. Fig. 4.11(c-iii) shows the envelope of Fig. 4.11(a) detected using a long time constant. This must be used instead of Fig. 4.11(c-i) for dynamic biasing. Increasing the time constant results in longer attack and decay times.

It can therefore be concluded that although a narrowband signal lends itself to fast envelope detection, the phase shifts of the envelope as it passes through a narrowband filter may force us to use a slower detection.

The time constant of the filter used in the envelope detector is therefore likely to be larger than the time constants in the filter with either narrowband or wideband input signals. Consequently smaller bias currents and larger capacitors have to be used. With low frequency filters, this may necessitate the use of external capacitors. Note however that high density nonlinear capacitors can be used since the linearity of the filter used in the envelope detector is unimportant.

Chapter 5

Log-Domain Filters in Pure CMOS Technologies

Log-domain filters, which represent the most promising form of instantaneous companding filters, need an exponential nonlinearity for their realization. The exponential nonlinearity of the bipolar junction transistor has been used successfully to implement log-domain filters in Bipolar/BiCMOS technologies. This chapters briefly discusses ways of implementing log-domain filters in pure CMOS[1] technologies.

5.1 MOS transistors operating in weak inversion

MOS transistors operating in weak inversion have a drain current that is exponentially related to the gate-source and the bulk-source voltages. The drain current I_D can be related to the gate-source voltage V_{GS} and the bulk-source voltage V_{BS} using the expression [45]

$$I_D = K\frac{W}{L}\exp\left(\frac{V_{GS}-V_{T0}}{nV_t}\right)\exp\left(\frac{n-1}{n}\frac{V_{BS}}{V_t}\right) \tag{5.1}$$

where K is a constant, V_{T0} is the threshold voltage and n is the subthreshold slope factor which is typically about 1.5.

Log-domain filters using MOS transistors in weak inversion have been investigated in the literature, notably in [46, 47, 48].

[1] The advantages enjoyed by plain digital CMOS processes over other technologies have been discussed ad nauseum in the present day literature on integrated circuits.

Log-domain filters can be realized substituting MOS transistors whose source and bulk are tied together in place of bipolar transistors in well known log-domain circuits (e.g. [6]). But such a technique requires MOS transistors with individual wells. In commonly available technologies, this implies that only pMOS transistors can be used. Also the relatively large parasitic capacitance from the well to the substrate that would be present at the signal carrying nodes limits the frequency response.

A better way to utilize the MOS transistors is to design circuits that make use of the bulk terminal advantageously as shown in [47, 49]. A log-domain filter topology in which the bulk nodes of all the transistors are grounded is presented in [48].

The range of currents over which the exponential law is obeyed is much smaller for an MOS transistor in weak inversion than for a bipolar junction transistor. Another disadvantage of exponential nonlinearity realized using MOS transistors when compared to its bipolar counterpart is the presence of the slope factor n in the argument of the exponential which reduces the transconductance (from the gate-source voltage to the drain current) for given current by the same factor. For example, if the log-domain filter in Fig. 2.7(c) were realized using MOS devices, the pole would be located at I_3/nCV_t instead of I_3/CV_t. This implies that a larger current is required to realize a given pole frequency of the filter.

The threshold voltage of MOS devices is an additional source of mismatch when compared to bipolar devices[2]. Mismatch introduces distortion in class-AB filters due to their internal nonlinearity. It has been demonstrated through simulations in [47] that the mismatch in threshold voltage is a dominant source of distortion in class-AB log-domain filters that use MOS devices in weak inversion.

[2]Bipolar transistors have mismatches in their saturation currents. This mismatch can be represented equivalently by an offset voltage source at the base. The offset voltage, or V_{BE} mismatch, is smaller than the mismatch in the threshold voltages of the MOS transistors.

At this point, not enough experimental results are available to judge the viability of using MOS devices in weak inversion to realize log-domain filters.

5.2 Lateral PNP transistors

Another alternative worth exploring for the realization of log-domain filters in CMOS technologies is the use of lateral bipolar transistors[3]. The most common use of lateral bipolar transistors in a CMOS technologies is in bandgap references [50]. But these are static circuits; i.e. the currents and voltages in the circuits are constant with time. Lateral bipolar transistors have been used previously to realize variable gain amplifiers [51, 52] and an assortment of translinear circuits such as multipliers [53].

5.2.1 Conventional lateral bipolar transistors

Fig. 5.1(a) shows the simplified cross section of a lateral PNP[4] transistor. The "source" and "drain" regions of a pMOS transistors form the collector and emitter and the n- well of the pMOS transistor forms the base. The gate is tied to the most positive voltage in the circuit (commonly done in order to push the carriers below the surface, to avoid imperfections associated with the latter) and has little influence on the operation of the transistor[5]. Fig. 5.1(b) shows the circuit equivalent of Fig. 5.1(a). A vertical transistor Q_S whose collector is the (grounded) substrate is inevitably present along with the desired lateral transistor Q_L.

This lateral bipolar transistor has the following disadvantages.

1. Due to the parallel connection of Q_L and Q_S, only a fraction of the emitter current

[3]The vertical PNP transistor available in plain CMOS technologies has its collector connected to the ground (p-substrate) and hence, its collector current is not accessible. Therefore they are not considered here.

[4]The PNP transistor is the only usable species of the lateral transistor since the base of the lateral NPN transistor is grounded in commonly available technologies.

[5]Note that the gate is almost always present. The transistor may be malformed in its absence [52].

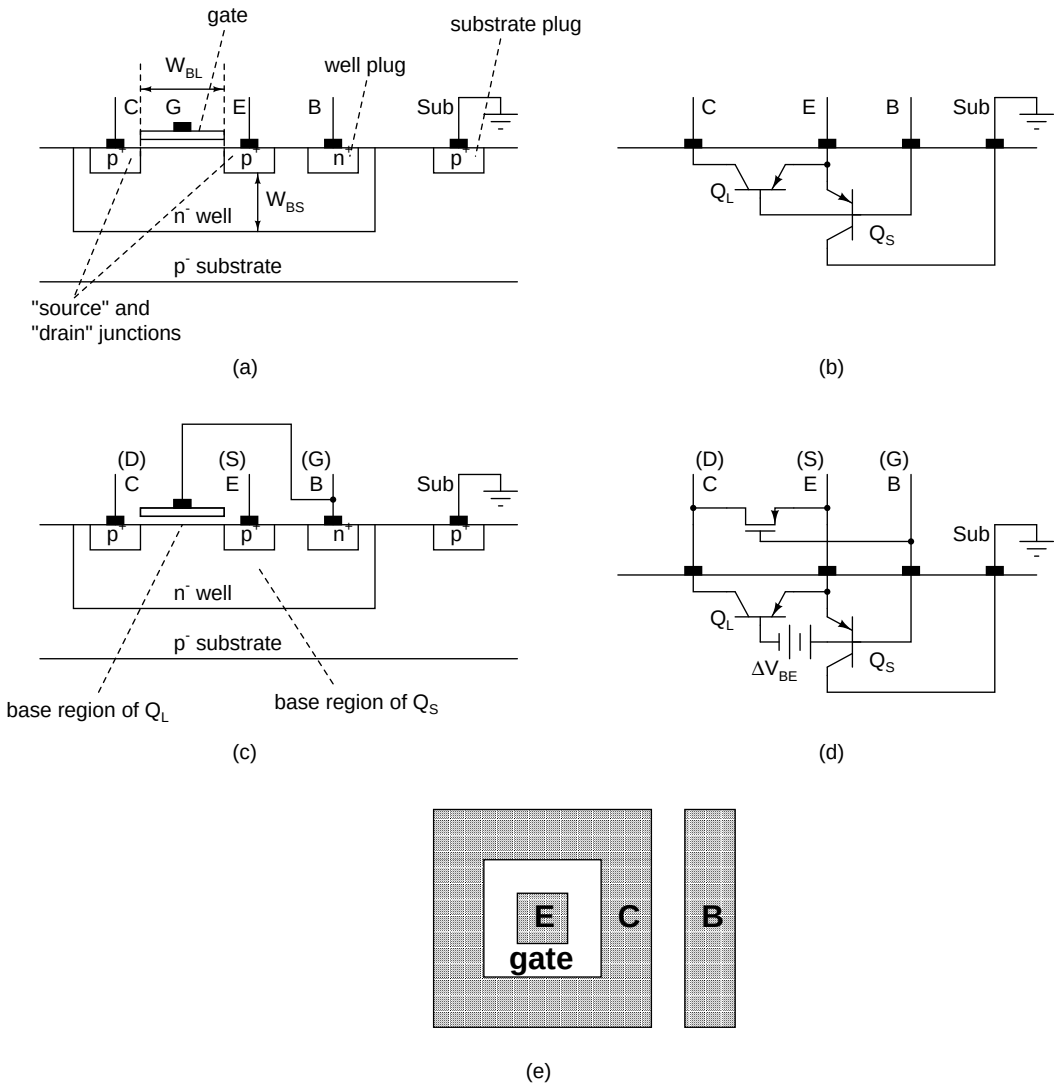

Figure 5.1: (a) Simplified cross section of a conventional lateral PNP transistor, (b) Electrical equivalent of (a), (c) Simplified cross section of an enhanced lateral PNP transistor, (b) Electrical equivalent of (c), (e) Top view of a practical lateral transistor.

reaches the collector. This results in a larger bias current (and hence, power consumption) to realize a given transconductance because the collector current of Q_S flows to the substrate and cannot be tapped.

2. While large values of the current gain β have been reported in the literature, the overall β (defined as the ratio of the collector current of Q_L to the sum of the base currents of Q_L and Q_S) of the structure in Fig. 5.1 in the technology available to this author was not very large. The current gain β of the bipolar transistor worsens with increasing base width. In the currently available short channel technologies, the lateral transistor Q_L has a base width W_{BL} that is small enough for it to have a current gain as high as a hundred. But the current gain of the vertical transistor Q_S is poor due to the large distance between the bottom of the emitter diffusion and the bottom of the n- well. Since the bases of Q_L and Q_S are connected together, the net current flowing into the base can be so large as to reduce the apparent β of the lateral transistor to be about ten.

   Such small values of β are inconvenient even for the β-immune log-domain filter shown in Fig. 2.16 [19, 20]. The compensating current shown in that figure has to be used. The nonlinearities due to variation of β with the collector current are worsened by a low value of β.

5.2.2 Enhanced lateral bipolar transistors

A solution to the problem of the poor current gain of bipolar transistors is given in [54, 55]. Fig. 5.1(c) shows the simplified cross section of the transistor presented in [54, 55]. The gate of the pMOS transistor is connected to the base. This is the exact opposite of the conventional practice of tying the gate to the most positive voltage. In this case, as a negative

V_{BE} is applied, the channel tends to invert. The transistor operates as a combination of a subthreshold device and a lateral bipolar device. The device can be thought of as a bipolar transistor with the assistance of the gate or as a pMOS transistor whose threshold is lowered by the application of a negative bulk-source voltage. The result of tying the gate to the base is the realization of very large dc current gains, and suppression of the substrate transistor action to a great extent [54, 55].

A crude explanation for the behavior of this device is as follows [56]. As shown in Fig. 5.1(c), the base region of the lateral transistor Q_L is near the surface of the device and the base region of the substrate transistor Q_S is deep inside the well underneath the p+ diffusion of the emitter. On applying a negative V_{GS} to the pMOS transistor, an electric field is created underneath the gate. This causes a voltage rise from the surface of the device to the region deep inside the well. This is denoted by ΔV_{BE} in Fig. 5.1(d). The voltage across the base of Q_L and the emitter is larger than the voltage across the base of Q_S and the emitter by ΔV_{BE}. Comparing Fig. 5.1(d) to Fig. 5.1(b), it can be seen that for given V_{BE}, the collector current of Q_S is smaller in the latter. Thus the substrate transistor Q_S is suppressed.

It is reported in [54, 55, 56] that the lateral PNP transistor shown in Fig. 5.1(c) has a dc current gain in the thousands for small collector currents.

Thus the enhanced lateral PNP transistor can be used in log-domain filters without the problem of base currents. One of the biggest disadvantages of lateral bipolar transistors is the large parasitic capacitance from the base (well) to the substrate. But when these transistors are used in the log-domain filter shown in Fig. 2.16 [19] they do not affect the performance seriously because this capacitance appears either across voltage sources or across a desired capacitance. In the latter case, the parasitic capacitance can be absorbed

into the desired integrating capacitance. The parasitics at the base do however limit the largest cutoff frequency that can be achieved.

Compared to MOS transistors operating in weak inversion, the enhanced lateral bipolar transistor has

1. A slope factor of unity, resulting in a larger transconductance for a given current.

2. A larger range of exponential behavior [54].

3. Hopefully better matching; because of its closer adherence to the classical bipolar operation, the threshold voltage may not have as much influence on the drain current.

Both the MOS transistor in weak inversion and the lateral bipolar transistor have a frequency response that is inferior to that of a conventional bipolar transistor. This is because the MOS transistor in weak inversion is operated at small current densities and the lateral PNP transistor has a large parasitic capacitance from the n- well that acts as its base.

Figs. 5.1(a, b) depicted a simplified cross section of the lateral PNP transistors in which the collector was on one side of the emitter. In practice, to maximize the collector efficiency, the emitter is surrounded by the collector as shown in Fig. 5.1(e). The base contact is placed outside the collector ring. A circular shape for the collector and emitter provides the best performance by minimizing the area of the substrate emission area, but an octagonal layout may be more convenient and not very sub-optimal.

5.2.3 Use of MOS capacitors in log-domain filters

Since the voltages across the capacitors in a log-domain filter have a limited peak-peak swing of a few V_ts, nonlinear capacitors can be used safely without causing distortion [47]. MOS capacitors in accumulation behave practically linearly over such small voltages [57].

This enables the use of large capacitors which are required if a large signal to noise ratio is demanded from the filter.

Chapter 6

Noise Analysis of Companding Filters

6.1 Introduction

Although companding filters are linear and time-invariant between their input and output, they are by construction internally nonlinear. The signal dependent output noise caused by such a nonlinearity is in fact the chief motivation for considering companding techniques to improve the dynamic range per unit power consumption of filters. Estimation of output noise of such filters is essential to determining their signal to noise ratio and dynamic range. But the presence of such nonlinearities between the internal noise sources and the output terminal implies that the traditional methods of noise calculation/simulation (linearization around the bias point, e.g. ".AC" in SPICE) do not suffice.

The behavior of noise in companding filters has been intuitively pointed out in [3]. Several papers have been published to date outlining the steps for calculating the noise in companding filters[32, 33, 36]. In this chapter we describe a generally applicable method for noise analysis of instantaneously companding filters using a first order filter as an example. While the method outlined in the next section is applicable to all instantaneous companding filters, log-domain filters are used as examples for the most part since they are the only types whose practical realization has been sufficiently investigated to date. Simulation methods for calculating the noise and possible simplifications for important

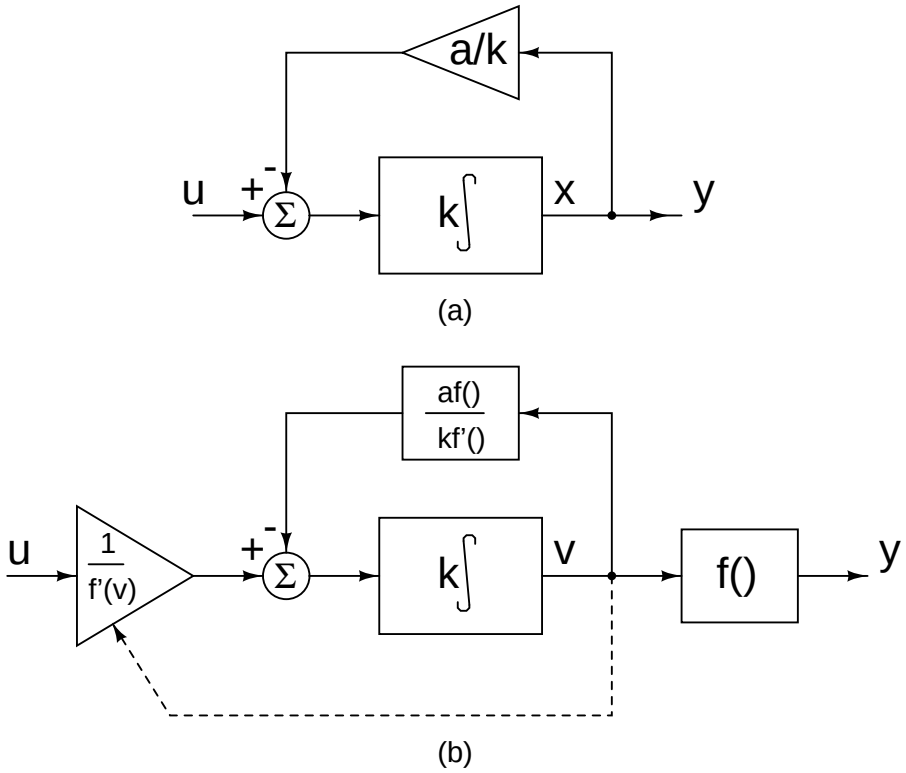

Figure 6.1: (a) Linear first order filter, (b) Instantaneous companding first order filter.

special cases are also pointed out.

6.2 General technique

6.2.1 Noise in instantaneously-companding filters

A first order companding filter is used as an example. Fig. 6.1(a) shows a linear first order filter whose state variable is denoted by x. Fig. 6.1(b) shows its companding equivalent which has an output nonlinearity $f()$ and a transformed state variable v as described in Chapter 2 [26].

The relation between the input u and the output y of these filters in the frequency domain is given by

$$H(s) = \frac{Y(s)}{U(s)} = \frac{k}{s+a} \tag{6.1}$$

where k and a are known constants. The state variable description of the first order com-

panding filter in terms of the transformed state variable v, input u and the output y is given by (2.12) and (2.13)

$$\dot{v} = -a\frac{f(v)}{f'(v)} + \frac{k}{f'(v)}u \tag{6.2}$$

$$y = f(v) \tag{6.3}$$

Noise can be added to various points in the system in Fig. 6.1(b). Noise added to the input will be processed in a linear time invariant fashion with the transfer function given in (6.1). Noise at the output of the nonlinearity $f()$ is simply added to the output. More interesting is the behavior of noise added to the internal nodes. Fig. 6.2 shows three such possibilities. Noise n_1 in Fig. 6.2(a) is added at the input of the output nonlinearity $f()$. Noise n_2 in Fig. 6.2(b) is added after the input amplifier. Noise n_3 in Fig. 6.2(c) is added to the feedback path. n_{o1}, n_{o2} and n_{o3} denote the output noise in Fig. 6.2(a), (b) and (c) respectively. Noise added to other points can be reduced to either one or a combination of these cases. These three possibilities will be separately treated below.

It is clear that in Fig. 6.2(a), the addition of noise n_1 does not alter the state variable v from its value in the noiseless case (Fig. 6.1(b)). Only the input to the nonlinearity $f()$ is changed. The new output of the filter y_1 is given by

$$y_1 = f(v + n_1) \tag{6.4}$$

Assuming that n_1 is sufficiently small, which presumably would be the case in a usable filter, the output can be approximated to

$$y_1 = f(v) + f'(v)n_1 \tag{6.5}$$

The output noise n_{o1} is the difference between the noisy output y_1 and the noiseless output y (given by (6.3)).

$$n_{o1}(t) = y_1(t) - y(t) \tag{6.6}$$

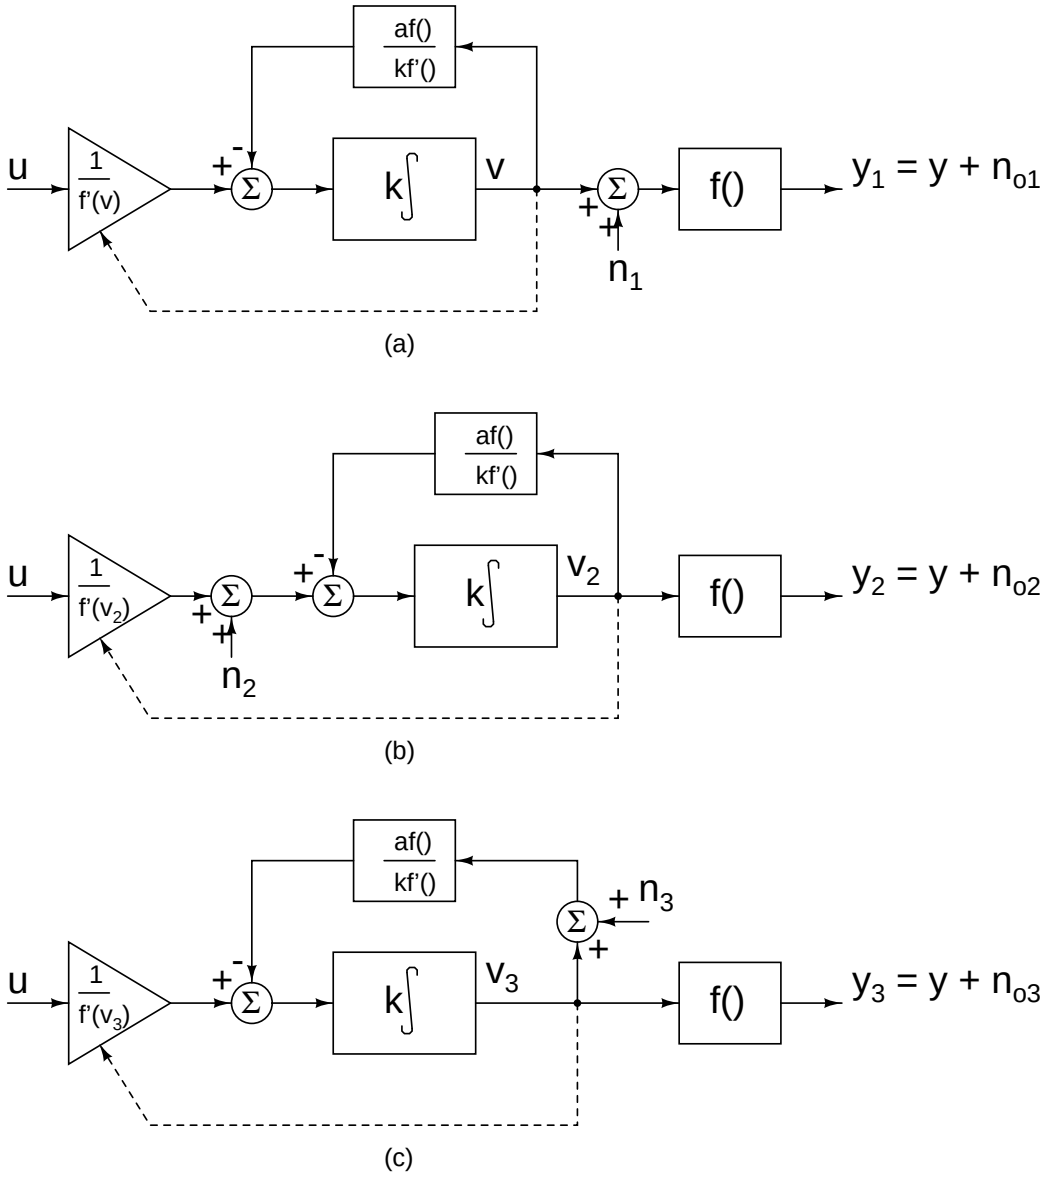

Figure 6.2: (a) Noise added at the input of the output nonlinearity, (b) Noise added after the input amplifier, (c) Noise added in the feedback path.

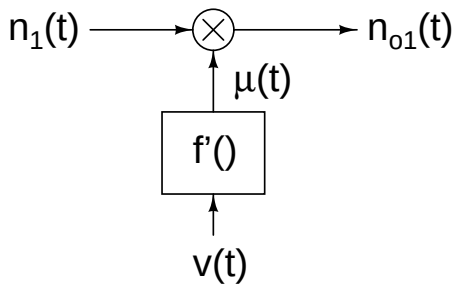

Figure 6.3: Noise equivalent circuit for Fig. 6.2(a).

$$= f'(v(t))\, n_1(t) \tag{6.7}$$

$$= \mu(t)\, n_1(t) \tag{6.8}$$

where $\mu(t)$ is defined as

$$\mu(t) = f'(v(t)) \tag{6.9}$$

$\mu(t)$ denotes the time-varying small-signal gain of the output nonlinearity $f()$. The resulting small signal model is shown in Fig. 6.3. The added noise is multiplied by a signal dependent gain. The calculation of the output power spectral density is given later in this section.

This result is hardly surprising. As described in the first two chapters, in the introduction to companding, a signal dependent gain is placed at the output of the filter to alter the overall output noise in accordance with the signal strength. It was also mentioned that the expanding nonlinearity $f(v)$ used at the output of instantaneously companding filter has a derivative that increases with increasing v. Therefore, the noise n_1 is multiplied by small gains when the output signal is small and by large gains when the output signal is large.

The case shown in Fig. 6.2(b) where noise n_2 is added to the output of the input block is slightly more involved. In this case, the state variable v is altered from its noiseless value to a new value v_2. The effect of the noise n_2 at the filter's output can be determined

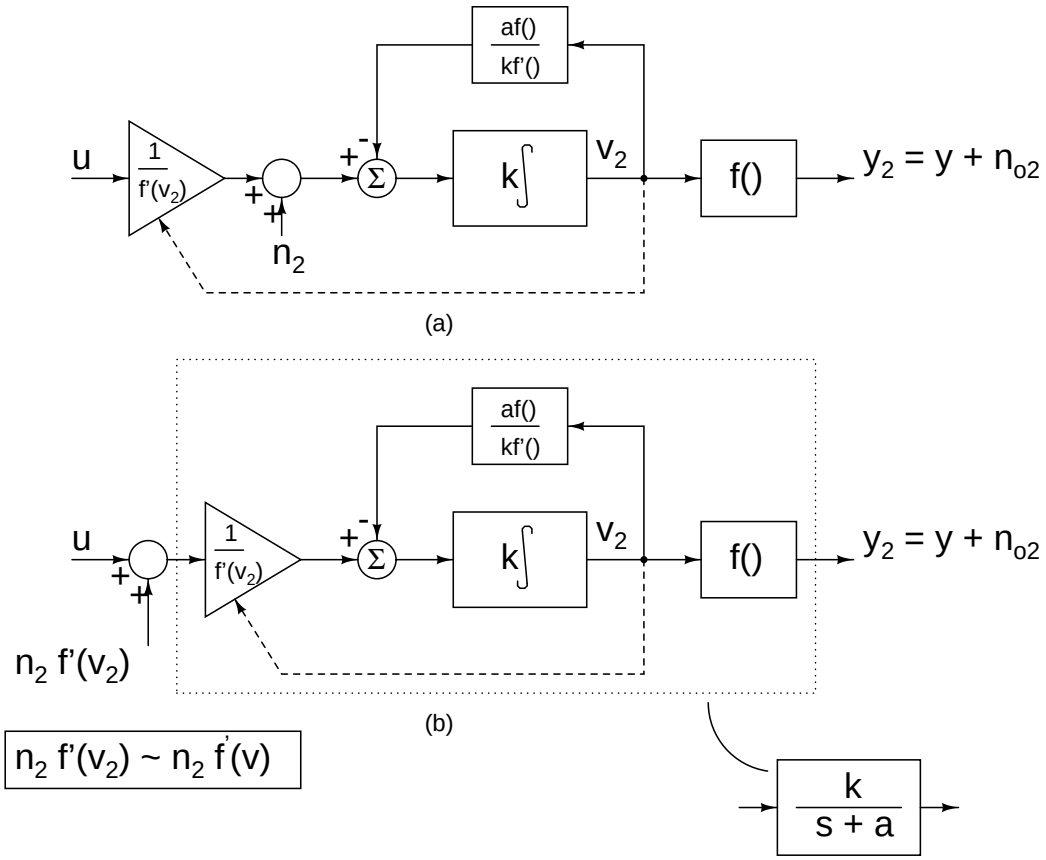

Figure 6.4: Transformation of Fig. 6.2(b) to a convenient form.

using the transformation shown in Fig. 6.4. Fig. 6.4(a) (same as Fig. 6.2(b)) shows noise n_2 being added after the input amplifier. As shown in Fig. 6.4(b), this can be transferred to the amplifier's input after dividing by the gain of that amplifier $1/f'(v_2)$. It can be easily seen that the part of the resulting circuit enclosed within dotted lines is the filter in Fig. 6.1 with the transfer function given by (6.1). Thus noise n_2 is multiplied by a signal dependent factor $f'(v_2)$, added to the input u and filtered. This multiplying factor itself depends on noise through v_2. But for small noise, which would be the case in a practical circuit, and a nonlinearity $f()$ with a continuous first derivative, the factor $f'(v_2)$ can be approximated by $f'(v)$ so that it can be calculated using the solution to the noiseless circuit. The filter enclosed by dotted lines, which is linear and time invariant, has an output $y + n_{o2}$ for an

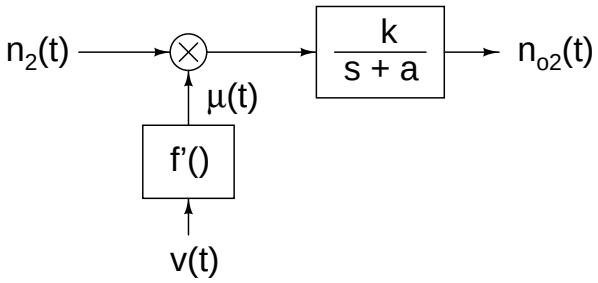

Figure 6.5: Noise equivalent circuit for Fig. 6.2(b).

input $u + n_2 f'(v)$. The same filter has an output y for an input u (Fig. 6.1(b)). Thus n_{o2} can be computed as the output of the linear time invariant filter shown in Fig. 6.1 for an input $n_2 f'(v)$ where v is the noiseless state variable of the filter for an input u. Fig. 6.4(b). Fig. 6.5 shows the resulting equivalent circuit (with $\mu(t)$ defined as before) from which the output noise n_{o2} can be computed.

The case in Fig. 6.2(c) in which noise is added in the feedback path can be treated similarly. Fig. 6.6(a) (same as Fig. 6.2(c)) shows noise n_3 added to the feedback path. In this case, the quantity fed back to the input of the integrator v_{3fb} is given by

$$v_{3fb} = \frac{a}{k}\frac{f(v_3+n_3)}{f'(v_3+n_3)} \tag{6.10}$$

$$= \frac{a}{k} f_1(v_3+n_3) \tag{6.11}$$

where a new symbol $f_1()$ is introduced to denote $f()/f'()$. Assuming that n_3 is sufficiently small, the following approximations can be made.

$$v_{3fb} \approx \frac{a}{k} f_1(v_3) + \frac{a}{k} f_1'(v_3) n_3 \tag{6.12}$$

$$= \frac{a f(v_3)}{k f'(v_3)} + \nu n_3 \tag{6.13}$$

where ν in the last equation is given by

$$\nu(v_3) = \frac{a}{k} f_1'(v_3) \tag{6.14}$$

$$= \frac{a}{k}\left(1 - \frac{f(v_3)f''(v_3)}{(f'(v_3))^2}\right) \tag{6.15}$$

Fig. 6.6(a) can be transformed into Fig. 6.6(b) where the two terms of the fed back quantity v_{3fb} given by (6.13) are added separately at the input of the integrator. Fig. 6.6(c) shows a further transformation whereby a noise $\mu_3(t)n_3(t)$ is added to the input. μ_3 is given by

$$\begin{aligned} \mu_3(t) &= f'(v_3(t))\nu(v_3(t)) && (6.16)\\ &\approx f'(v(t))\nu(v(t)) && (6.17)\\ &= \frac{a}{k}\left(f'(v(t)) - \frac{f(v(t))f''(v(t))}{f'(v(t))}\right) && (6.18) \end{aligned}$$

where, as before, the functions are evaluated at the noiseless state variable v instead of the noisy state variable v. The circuit inside the dotted lines is again seen to be the filter in Fig. 6.1(b) with the transfer function in (6.1). n_{o3} can thus be computed as the output of a linear time invariant filter to an input $-\mu_3(t)n_3(t)$. The equivalent circuit so obtained is shown in Fig. 6.7.

A feature common to the three cases is multiplication of noise by a signal dependent factor which results in signal-dependent output noise.

The equivalent circuits obtained above can also be derived analytically as demonstrated in [33] for the cases in Fig. 6.2(a, b).

6.2.2 Response to stationary white noise in the presence of a periodic input

While an analytical solution for the output noise spectral density is impossible to obtain in the general case with arbitrary inputs and random noise, the special case of periodic inputs and stationary white noise is tractable and is dealt with below.

In steady state, if the input to the circuit in Fig. 6.1(b) is periodic, the state variable v and the output y will be periodic as well. This means that the multiplying factors in Figs. 6.3, 6.5 and 6.7 will be periodic. Therefore the equivalent circuits shown will be linear periodically time varying systems. The output noise in such circuits can be computed

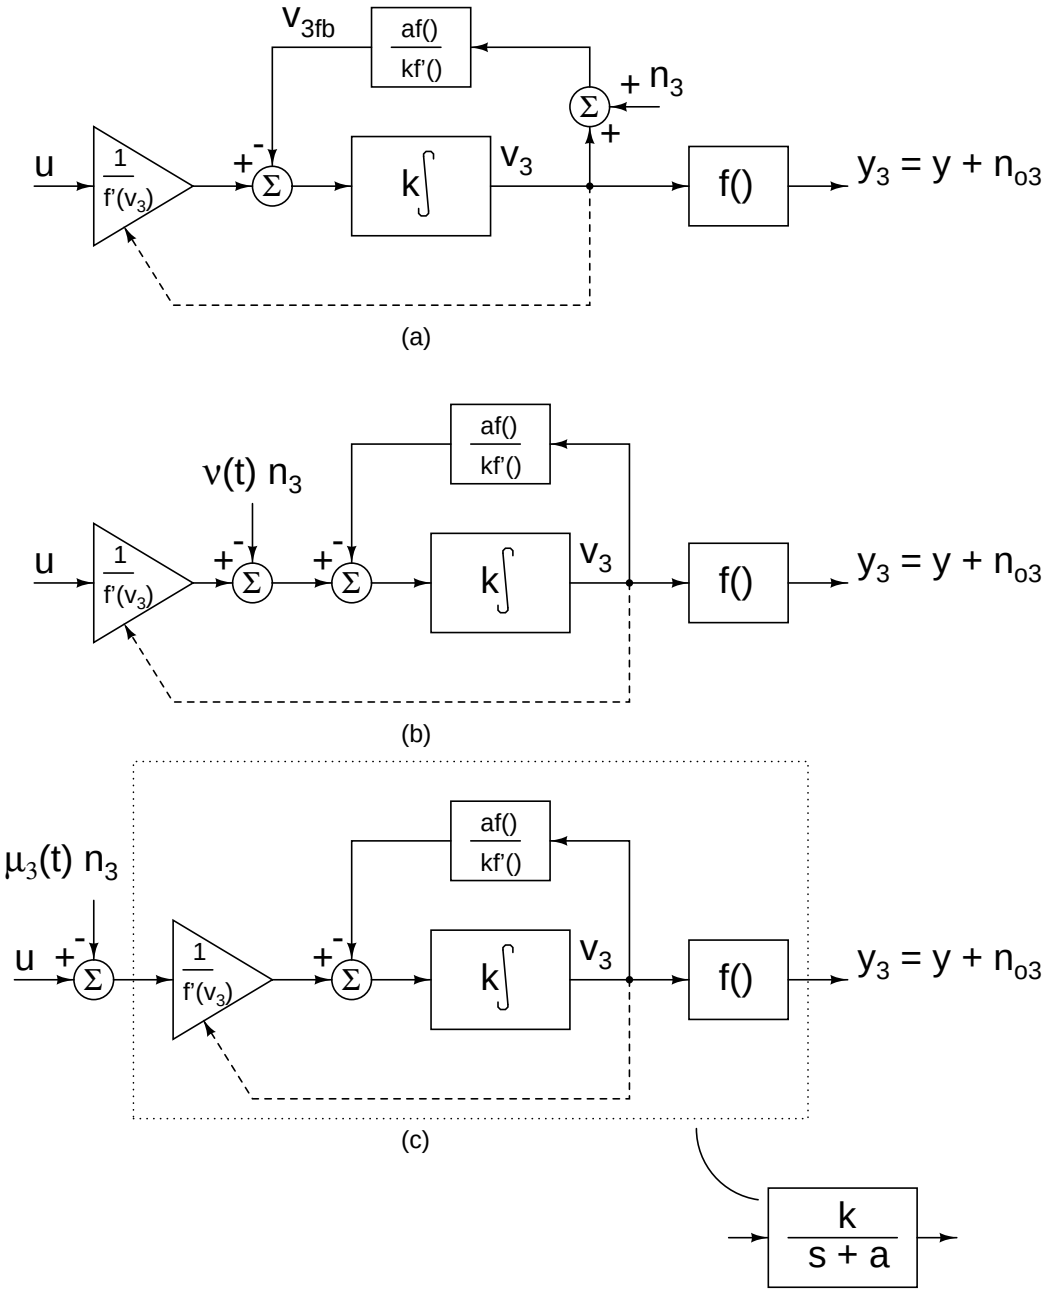

Figure 6.6: Transformation of Fig. 6.2(c) to a convenient form.

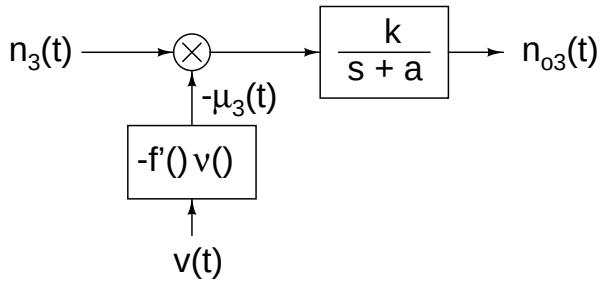

Figure 6.7: Noise equivalent circuit for Fig. 6.2(c).

using the response of the system to delayed delta functions at the source of noise [58, 59].

A linear periodically time varying circuit is characterized by its response to a delayed delta function $\delta(t-\tau)$. In Fig. 6.3, if $n_1(t)$ is set to $\delta(t-\tau)$, the output $h_1(t,\tau)$ would be

$$h_1(t,\tau) = \mu(t)\delta(t-\tau) \tag{6.19}$$

$$= \mu(\tau)\delta(t-\tau) \tag{6.20}$$

The response $h_2(t,\tau)$ of the circuit in Fig. 6.5 to $\delta(t-\tau)$ is given by

$$h_2(t,\tau) = (\mu(t)\delta(t-\tau)) * h(t) \tag{6.21}$$

$$= \mu(\tau)(\delta(t-\tau) * h(t)) \tag{6.22}$$

$$= \mu(\tau)h(t-\tau) \tag{6.23}$$

where $h(t)$ is the impulse response of the linear time invariant filter shown in Fig. 6.1 (equivalently, the inverse Laplace transform of (6.1)) and $*$ denotes convolution.

Similarly, the response $h_3(t,\tau)$ of the circuit in Fig. 6.7 to $\delta(t-\tau)$ is given by

$$h_3(t,\tau) = (-\mu_3(t)\delta(t-\tau)) * h(t) \tag{6.24}$$

$$= -\mu_3(\tau)(\delta(t-\tau) * h(t)) \tag{6.25}$$

$$= -\mu_3(\tau)h(t-\tau) \tag{6.26}$$

If white noise with spectral density S_n is fed to a linear periodically time varying (LPTV) system whose response to $\delta(t-\tau)$ is $h(t,\tau)$ the output spectral density can

Table 6.1: Output noise PSDs in Fig. 6.2

	Source	$h(t,\tau)$	$H(\omega,\tau)$	Output PSD
Fig. 6.2(a)	n_1	$\mu(\tau)\delta(t-\tau)$	$\mu(\tau)e^{-j\omega\tau}$	$S_{o1} = \frac{S_{n1}}{T}\int_0^T \mu^2(\tau)d\tau$
Fig. 6.2(b)	n_2	$\mu(\tau)h(t-\tau)$	$\mu(\tau)H(j\omega)e^{-j\omega\tau}$	$S_{o2} = \frac{S_{n2}}{T}\frac{k^2}{a^2+\omega^2}\int_0^T \mu^2(\tau)d\tau$
Fig. 6.2(c)	n_3	$-\mu_3(\tau)h(t-\tau)$	$-\mu_3(\tau)H(j\omega)e^{-j\omega\tau}$	$S_{o3} = \frac{S_{n3}}{T}\frac{k^2}{a^2+\omega^2}\int_0^T \mu_3^2(\tau)d\tau$

be evaluated using the following expression [58, 59].

$$S_o(\omega) = S_n\frac{1}{T}\int_0^T |H(\omega,\tau)|^2 d\tau \tag{6.27}$$

where T is the period of the periodic variation and $H(\omega,\tau)$ is the Fourier transform of $h(t,\tau)$ w.r.t. the variable t.

Thus the PSD of output noise n_{o1}, n_{o2} or n_{o3} can be computed analytically when n_1, n_2 or n_3 in Fig. 6.2 is white with a PSD S_{n1}, S_{n2} or S_{n3} respectively. For each of n_{o1}, n_{o2} or n_{o3}, the corresponding $H(\omega,\tau)$ can be calculated using (6.20), (6.23) or (6.26) and substituted in (6.27). The results so obtained are summarized in Table 6.1.

6.2.3 Computing the total output noise of a circuit

The previous section outlined the methods for calculating the contribution of noise n_1, n_2 or n_3 in Fig. 6.1 when each noise source is acting alone. In a practical circuit, all these sources are simultaneously present and the output noise is the sum of the output contributions due to each noise source. If the sources are uncorrelated, the output PSD can be computed simply by summing the output noise PSDs due to each source (Table 6.1).

The situation is complicated if the sources are correlated. Correlations could exist between n_1, n_2 and n_3 even though noise at the device level is uncorrelated. The correlation could be due to a single device contributing noise to two different points of the circuit. An

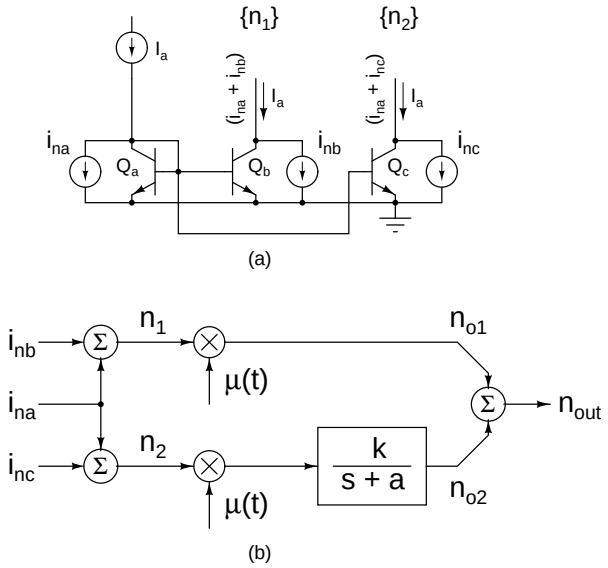

Figure 6.8: (a) Current mirror with correlated output noise, (b) Equivalent circuit.

Table 6.2: Output noise PSDs in Fig. 6.8

Source	$h(t,\tau)$	$H(\omega,\tau)$	Output PSD
i_{na}	$\mu(\tau)(\delta(t-\tau)+h(t-\tau))$	$\mu(\tau)(1+H(j\omega))e^{-j\omega\tau}$	$S_{oa}=\frac{S_{na}}{T}\frac{(k+a)^2+\omega^2}{a^2+\omega^2}\int_0^T\mu^2(\tau)d\tau$
i_{nb}	$\mu(\tau)\delta(t-\tau)$	$\mu(\tau)e^{-j\omega\tau}$	$S_{ob}=\frac{S_{nb}}{T}\int_0^T\mu^2(\tau)d\tau$
i_{nc}	$\mu(\tau)h(t-\tau)$	$\mu(\tau)H(j\omega)e^{-j\omega\tau}$	$S_{oc}=\frac{S_{nc}}{T}\frac{k^2}{a^2+\omega^2}\int_0^T\mu^2(\tau)d\tau$
Total output noise PSD			$S_{out}=S_{oa}+S_{ob}+S_{oc}$

example of this is shown in Fig. 6.8(a) in which a reference current I_a is mirrored to obtain two current sources. Noise currents i_{na}, i_{nb} and i_{nc} from the three transistors Q_a, Q_b and Q_c are uncorrelated. The output noise from the two current source transistor Q_b and Q_c are $i_{na}+i_{nb}$ and $i_{na}+i_{nc}$ respectively. Clearly, they are correlated.

These current sources could be used in the implementation of Fig. 6.1(b) such that they contribute noise in the form of n_1 and n_2 (see Fig. 6.2) respectively. In this case, the output noise PSD cannot be obtained simply by summing S_{o1} and S_{o2} given in Table 6.1. Fig. 6.8(b) shows the equivalent circuit for this situation. n_{o1} and n_{o2} are related to n_1 and

n_2 via the equivalent circuits in Figs. 6.3 and 6.5. n_1 and n_2 in turn are linear combinations of i_{na}, i_{nb} and i_{nc}. Thus the resulting circuit is yet another LPTV system with uncorrelated noise inputs i_{na}, i_{nb} and i_{nc} and an output n_{out}. Assuming they are stationary and white, the output power spectral density due to each of these three sources can be determined as before from the corresponding responses to delayed delta functions. Since the three sources are uncorrelated, the total output noise PSD can be obtained by summing them. The results are summarized in Table 6.2.

Fig. 6.8(b) showed n_1 and n_2 as a linear combination of uncorrelated noise sources i_{na}, i_{nb} and i_{nc}. There may be cases where dependence of n_1, n_2 and n_3 in Fig. 6.2 on noise sources at the device level (such as i_{na}, i_{nb} and i_{nc}) is not known. In such cases, the power spectral density matrix of n_1, n_2 and n_3, which would be the complete description of the noise sources including their correlations can be used to calculate the output noise. The power spectral density matrix of a set of k noise sources is an $k \times k$ matrix whose k diagonal terms are the power spectral densities of each of the noise sources and off-diagonal terms are the cross spectral densities [60, 4].

The k correlated noise sources can be generated as a linear combination of k uncorrelated noise sources.

$$\vec{n_c} = \mathbf{M}\vec{n_{uc}} \tag{6.28}$$

where $\vec{n_c}$ and $\vec{n_{uc}}$ are the correlated and uncorrelated noise vectors respectively and $\mathbf{M} = [m_{ij}]$ is an $k \times k$ square matrix. $\mathbf{M}$ is not unique; only $k(k+1)/2$ of its k^2 elements are independent. Extra constraints can be placed in order to fill in the rest of the matrix. For example, $\mathbf{M}$ could be constrained to be symmetric or top diagonal[1] and so on.

Fig. 6.9 shows n_1, n_2 and n_3 being generated from a set of three uncorrelated noise

[1] This would be akin to Gram-Schmidt orthonormalization of vectors[61].

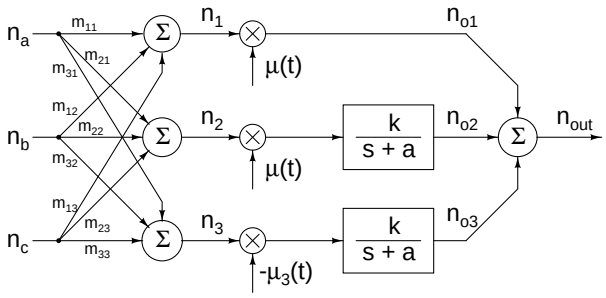

Figure 6.9: Model for calculating the output noise due to correlated noise sources.

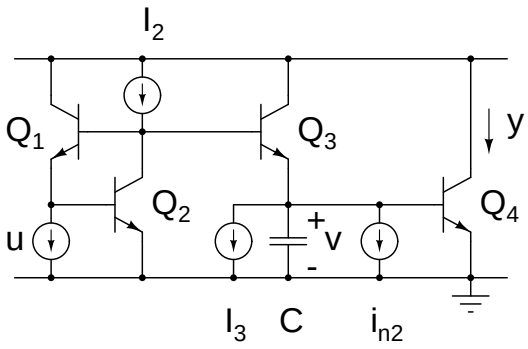

Figure 6.10: First order class-A log-domain filter.

sources n_a, n_b and n_c. n_1, n_2 and n_3 are fed to the equivalent circuits shown in Figs. 6.3, 6.5 and 6.7 to produce the output noise. This system is very similar to that in Fig. 6.8(b) which was obtained using the known dependence of n_1 and n_2 to noise from the individual devices. Again we have a linear periodically time-varying system with uncorrelated noise sources as inputs. The response of the system to a delayed impulse from each of the noise sources n_a, n_b and n_c is a linear combination of the responses to delayed impulse from the sources n_1, n_2 and n_3. The output noise can be obtained by summing the contributions from each source which in turn can be obtained by substituting the appropriate $H(\omega, \tau)$ in (6.27).

6.2.4 Experimental verification

Measurements were performed on the discrete component versions of the first order log-domain circuits of Fig. 6.10 and Fig. 6.11 using $C = 10\,\text{nF}$ and $I_2 = I_3 = 20\,\mu\text{A}$. These

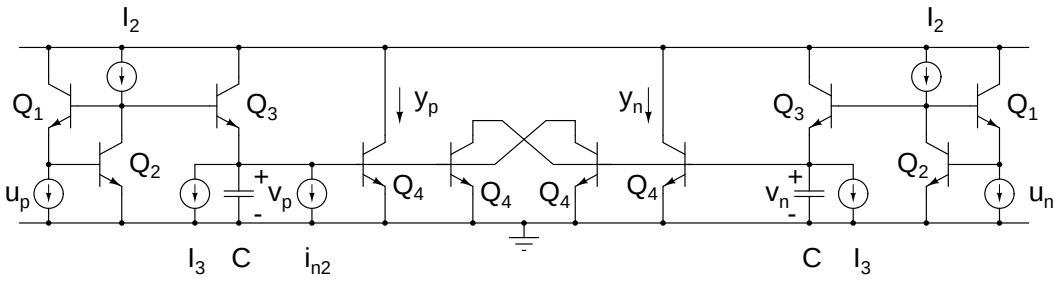

Figure 6.11: First order class-B log-domain filter.

values result in $k = a = 2\pi \times 12.3\,\text{krad/s}$ in Fig. 6.1. Current source i_{n2} in these figures denotes the externally injected interference and it corresponds to n_2 in Fig. 6.2(b). Experiments were performed both with a deterministic sinusoidal interference and random noise. i_{n2} was deliberately made large enough to be the dominant noise source in the circuit so as to enable easy comparison to analytical results. The output current was converted to a voltage using a transresistance amplifier, and the noise was measured using a spectrum analyzer. The experiments performed are described below.

1. Class–A circuit (Fig. 6.10):

   - i_{n2} was a sinusoidal interference of $0.1\,\mu\text{A}\cos(2\pi \times 11\text{kHz} \times t)$. The input to the filter was $20\,\mu\text{A} + 5\,\mu\text{A}\cos(2\pi \times 2\text{kHz} \times t)$. From the equivalent circuit of Fig. 6.5, the intermodulation components were calculated. The calculated values and the measured points are indicated on Fig. 6.12.
   - White noise with a PSD of $0.35\,\text{nA}/\sqrt{\text{Hz}}$ was injected as shown in Fig. 6.10. The input to the filter was $20\,\mu\text{A} + 5\,\mu\text{A}\cos(2\pi \times 6\text{kHz} \times t)$. The output noise power spectral density was calculated using the last equation in the second row of Table 6.1. The calculated curve and the measured points are shown in Fig. 6.13.

2. Class–B circuit (Fig. 6.11)

   - A noise current with a PSD of $0.35\,\text{nA}/\sqrt{\text{Hz}}$ was injected as in Fig 6.11. Again,

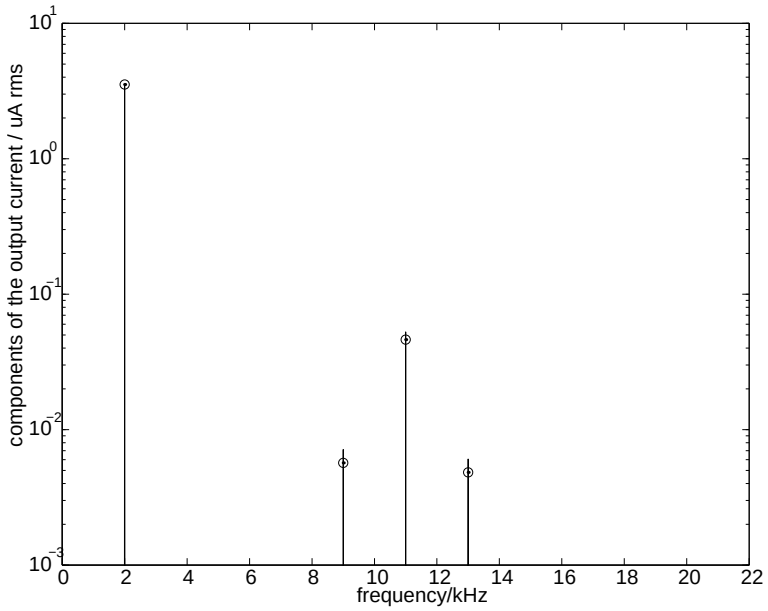

Figure 6.12: Intermodulation due to sinusoidal interference—lines: calculated, circles: measured.

the last equation in the second row of Table 6.1 is used to compute the output noise PSD. The calculated curves and measured points are shown in Fig 6.11. It is apparent that the noise increases with an increase in the input signal.

The analytical calculation of the output PSD in the above cases is based on the equivalent circuit of Fig. 6.5. Good agreement is observed between calculations and measurements. The details of calculation of the signal dependent multiplier $\mu(t)$ for the particular cases of Fig. 6.10 and Fig. 6.11 can be found in [33].

6.3 Simulation methods

6.3.1 Transient noise analysis

A brute force method of simulating noise in nonlinear circuits is to run transient simulation of the circuits including random noise sources with appropriate statistics. Two versions of the circuit, one with noise sources and one without them are simulated and their difference

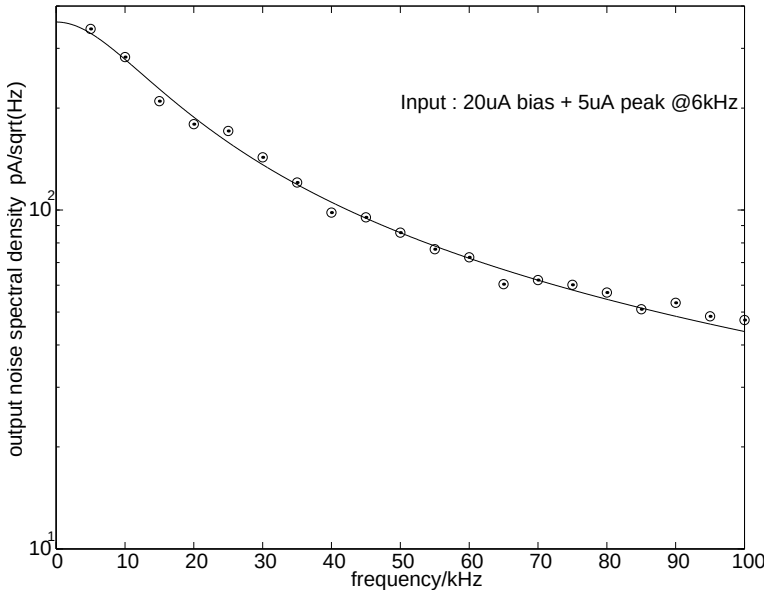

Figure 6.13: Output noise of a first order class-A filter—lines: calculated, circles: measured.

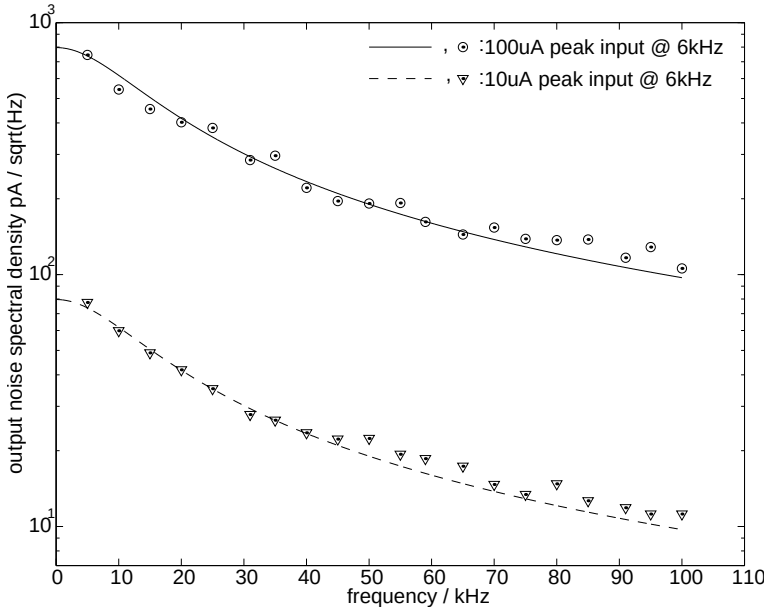

Figure 6.14: Output noise of a first order class-B filter—lines: calculated, circles: measured.

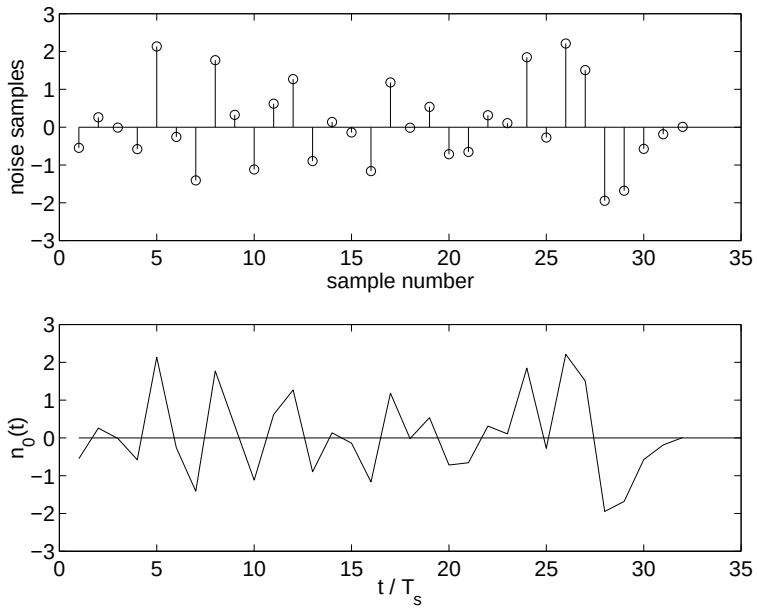

Figure 6.15: (a) Discrete time random sequence, (b) Linearly interpolated continuous time waveform.

taken. This difference is the output noise of the circuit for one set of noise waveforms. This simulation is repeated several times with new uncorrelated waveforms for the noise sources. The average of the magnitude squared of the Fourier transform of the resultant noise waveforms is an estimate of the output power spectral density. The estimate gets better as more waveforms are included in the averaging.

In order to simulate noise in log-domain filters, the shot noise of bipolar transistors (and possibly the thermal noise from ohmic resistances) have to be modeled with appropriate waveforms. HSPICE, which was used for the work in this thesis, as well as several other simulators provide piecewise linear sources whose values at certain time points can be specified. The value of the waveform between these time points is determined by linear interpolation. This facility was used to generate the noise waveforms.

To model white noise, a sequence of independent identically distributed (i.i.d.) gaussian samples with unit variance is generated externally (e.g. using MATLAB). The contin-

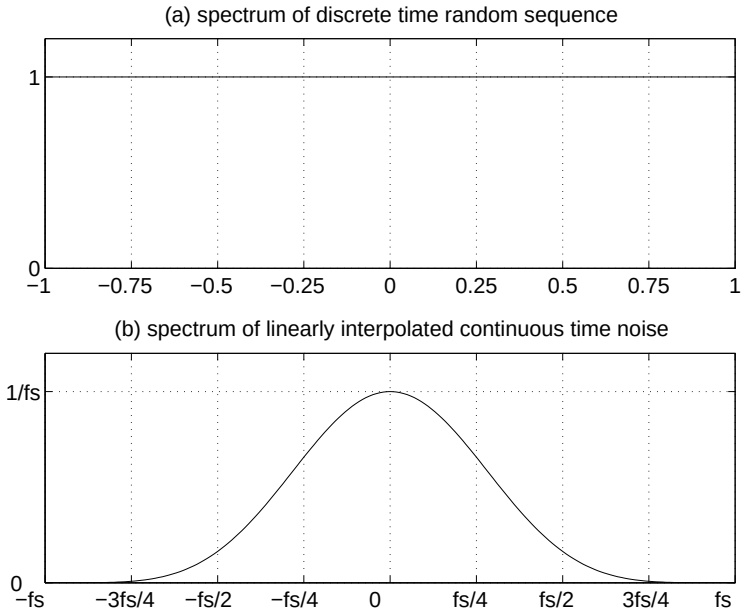

Figure 6.16: (a) Spectral density of discrete time uncorrelated sequence, (b) Spectral density of linearly interpolated continuous time waveform.

uous time waveform that represents the noise takes the value of these samples at integer multiples of a certain period $T_s = 1/f_s$ and successive samples of the waveform are connected by straight lines (Fig. 6.15). The discrete time random sequence has a white power spectral density (double sided) as shown in Fig. 6.16(a) [62]. Linear interpolation has a low pass filtering effect and the interpolated waveform has the double sided power spectral density shown in Fig. 6.16(b). For frequencies much smaller than the sampling frequency, it appears "white".

The subcircuit shown in Fig. 6.17(a) is used for each transistor in order to include the shot noise. The value of the noise current source in parallel with the transistor is given by

$$i_n(t) = n_0(t)\sqrt{q f_s i_c(t)} \tag{6.29}$$

where $n_0(t)$ is the piecewise linear waveform formed from a gaussian distributed random sequence with unit variance, q is the electron charge, f_s is the sampling frequency at which the noise samples are specified, and $i_c(t)$ is the instantaneous collector current. When

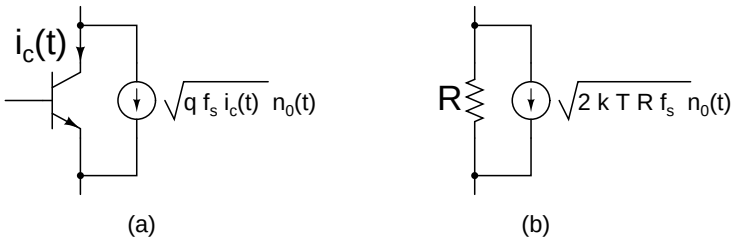

Figure 6.17: (a) Noise model of a transistor, (b) Noise model of a resistor.

$i_c(t) = I_c$, a constant current, it can be seen from the scaling factor in Fig. 6.16 and (6.29) that the double sided spectral density at low frequencies reduces to qI_c as it should. The square root dependence on the instantaneous collector current is a common way (e.g. [63]) of taking into account the non-stationarity of the noise with time varying collector current.

In a similar way, a resistor's thermal noise can be represented by a parallel current source as shown in Fig. 6.17(b). The value of the current source is

$$i_n(t) = n_0(t)\sqrt{2kTRf_s} \tag{6.30}$$

where $n_0(t)$ is the piecewise linear waveform formed from a gaussian distributed random sequence with unit variance, k is the Boltzmann's constant, T is the absolute temperature R is the value of the resistor and f_s is the sampling frequency at which the noise samples are specified.

This technique has been used to obtain the signal dependent output noise of dynamically biased first order filters described in Chapter 2 and Chapter 3. Fig. 2.23 shows the output noise spectral density of the filter in Fig. 2.21. Fig. 3.4 shows the output noise waveform in time domain of the filter in Fig. 3.1. In both these cases, dependence of output noise on input signal strength is clearly seen.

The technique described in this section is time consuming due to the fact that a large number of transient runs are required to obtain the averaged power spectral density curve.

Also, the sampling frequency used to represent the samples has to be high enough to ensure that the power spectral density appears white in the frequency range of interest. This implies long transient runs, but the technique is conceptually simple. Non-stationary noise can be easily dealt with and no analytical calculations are required. A one time effort to setup noise models (Fig. 6.17) and to automate the process of repeating transient analyses and averaging their Fourier transform magnitude squared is involved. Once this is done, simulation of a new circuit involves only the generation of an appropriate number of noise vectors.

The method described in this section enables transient simulation of circuit that contains white noise sources. For other types of noise, e.g. $1/f$, shaping filters could be used to modify the white sequence (Fig. 6.15(a)) appropriately. Also, the power spectral density shown in Fig. 6.16(b) could be made "whiter" by using a filter with an appropriate frequency response. However, this would further increase the computations required.

6.3.2 LPTV simulations

As mentioned in section 6.2.2, the noise equivalent circuits for companding filters reduce to LPTV systems whose noise can be computed using the response of the system to delayed delta functions at the source of noise.

Another class of LPTV systems are mixers used for frequency translation in radios. Their periodic nature arises from the local oscillator (frequency translating) signal being periodic. A method for simulating noise in mixers using a SPICE-like simulator is described in [64]. In this approach, the response at the output of the mixer $h_k(t, \tau)$ to delayed delta functions $\delta(t - \tau)$ from a particular noise source n_k in the circuit is simulated. The response from one noise source can be simulated at a time. $h_k(t, \tau)$, which is a function of two variables and is periodic in τ, is approximated by a finite number of functions $h_{kj}(t, \tau_j)$

of one variable t with a fixed τ_j. τ_j is incremented in steps from zero to T, where T is the period of the LPTV system. Triangular pulses of narrow width are used to approximate the impulses. Each simulation results in a function $h_{kj}(t, \tau_j)$ for a particular value τ_j[2]. An approximate form of $h_k(t, \tau)$ is thus determined. This can be used in (6.27), with the integral replaced by a summation, to determine the output noise due to the noise source n_k. The process is repeated for all noise sources in the circuit.

Compared to the method described in the previous section, the individual transients to be simulated in this method are much shorter. However, a large number of them[2], for different values of τ_j and different noise sources in the circuit, need to separately simulated. Still, for a complicated circuit, it would be less time consuming than the brute force method.

6.3.3 Traditional ".AC" simulation

While ".AC" simulation in SPICE which computes noise by linearizing the circuit around the dc operating point is not applicable to the most general case of companding, it is nonetheless useful in several special cases. The class-A log-domain filter referred to in section 6.2 is an example. Although internally nonlinear, since the signal is restricted to be less than the bias, the nonlinearities are not greatly exercised. The difference between the noise in quiescent condition, which is what is computed by a ".AC" simulation and the noise when the signal swings as much as the bias is less than 2 dB [32, 33]. As a corollary, noise in dynamically biased class-A log-domain filters can be computed in this manner as well, by setting the bias point to various values. The result so obtained for the dynamically biased filter in Fig. 2.21 is shown in Fig. 2.23 along with the result from transient noise anal-

[2] $h_{kj}(t, \tau_j)$ for m values of τ_j can be determined either from m separate simulations or a single long simulation with a train of m impulses. The interval between the impulses is made equal to an integer multiple of T plus the increment in τ_j and must be long enough for the impulse response to die out sufficiently.

ysis described in section 6.3.1. The negligible difference between the two is evident from the plots. In general, this principle is applicable to syllabic companding filters where the circuit parameters change slowly with signal. It would be possible in most cases to "hold" the filter in one particular state in which it can be treated as a linear and time invariant system and apply traditional ".AC" analysis. Phenomena like envelope transient noise [3], however, are caused by the dynamics of syllabic companding and cannot be determined in such a fashion. The dependence of noise on the signal has to be taken care of in the simulation.

Chapter 7

Design of the Prototype Chips

7.1 Introduction

This chapter presents the design of several prototype chips used to evaluate the ideas presented in previous chapters. The technology available was a $0.25\,\mu$m BiCMOS technology from Lucent Technologies which provided nMOS and pMOS transistors with a drawn minimum length $0.25\,\mu$m and threshold voltages of 0.6 V and 0.9 V respectively and bipolar transistors with minimum emitter area of $(0.4\,\mu\text{m})^2$ and an Early voltage of about 12 V. Metal-metal capacitors with a density of $0.8\,\text{fF}/\,\mu\text{m}^2$ were also available in this technology.

To test the ideas presented in section 3.2, a third-order pseudo-differential Butterworth filter with a -3 dB frequency of 1 MHz and a peak detector to generate the dynamic bias were designed. Sections 7.2 to 7.5 deal with various aspects of this design.

Section 7.6 discusses the design of a second-order filter using enhanced lateral bipolar transistors and pMOS accumulation capacitors which is used to test the feasibility of log-domain filters in a pure CMOS technology as described in section 5.2.

The results of simulation of each of the designs are described in the respective sections.

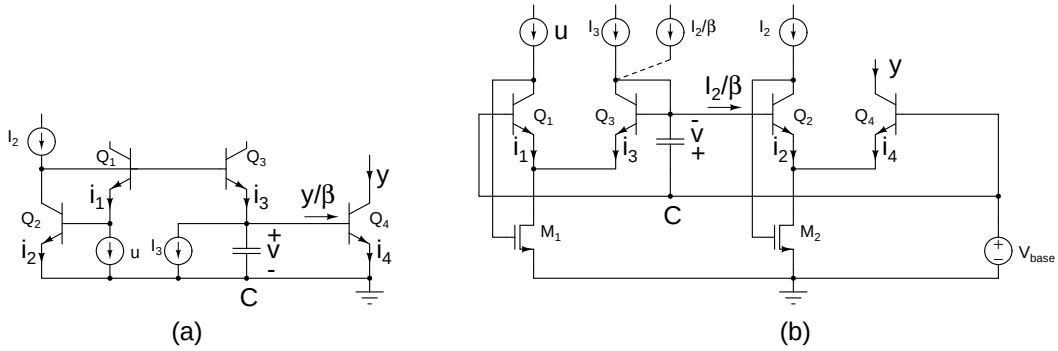

Figure 7.1: First order log-domain filter topologies (a) From [6], (b) From [19].

7.2 Third-order Butterworth ladder filter

The first order log-domain filters in Figs. 2.13 and 2.16 are redrawn in Fig. 7.1. As mentioned in section 2.2.1, the base current of Q_4 introduces nonlinearities in the topology in Fig. 7.1(a). With dynamic biasing, the collector current y in the output transistor Q_4 varies over a wide range and can be very large. In that case, the effect of the base current is particularly disastrous. Therefore, the topology in Fig. 7.1(b) [19, 20] which is immune to base currents is chosen for the dynamically biased filter in this work. The current source I_2/β used to compensate the reduction in pole frequency due to the base current of Q_2 in Fig. 7.1(b) is not used in our design. In the BiCMOS technology used for the current design, the current gain β of the bipolar transistors is large enough for this change in pole frequency to be tolerable.

7.2.1 Filter synthesis

The first-order log-domain filter in Fig. 7.1(b) is redrawn in Fig. 7.2(a). Here onwards the log-domain filters will be drawn in this fashion without the feedback loops to avoid clutter. The presence of feedback loops will be indicated by a dashed line from the collector.

Fig. 7.2(b) shows a cascade of two first order log-domain filters. The output y_1 of the first filter is mirrored using a current mirror and fed to the second filter as the input

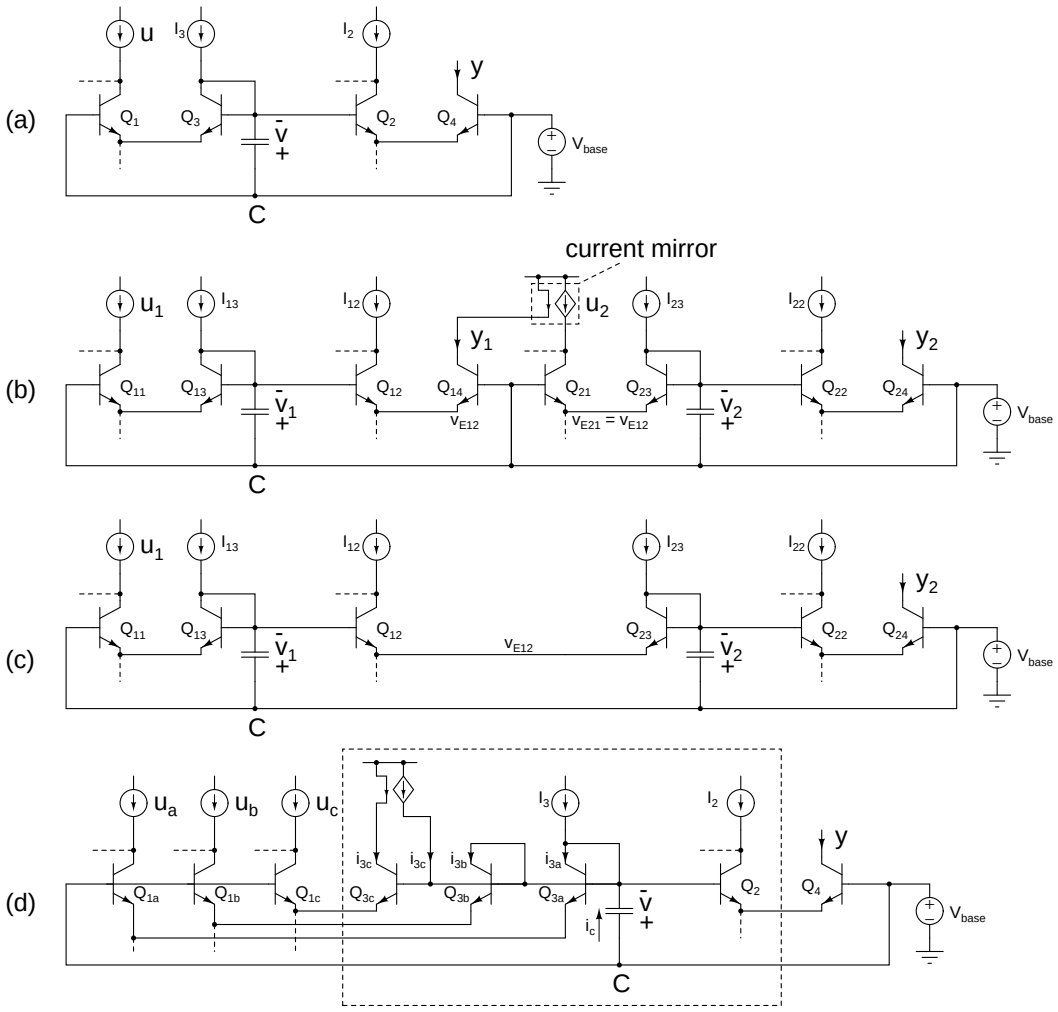

Figure 7.2: (a) First order log-domain filter in Fig. 2.16, (b) Cascading two stages, (c) Eliminating redundancies in (b), (d) Modification to accept multiple inputs with positive and negative weights.

u_2. Q_{14} converts the voltage v_{E12} into a current through an exponential nonlinearity. This current is converted back to a voltage v_{E21} by Q_{21}. Since the transistors Q_{14} and Q_{21} carry equal current and have their bases tied together, their emitter voltages v_{E12} and v_{E21} are identical. Therefore a cascade of the two stages can be realized in the voltage domain as shown in Fig. 7.2(c) by eliminating Q_{14} and Q_{21}. Thus a cascade of log-domain filters can be realized without voltage to current and current to voltage conversion at each interface.

Fig. 7.2(d) shows a modification of Fig. 7.2(a) to accommodate multiple inputs with

both positive and negative weights ([19, 20]). The transistor pair $Q_{1,3}$ in Fig. 7.2(a) is replaced by three pairs $Q_{1a,3a}$, $Q_{1b,3b}$, and $Q_{1c,3c}$. Inputs u_a, u_b, and u_c are fed to Q_{1a}, Q_{1b}, and Q_{1c} respectively. Q_{3a} and Q_{3b} are diode connected. The collector areas of Q_{3a}, Q_{3b}, and Q_{3c} are A_{3a}, A_{3b}, and A_{3c} respectively. These areas determine the gain of the filter from the respective inputs as will be shown below. The rest of the transistors in the circuit have unit areas. The collector current of Q_{3c} is fed back to its base through a current mirror to accomplish a sign inversion, as will be seen below. This circuit can be analyzed by writing the equations for the products of currents around the translinear loops (see (2.18), (2.19) and the associated discussion).

$$i_{3a} = \frac{u_a I_2 A_{3a}}{y} \qquad (Q_{1a}\text{-}Q_{3a}\text{-}Q_2\text{-}Q_4) \tag{7.1}$$

$$i_{3b} = \frac{u_b I_2 A_{3b}}{y} \qquad (Q_{1b}\text{-}Q_{3a}\text{-}Q_2\text{-}Q_4) \tag{7.2}$$

$$i_{3c} = \frac{u_c I_2 A_{3c}}{y} \qquad (Q_{1c}\text{-}Q_{3a}\text{-}Q_2\text{-}Q_4) \tag{7.3}$$

$$i_c = i_{3a} + i_{3b} - i_{3c} - I_3 \qquad \text{(total current into } C\text{)} \tag{7.4}$$

The capacitor current can be related to the output current using the following equations.

$$\begin{aligned} v &= V_{BE4} - V_{BE2} \\ &= V_t \ln \frac{y}{I_2} \end{aligned} \tag{7.5}$$

$$\begin{aligned} i_c &= C\frac{dv}{dt} \\ &= \frac{CV_t}{y}\frac{dy}{dt} \end{aligned} \tag{7.6}$$

(7.4) and (7.6) can be combined to obtain

$$\frac{CV_t}{y}\frac{dy}{dt} = \frac{u_a I_2 A_{3a}}{y} + \frac{u_b I_2 A_{3b}}{y} - \frac{u_c I_2 A_{3c}}{y} - I_3 \tag{7.7}$$

$$\frac{dy}{dt} = -\frac{I_3}{CV_t} + \frac{A_{3a} I_2}{CV_t} u_a + \frac{A_{3b} I_2}{CV_t} u_b - \frac{A_{3c} I_2}{CV_t} u_c \tag{7.8}$$

(7.8) is the differential equation of a first-order filter with an output y and inputs u_a, u_b and u_c. The relative weights of the inputs are determined by the areas A_{3a}, A_{3b} and A_{3c}. u_c is filtered with a negative weight. More inputs can be added to the filter in a similar fashion. Thus the circuit in Fig. 7.2(d) is seen to be a first-order stage with multiple inputs whose weights can be set to arbitrary positive or negative values. This stage can be used to synthesize higher order filters. As in Fig. 7.2(c) the input and output transistors can be omitted when cascading several stages. The circuit enclosed within the dashed lines forms the basic building block of a higher-order log domain filter.

The doubly terminated LC ladder filter in Fig. 7.3(a) is used as the prototype for the desired Butterworth filter. With the values shown, the Butterworth filter has a -3 dB bandwidth of 1 rad/s. The ladder filter is converted to the block diagram in Fig. 7.3(b) which uses two lossy integrators and a lossless integrator. ω_p in the figure is the desired bandwidth of the filter in rad/s. The LC ladder prototype has an attenuation of 6 dB in the passband. A gain of 2 is introduced at the input so that a passband gain of unity is realized. The three stages are marked on the figure. The first is a lossy integrator with an inverting and a non-inverting input. The second is a lossless integrator with an inverting and a non-inverting input. The third is a lossy integrator with a single non-inverting input.

The log-domain version[1] of this filter is shown in Fig. 7.3(c). The first-order log-domain filter with multiple inputs shown in Fig. 7.2(d) is used for each of the stages. All transistors are of unit areas except Q_2 whose area is doubled to implement the input gain of 2 mentioned above (see (7.8)). Note that the damping current is absent from the second stage. The three stages are cascaded in a manner similar to Fig. 7.2(c). $I_0 = 5\,\mu$A and $C_1 = C_3 = 30$ pF, $C_2 = 60$ pF are used to obtain the desired cutoff frequency of 1 MHz for

[1] Details of the procedure for the synthesis of log-domain filters can be found in [7, 14, 27].

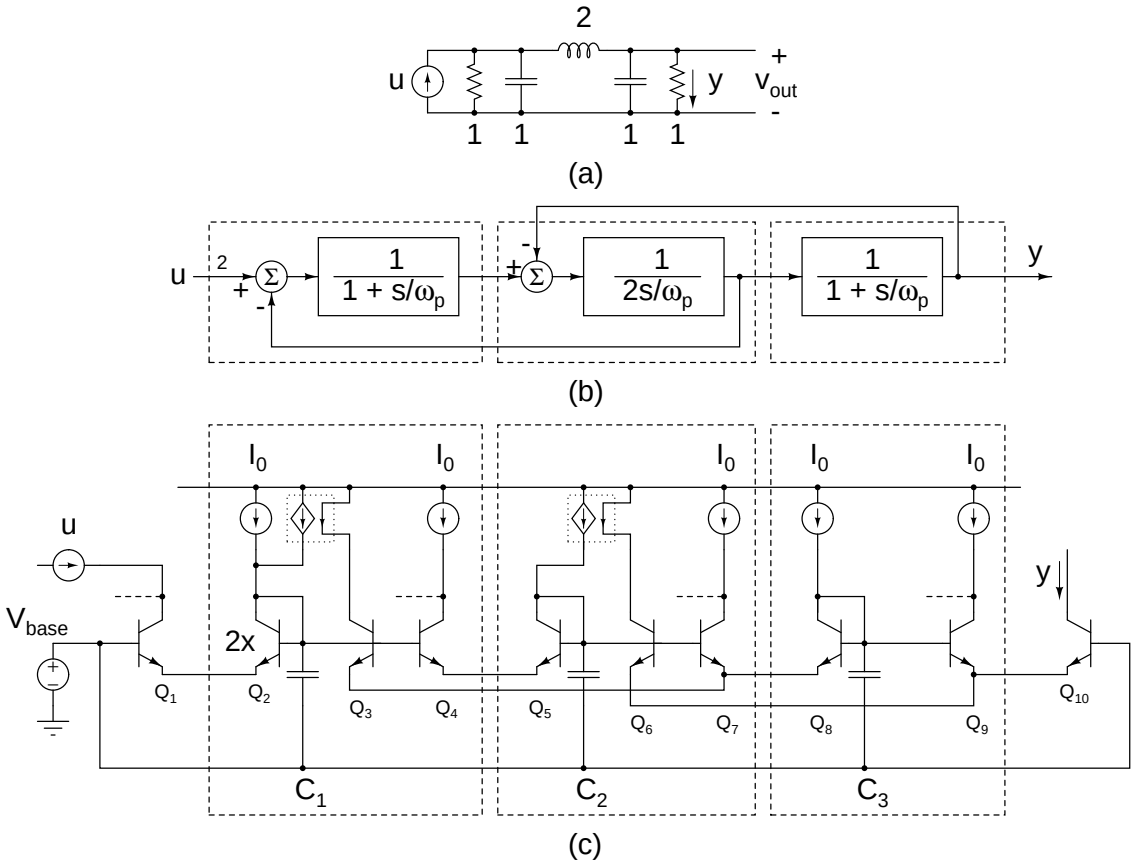

Figure 7.3: Third-order Butterworth filter: (a) RLC prototype, (b) Block diagram using first order stages, (c) Log-domain realization.

the Butterworth response.

In the initial stages of the design, it was envisioned that the input currents would vary in a large range of about 50 nA to 500 μA whose center (on a log scale) is the internal bias current $I_0 = 5\,\mu$A. A supply voltage of 2.5 V was chosen for the circuit. Since the chief aim was the demonstration of dynamic biasing, no attempt was made to minimize the operating voltage of the circuits.

In order to distinguish the dynamic bias applied to the input from the internal bias current I_0 (Fig. 7.3) used to set the bandwidth of the filter, here onwards, the latter is referred to as the tuning current.

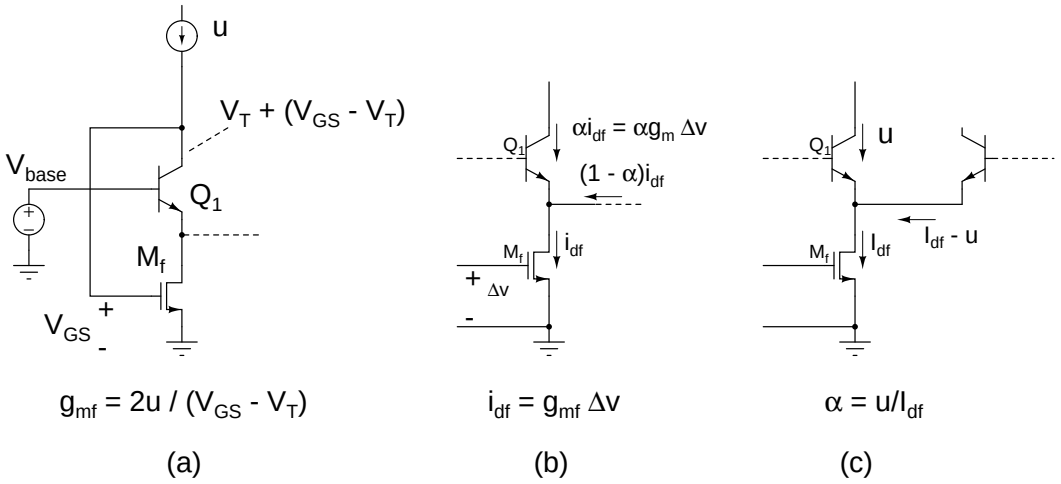

Figure 7.4: Feedback to establish the collector currents (a) From [19], (b, c) Small signal picture of (a).

7.2.2 Feedback circuit used to establish the collector currents

Fig. 7.4(a) shows the feedback circuit used in [19] (Fig. 2.15) to force the input current u into the collector of Q_1 and tuning currents I_0 into the collectors of Q_4, Q_7 and Q_9 in Fig. 7.3(c). The following discussion uses the input stage with Q_1 as an example, but it applies equally well to the other cases.

In equilibrium, M_f's drain current equals the sum of emitter currents of all the transistors connected to it. Its gate voltage is adjusted appropriately by the feedback loop.

Fig. 7.4(b) shows the circuit in Fig. 7.4(a) with the feedback loop opened. The transconductance of this circuit between the gate of M_f and the collector current of Q_1 is the small signal conductance presented to the input u by the circuit in Fig. 7.4(a). If α is the fraction of the small signal drain current i_d of M_f that flows through the transistor Q_1 and g_{mf} is the transconductance of M_f, the small signal transconductance between gate of M_f and the collector of Q_1 is αg_{mf}. If bipolar transistors connected to the drain of M_f as shown in Fig. 7.4(c), the fraction of the drain current of M_f that flows through Q_1 is the same for

large and small signals[2]. Thus the conductance g_{in} presented to the input u by the circuit in Fig. 7.4(a) is given by

$$g_{in} = \frac{u}{I_{df}} g_{mf} \tag{7.9}$$

where I_{df} is the total (large signal) drain current of M_f. Assuming that M_f operates as a square law device in saturation, its transconductance is [45]

$$g_{mf} = \frac{2I_{df}}{V_{GS} - V_T} \tag{7.10}$$

and the transconductance g_{in} presented to the input is

$$g_{in} = \frac{2u}{V_{GS} - V_T} \tag{7.11}$$

The voltage at the input node is V_{GS} and is the sum of the voltages across Q_1 and M_f. If this voltage is too small, Q_1 may be pushed into saturation where its log conformity degrades. This is likely to happen with short-channel technologies where V_T is small. The only way to increase V_{GS} is to reduce the W/L ratio of M_f which reduces g_{mf}. Also, increasing its V_{GS} for the same drain current would tend to push M_f into the triode region, further reducing g_{mf}. The effect of reducing g_{mf} is to increase the signal voltage swing at the input node. As will be shown later, the finite Early voltage of the transistors is one of the main sources of distortion in this log-domain filter. In view of this, a small voltage swing at the collector node and thus a large g_{mf} is desirable.

In a dynamically biased filter, the currents in the input and output transistors vary over a wide range. This means that M_f's current varies over a large range. Consequently, its V_{GS} would experience wide variations and it is likely that either or both of Q_1 and M_f would be operating in undesirable regions at the extremes of input current variations.

[2]The small signal conductance of bipolar transistors is proportional to their large signal collector current, owing to their exponential characteristic.

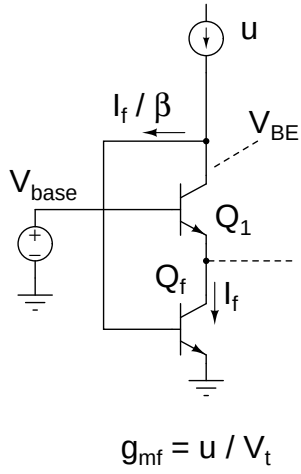

Figure 7.5: Feedback circuit used in [22] to establish the collector currents.

It may be possible in a technology with an appropriate threshold voltage to maintain Q_f in the active region and size M_f so large that it would be operating in the subthreshold region even for the largest input current. But this would imply a large gate capacitance at the input node which would be unacceptable at the frequencies considered in this design. A large capacitance has the following undesirable effects on the operation of this circuit.

1. The pole at the input node moves to lower frequencies and causes a deviation in the small signal frequency response of the filter.

2. The input capacitance is across a nonlinear input conductance. This causes distortion in the current that flows through Q_1.

[22] presents a modified version of this circuit in which the MOS transistor M_f is replaced by a bipolar transistor Q_f as shown in Fig. 7.5. The input conductance in this case is larger and the input voltage variations are limited to a few V_t. But the voltage at the input node is V_{BE} and could be smaller than in the previous case. This means that Q_1 or Q_f could enter saturation region if their $V_{CE,SAT}$ is larger than about 0.3 V.

A serious problem with this circuit is the base current of Q_f which subtracts from the

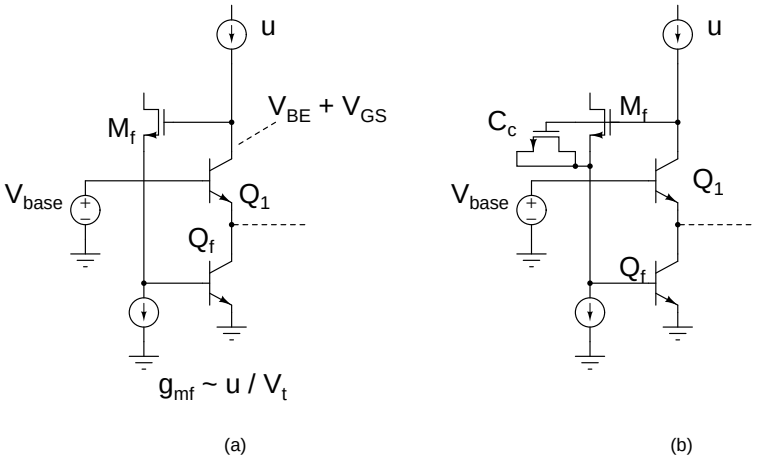

Figure 7.6: (a) Proposed feedback circuit for establishing collector currents in Fig. 7.3(c), (b) With frequency compensation.

input current flowing into Q_1 as shown in the figure. This current is non-linearly related to the input and consequently results in distortion of the signal flowing into Q_1's collector. This circuit too is therefore unsuitable for large dynamic range applications.

The circuit in Fig. 7.6(a) remedies the above problems. A source follower M_f is used to drive the bipolar transistor Q_f. Unlike the circuit in Fig. 7.5, no current is drawn by the feedback network at low frequencies. The quiescent voltage at the input node is $V_{BE}+V_{GS}$ ensuring that Q_1 stays in its active region. The source follower attenuates to some extent due to the body effect of M_f. But the input conductance is still close to that in Fig. 7.5.

Insertion of the source follower adds another pole to the loop gain function and hence the circuit in Fig. 7.6(a) needs frequency compensation to be assured of stability over the entire range of input currents. A capacitor (nMOS transistor in inversion region) is connected between the gate and source of M_f as shown in Fig. 7.6(b) to provide a feed-forward path at high frequencies and eliminate the phase lag due to the source follower. From simulations, this circuit is determined to be stable over the desired range of input current. This feedback arrangement is used around Q_1, Q_4, Q_7 and Q_9 in Fig. 7.3(c).

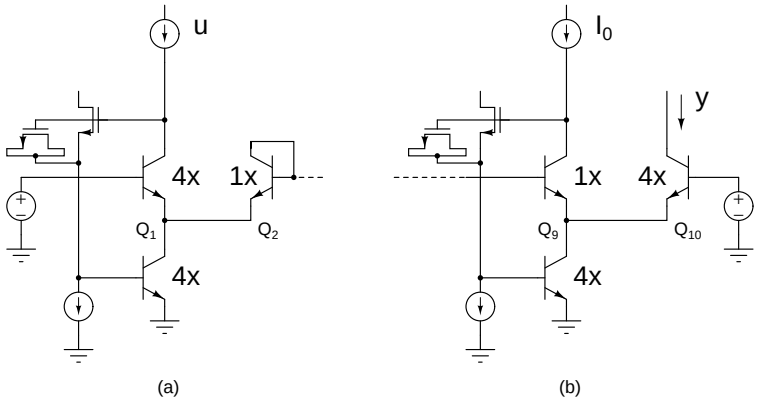

Figure 7.7: Transistor sizing at the (a) input, (b) output.

7.2.3 Transistor sizing at the input and the output

The exponential relation between the base-emitter voltage and collector current of a transistor deteriorates at high current densities due to increased voltage drop across parasitic series resistances and high-level injection effects. At very low current densities—i.e. with a very large transistor for a given current—the frequency response of a transistor deteriorates due to larger parasitic capacitances. The transistors in the log domain filters should be sized to strike a compromise between these two conflicting requirements. As mentioned above, the currents in the input and the output transistors (Q_1 and Q_{10} in Fig. 7.3(c)) can be much larger than the currents in the transistors comprising the filter core (Q_{2-9}). For this reason, the input and output transistors in the filter are made four times larger than the transistors in the filter's core as shown in Fig. 7.7. Increasing the size of both the input and the output transistors in the same proportion leaves the transfer function of the filter unaffected.

7.2.4 Feedback and feed-forward paths in the filter

In the Butterworth filter of Fig. 7.3, emitter followers (Q_4, Q_7 and Q_9) are used for inter-stage coupling. Of these, Q_7 and Q_9 drive two transistors each; one in the forward path

and another in the feedback path. These paths are shown using thick lines in Fig. 7.8(a). Consider Q_9 in Fig. 7.8(a): This drives the output transistor Q_{10} and a feedback transistor Q_6. The current in Q_{10} varies over a wide range due to dynamic biasing. This means that the emitter follower Q_9 is loaded by a widely varying impedance causing its input-output relationship (both gain and phase) to change as the output current changes. Thus the feedback branch also suffers this variation in gain and phase. Any phase change in the feedback path has a very severe effect on the overall frequency response. This was confirmed through simulations. Using separate transistors for the feedback and feed-forward paths rectified this problem and the frequency response of the filter remained invariant with changing bias current at the input. The resulting circuit is shown in Fig. 7.8(b) where two separate emitter followers Q_{9a} and Q_{9b} are used. Emitter follower Q_{9a} still suffers from a changing load and consequent gain and phase changes, but these are now outside the filter's feedback loop and do not affect the frequency response as much. The small gain and phase variations are added to the frequency response. The problem is not as severe with Q_7 driving Q_8 and Q_3 (Fig. 7.8(a)), both of which are inside the filter's core and carry currents that do not change with dynamic biasing. Nevertheless, this device was separated as well into Q_{7a} and Q_{7b} as shown in Fig. 7.8(b). A slight improvement in the behavior of frequency response with changing bias was seen as a result.

7.2.5 Effect of nonidealities

To determine the chief impediments to good performance of this filter, the single ended circuit in Fig. 7.3(b) was simulated[3]. Idealized components with various non-idealities added one at a time were used. Although the series resistance in the emitter causes the transistor to deviate from the ideal exponential law, it was found that the distortion in-

[3] At the time of these simulations, the parameters of the technology described in section 7.1 were not available and device models from a different bipolar technology were used.

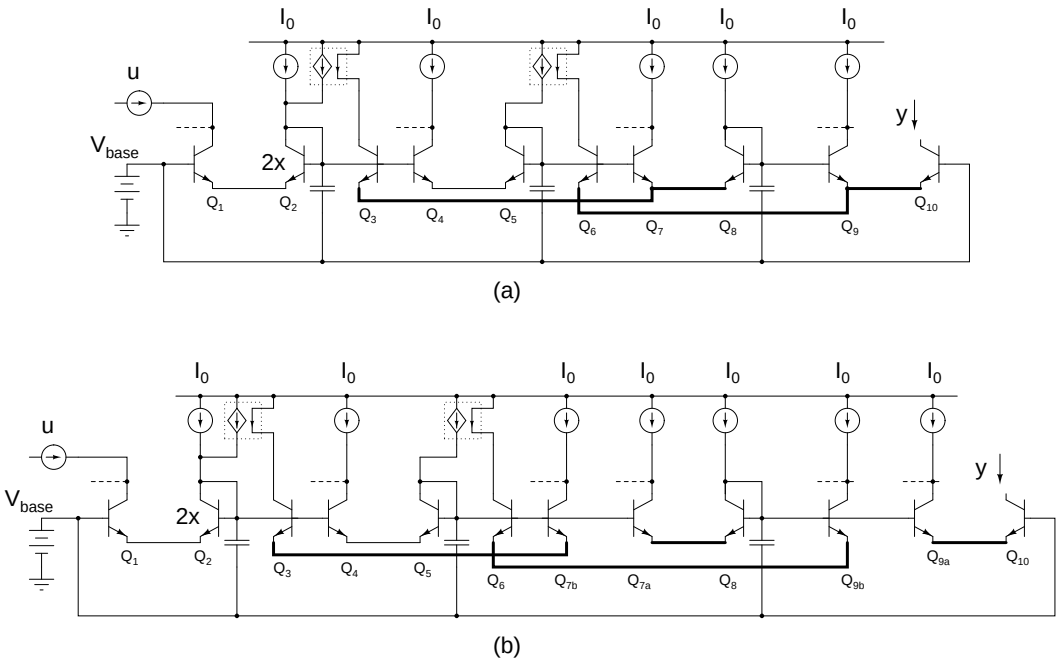

Figure 7.8: Separating the feedback and the feed-forward paths.

troduced by practical values of emitter series resistance was insignificant. The finite Early voltage (V_A) of the transistors was the most important contributor to distortion. Fig. 7.9 plots the second and the third harmonic distortion of the single ended circuit of Fig. 7.3(c) vs. bias current for Early voltages of 25 V and 10 V. The input tone used had a frequency of one third the bandwidth and an amplitude that was 90% of the bias. The increase in distortion with decreasing Early voltage is evident. Also notable is the deterioration in the distortion performance for large values of the input bias current. The reason for this is the large voltage swing that occurs at the collectors of Q_1 and Q_9 (in Fig. 7.3(c)) when large signal swings accompany a large input bias. For this reason, it is necessary to have as large a transconductance as possible in the feedback loop used to force the input current u into Q_1. Holding the collector voltages of these transistors constant by cascoding is a possibility, but this requires a larger supply voltage and further complicates the circuit, and therefore, it was not used in this implementation.

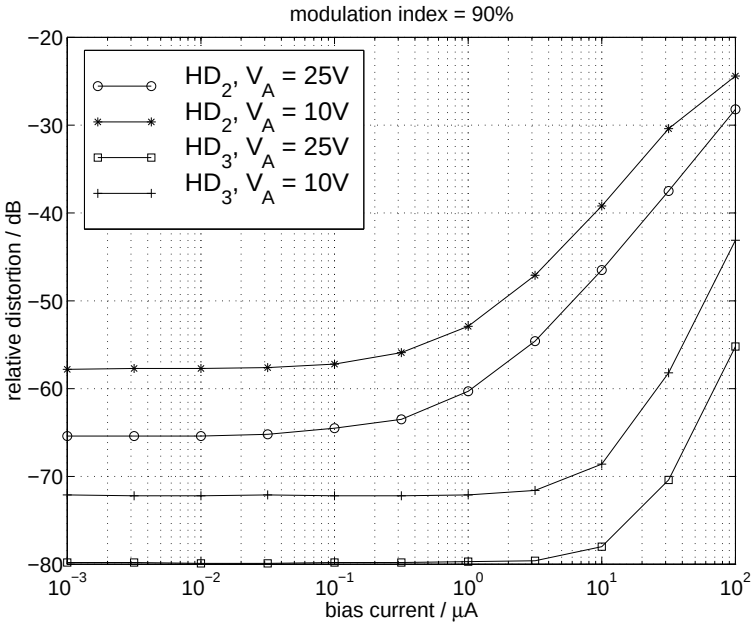

Figure 7.9: Distortion vs. bias for a constant modulation index and two Early voltages.

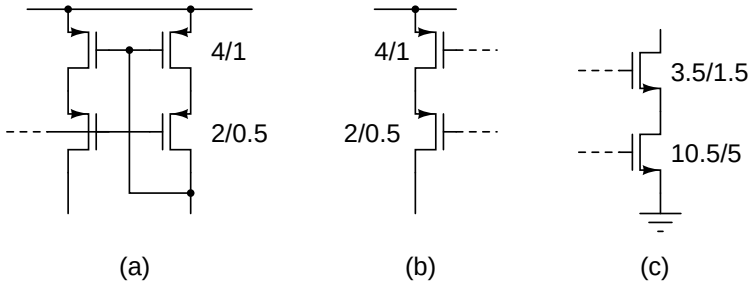

Figure 7.10: (a) Current mirror, (b) Current source, (c) Current sink.

It is clear that the finite output conductances of the current sources and the current mirrors in Fig. 7.3(c) affect the distortion in the same manner as finite output conductances of the transistors Q_{1-10}. Therefore, cascode current sources and mirrors as shown in Fig. 7.10 were used.

Aside from the noise contributed by the bipolar transistors which form the core of the filter, extra noise is contributed by the current sources in Fig. 7.3(c). The noise from the current sources can be particularly significant at low frequencies due to $1/f$ noise. To minimize their thermal noise contribution, these current sources are operated with as large

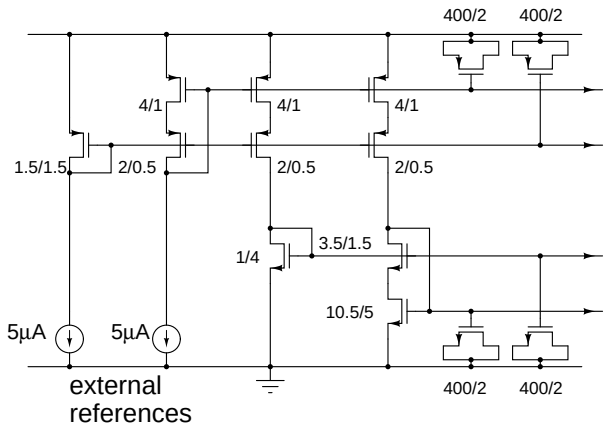

Figure 7.11: Bias circuitry.

a $V_{GS} - V_T$ as possible [45] within the power supply voltage available. $V_{GS} - V_T \approx 200\,\text{mV}$ was used in this design[4]. To minimize $1/f$ noise, long channel transistors [45] are used for the current sources as shown in Fig. 7.10. Due to their large 1/f noise, nMOS transistors with very long channels are used in the current sources.

7.2.6 Bias generation

Fig. 7.11 shows the circuit used to generate the bias voltages for the current sources and mirrors. The bias lines are bypassed for high frequencies using the gate capacitance of large transistors operating in strong inversion.

7.2.7 Pseudo differential version of the filter

For pseudo-differential operation, the filter in Fig. 7.8(b) is duplicated and laid out symmetrically.

7.2.8 Automatic tuning of the filter

The bandwidth of the third-order Butterworth filter discussed in the previous sections is given by $I_0/(C_1V_t)$ where I_0 and C_1 are the tuning current and the capacitor used in

[4] Assuming square law operation in saturation, the transconductance of a MOSFET with a drain current I_d is $2I_d/V_{GS} - V_T$ and its drain current noise is $8kTg_m/3$ [45].

Fig. 7.3(c) respectively, and V_t is the thermal voltage. The tuning current I_0 is the electronic control available to change the bandwidth as desired.

Two sources of variation of bandwidth of the filter are apparent:

1. The temperature dependence of the thermal voltage V_t causes the bandwidth to decrease with increasing temperature[5]. This effect can be countered by having I_0 vary in proportion to the absolute temperature (PTAT).

2. Process variations in the values of the capacitors used in the filter cause a corresponding variation in the bandwidth. This can be corrected either by using a PTAT source in combination with trimming for the correct bandwidth or using more sophisticated automatic tuning techniques [35].

Since the chief aim of this design was the demonstration of the dynamic range capabilities of dynamically-biased filters, none of these bandwidth correction schemes was incorporated in the fabricated chip.

7.2.9 Simulation results

Various simulations were carried out on the pseudo-differential third-order Butterworth filter including the padframe. Most of the simulations were carried out with the bias currents in the range[6] from 50 nA to 500 μA. However, 500 μA is not the largest bias current that can be applied and it is expected that the experimental prototype will function even beyond this bias level. As the bias current at the input is increased, the voltages at the collectors of Q_1 and Q_{10} in Fig. 7.7(a) and Fig. 7.7(b) respectively increase due to the following

[5] $V_t = kT/q$ where k, T, and q are the Boltzmann constant, the absolute temperature, and the electron charge respectively.

[6] Due to the limited time available in a summer term for the design and layout of the circuits described in this chapter, the simulations carried out were limited in their scope. The simulations were extensive enough to verify the satisfactory performance in the expected regions, but not to prcisely determine the boundaries of performance of the chips.

reasons.

1. As the current carried by the feedback transistor used at the bottom (Fig. 7.7) increases, its base-emitter voltage increases.

2. The V_{GS} of the source follower transistor increases due to (i) the body effect and (ii) an increase in the base current supplied to the base of the bottom transistor as the bias current is varied over a wide range.

A sufficient increase in the voltage at the collector of Q_{10} causes the current source feeding the collector to operate in the triode region. This results in a reduced current I_0 fed to the collector which in turn results in a change in the frequency response of the filter. With the transistor sizes used in this design, it is expected that this limiting will occur at a few milliamperes of bias current.

Frequency response

Fig. 7.12 shows the simulated magnitude response of the filter with input bias currents varying from 500 nA to 500 μA. In these simulations, the current I_0 in Fig. 7.3(c) and Fig. 7.8 was set to 5 μA. The response stays practically constant over such a wide range of input bias currents. Below 500 nA, the response exhibits slight peaking. This is because the frequency response of the feedback loop (Fig. 7.6) used to force the input current is poor at small bias currents. Fig. 7.13 shows the passband detail of the magnitude response. The variation in dc gain over the entire range is about 0.6 dB. Note that the curves in Fig. 7.13, with the exception of the topmost one, are equally separated in the vertical direction fromeach other. This implies that the bandwidth of the filter is unaltered with changes in the bias current.

One possible reason for this gain variation can be understood by examining Fig. 7.7.

Consider Fig. 7.7(a). The collector-emitter voltage V_{CE1} of Q_1 increases with increasing input bias current. The base-emitter voltage V_{BE1} of transistor Q_1 is given by the following relation.

$$V_{BE1} = V_t \ln\left(\frac{u}{I_s(1 + V_{CE1}/V_A)}\right) \tag{7.12}$$

where V_A is the Early voltage of the bipolar transistor. For a given input u, V_{BE1} decreases as V_{CE1} increases. Thus, the effect of finite Early voltage is an apparent reduction in the input signal by a factor $(1 + V_{CE1}/V_A)$, or equivalently, a reduction in the gain of the filter. As V_{CE1} increases with increasing input bias current, the gain of the filter decreases. A similar effect occurs in the output stage shown in Fig. 7.7(b) due to an increase in the collector-emitter voltage of the transistor Q_9.

Assuming an increase ΔV_{CE} of 300 mV in the collector-emitter voltages of Q_1 and Q_9, and an Early voltage of 12 V the reduction in the gain of the filter can be calculated as

$$\Delta gain = 20 \log \frac{1}{(1 + \Delta V_{CE}/V_A)^2} \tag{7.13}$$

$$= 0.43 \text{ dB} \tag{7.14}$$

The observed reduction in the gain is slightly larger than the calculated value.

While the variation of dc gain with the input bias current reflects a non-ideality, its effect on the performance of the dynamically biased filter is difficult to quantify. When used in association with a peak detector, the bias at the input varies slowly due to the long time constant of the peak detector. There is no gain variation within a cycle of the input signal, and hence, no distortion. The gain variation in such cases can be eliminated through automatic gain tuning, if desired. The gain variation would manifest itself as nonlinear components in the output during envelope transients.

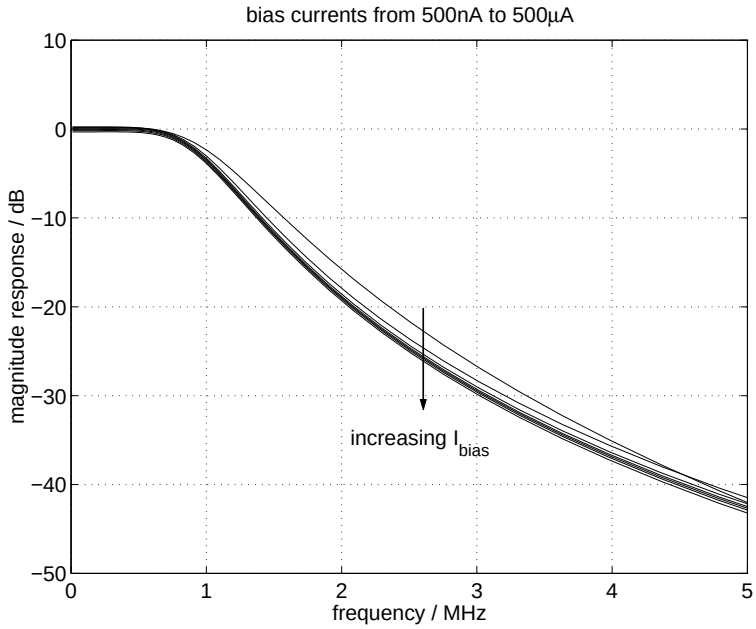

Figure 7.12: Frequency response with the bias current varied from 500 nA to $500\ \mu$A.

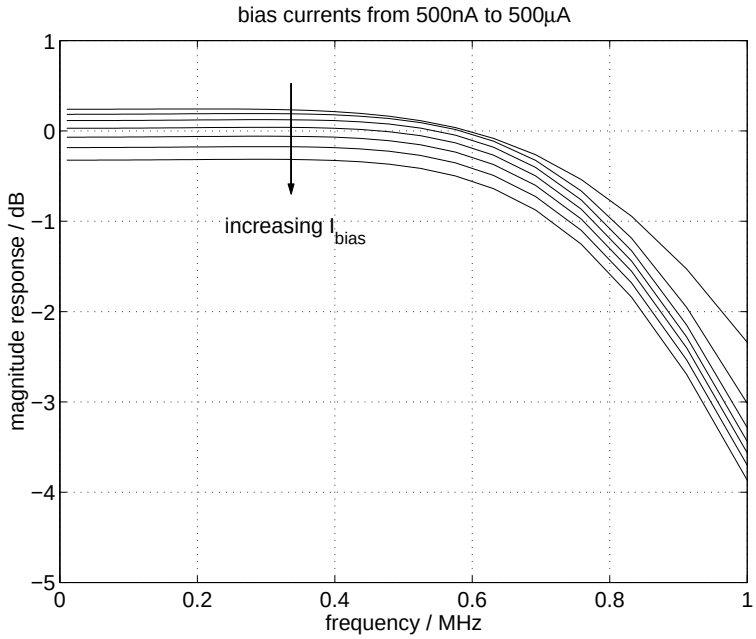

Figure 7.13: Passband detail of the frequency response with the bias current varied from 500 nA to $500\ \mu$A.

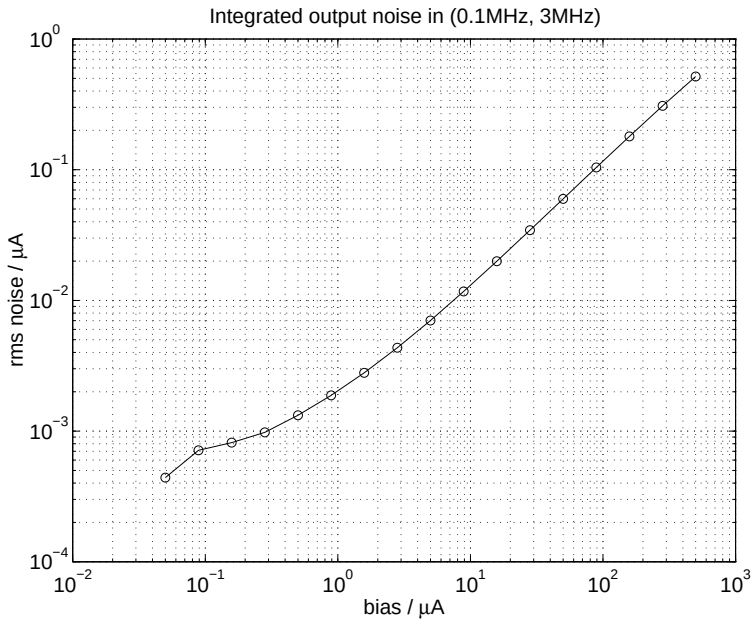

Figure 7.14: Output noise integrated in 100 kHz to 3 MHz band.

Output noise and signal to noise ratio

Fig. 7.14 shows the output noise of the filter integrated from 100 kHz to 3 MHz. The output noise increases as the square root of the bias for small bias currents and in proportion to the bias for large bias currents as expected (section 2.2.3, [30]).

The output signal to noise ratio, assuming a single ended input peak current that is 50% of the bias current is plotted in Fig. 7.15. Since the single ended input current is 50% of the input bias I_{bias}, the differential input peak is equal to I_{bias}. Therefore, the X-axis denotes both the bias current used at the input and the differential input peak current under this assumption. The S/N increases at 0.5 dB per dB of increase in the bias for small currents and saturates for large currents.

It was pointed out earlier that the filter's frequency response deteriorates to some extent below a bias current of 0.5 μA. Assuming this to be the smallest usable bias, and taking 500 μA as the upper limit[7], the dynamic range of the filter is 60 dB. In this 60 dB

[7] As mentioned earlier, simulations were done only up to a bias current of 500 μA. The conclusions that follow are based on this number.

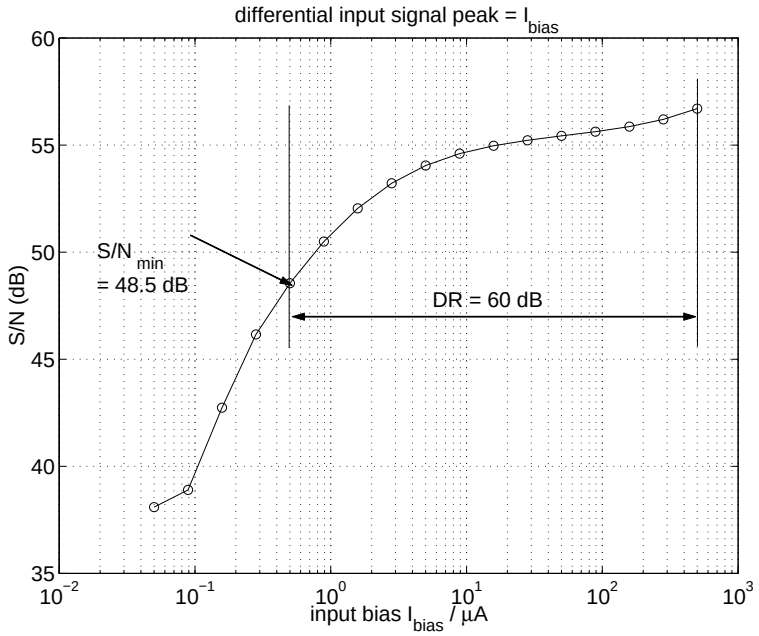

Figure 7.15: Signal to noise ratio with single ended signal peak at 50 % of the bias.

range, a S/N of 48.5 dB is maintained as can be seen from Fig. 7.15. A conventional linear filter would be required to have a peak S/N of $48.5 + 60 = 108.5$ dB to satisfy the same criterion!

Distortion

The distortion performance of the filter was simulated using harmonic balance. This is much less time consuming than taking Fourier transforms of long transient runs. A single tone input at a frequency of 312.5 kHz was used. Fig. 7.16 shows the third harmonic distortion as a function of the modulation index (ratio of peak signal current to bias current) for different values of input bias current. The distortion is worse at low bias currents due to the poor frequency response of the feedback loop used to force the input current (Fig. 7.6). For bias currents larger than 500 nA, the distortion is the same for a given modulation index regardless of the value of the bias current. As explained in section 3.2, the peak-peak voltage signal applied to the voltage mode filter embedded between the input and output

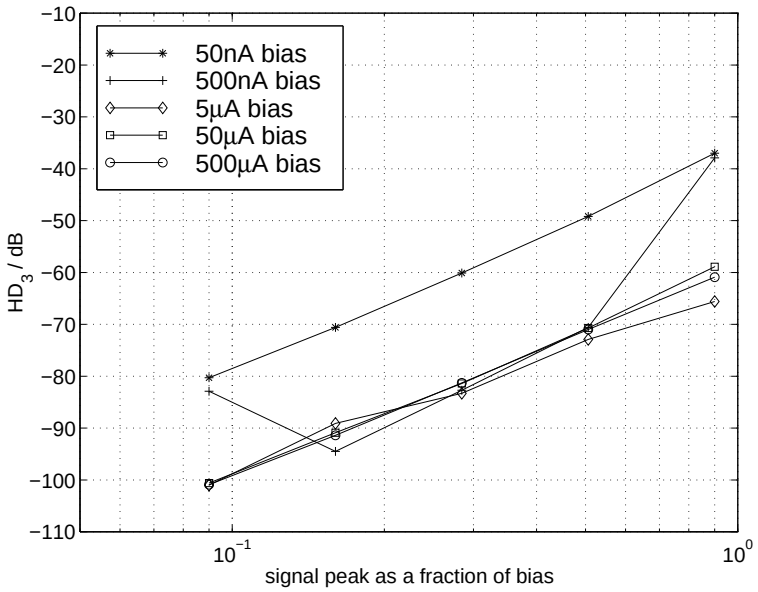

Figure 7.16: Third harmonic distortion as a function of bias.

"amplifiers" (Fig. 3.1(b)) is a function of the modulation index and does not depend on absolute values of the input bias or signal currents. The distortion relative to the fundamental is below 60 dB for the most part and is lower than the noise level. Fig. 7.17 shows the variation of distortion with input frequency when the single ended input signal peak is 50% of the bias current for various bias currents. It can be seen that the distortion increases slowly with frequency.

Power consumption

The filter consumes 250 μA when the input bias is set equal to the internal bias level of 5 μA. When the input bias is decreased to very small values, it drops to 210 μA. With a 500 μA bias at the input, the total current drawn from the supply is 2.23 mA.

Power consumption and dynamic range: comparison to passive filters

As described in section 2.1.1 and Appendix A, the power consumption of the passive RC filter varies with the input signal. In order to compare the pseudo-differential filter to the RC filter, the dependence of the power consumption (from the power supply voltage) of

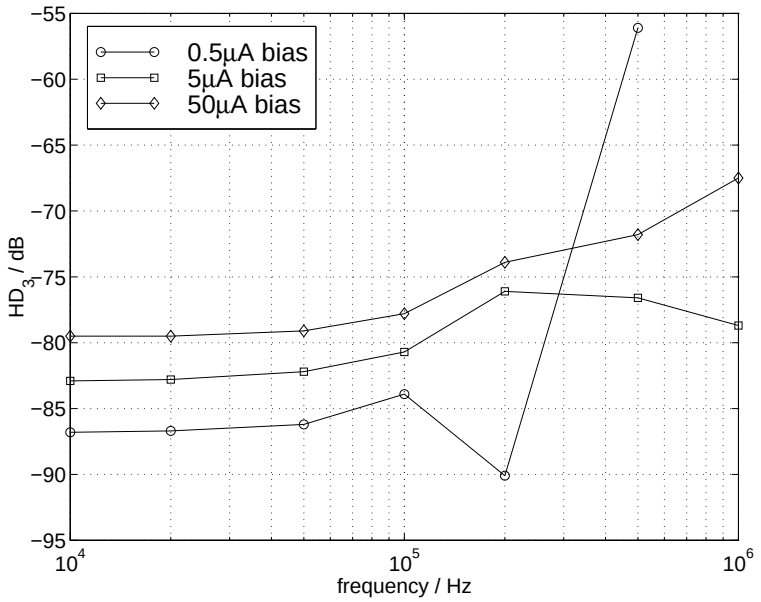

Figure 7.17: Third harmonic distortion as a function of frequency.

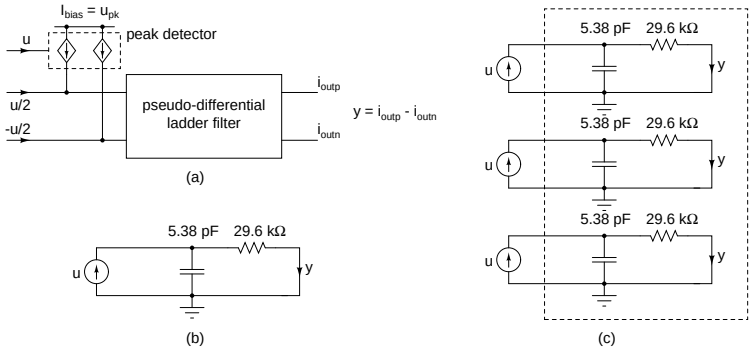

Figure 7.18: (a) Pseudo differential filter biased from an ideal peak detector, (b) Current mode RC filter with $S/N = 48.5$ dB for a 0.5 μA peak input, (c) Three first-order RC filters.

the former on the input signal has to be determined.

In a dynamically biased filter, the input bias I_{bias} is varied as function of the peak value of the input signal. Assume that the bias required for a given input signal is twice the peak value of that signal. Fig. 7.18(a) shows the pseudo-differential filter biased from an ideal peak detector. A bias I_{bias}, which is equal to the peak value of u, is added to the differential ac inputs $u/2$ and $-u/2$ and fed to the filter. The system in Fig. 7.18(a) with the ideal peak detector will be compared with the RC filter in this section.

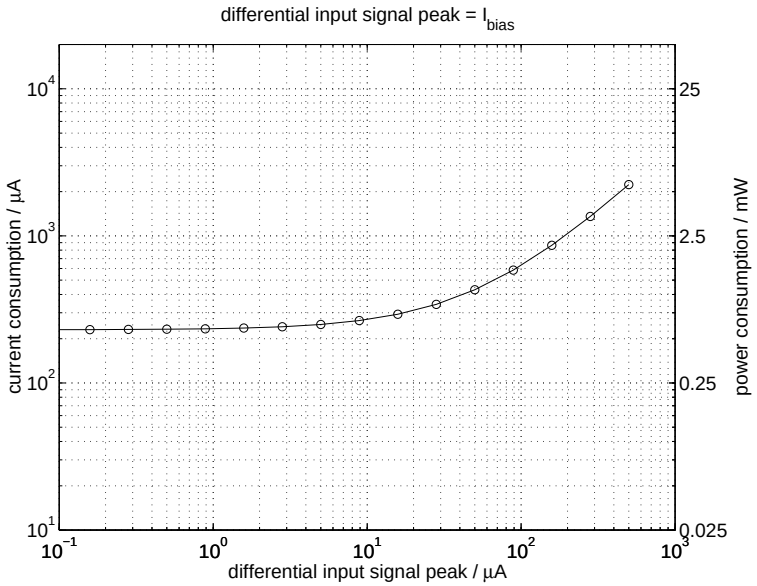

Figure 7.19: Power consumption of the pseudo-differential log-domain filter.

The left and right Y-axes of Fig. 7.19 respectively show the current and the power consumed by the Butterworth ladder filter as a function of the peak value of the input assuming the arrangement shown in Fig. 7.18(a). As with Fig. 7.15, the X-axis denotes both the bias current used at the input and the differential peak input current. This curve depicts the dependence of the power dissipation of the pseudo-differential filter on the input signal.

The signal to noise ratio of the pseudo-differential filter was shown in Fig. 7.15. The minimum signal to noise ratio was 48.5 dB assuming that the smallest input had a differential peak value of 0.5 μA (with an input bias I_{bias} of 0.5 μA). Fig. 7.18(b) shows a current mode first-order passive RC filter. The values of R and C are chosen such that:

1. The filter has a 1 MHz bandwidth, the same as the log-domain filter in consideration.

2. The output signal to noise ratio for a 0.5 μA peak input sinusoidal signal at 1 MHz is 48.5 dB, that same as that for the dynamically biased filter.

The output signal to noise ratio of both the pseudo-differential filter (Fig. 7.18(b)) and the

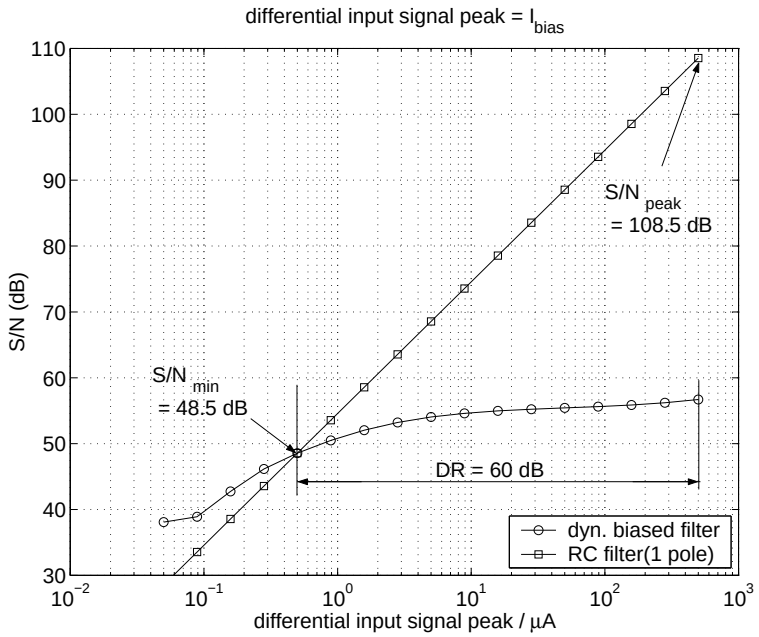

Figure 7.20: Signal to noise ratio of the log-domain and passive RC filters.

filter in Fig. 7.18(b) are plotted versus the input signal peak in Fig. 7.20. As expected, the signal to noise ratio of the passive RC filter increases linearly with the input signal whereas the signal to noise ratio of the log-domain filter remains nearly a constant. Assuming the minimum and the maximum input signal peaks to be 0.5 μA and 500 μA the usable dynamic range of both filters is 60 dB. At the top end of this dynamic range, the signal to noise ratio of the RC filter is 108.5 dB, which is wasteful, if the required S/N is only 48.5 dB.

The first order RC filter in Fig. 7.18(b) is repeated thrice as shown in Fig. 7.18(c) to form a 3 pole system. Comparisons can now be made between the power dissipation of the dynamically biased log-domain filter and the system enclosed by dashed lines in Fig. 7.18(c), both of which have three poles. The power drawn from the signal source by the first-order RC filter in Fig. 7.18(b) at the largest value of the input is 1.85 mW (see (2.1), (A.10)). The three RC filters in Fig. 7.18(c) consume three times this power—5.55 mW. The

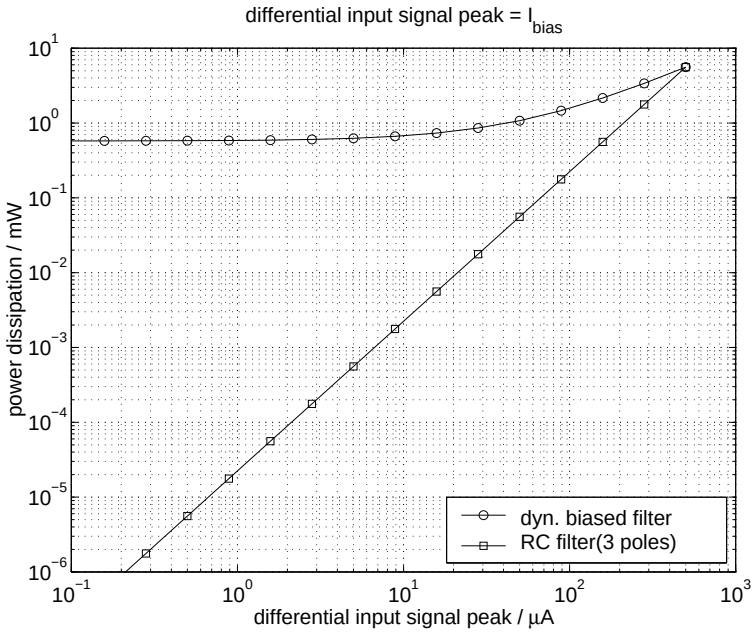

Figure 7.21: Power dissipation versus the input signal: Pseudo-differential ladder filter (Fig. 7.18(a)) and passive RC filter (Fig. 7.18(c)).

power dissipation of the pseudo-differential filter at the highest input is 5.575 mW which is barely larger than that of the passive RC filter with three poles.

The proposed filter achieves a much higher dynamic range per unit power consumption that any of the filters listed in Table A.2, the best of which consumes about a hundred times more power (normalized to the product of the order, the bandwidth, and the dynamic range) than the passive RC filter.

Fig. 7.21 shows the power dissipation of the pseudo differential filter and the passive RC filter in Fig. 7.18 plotted against the input signal. For very small inputs, the passive filter consumes several orders of magnitude smaller power than the dynamically-biased filter. As the input signal increases, the power consumption of the RC filter approaches that of the dynamically biased filter. For very large input signals, the power consumption of the RC filter increases as the square of the input signal (see (A.7)) and that of the dynamically biased filters increases linearly with the input signal. If dynamic biasing is extended to

higher values, the two curves cross and the dynamically biased log-domain filter will be more efficient than the passive RC filter.

As mentioned earlier, 500 μA is not the absolute maximum bias current that can be applied to this filter. The fabricated chip was Bias currents upto 2.5 mA could be applied to the fabricated chip while maintaining acceptable distortion performance for single ended peak inputs that were 50% of the input bias.

7.3 Third-order Butterworth cascade filter with MOS capacitors

7.3.1 Filter synthesis

An alternative to the ladder (Fig. 7.3) realization of the third-order Butterworth filter is the cascade realization based on the prototype in Fig. 7.22(a). A second-order filter with a quality factor of 1 is cascaded with a first-order filter. Fig. 7.22(b) shows this filter constructed using lossless and lossy integrators. The log-domain version of this filter is shown in Fig. 7.22(c). As before, the first-order log-domain filter with multiple inputs shown in Fig. 7.2(d) is used for each of the stages. The desired bandwidth of 1 MHz is obtained by setting $C_1 = C_2 = C_3 = 30$ pF and $I_0 = 5\,\mu$A.

There are minor differences between the cascade (Fig. 7.22(b, c)) and the ladder (Fig. 7.3(b, c)) realizations of the third-order Butterworth filter.

1. The gain of 2 used at the input of the filter in the ladder realization to obtain a unity passband gain is not required in the cascade realization.

2. The feedback from the second stage to the first is absent in the cascade realization.

It was suggested in section 5.2.3 that the low voltage swings in log-domain filters enable the use of non-linear capacitors without additional distortion. It was decided to use pMOS accumulation capacitors for C_1, C_2 and C_3 in Fig. 7.22(c) in the cascade realization of

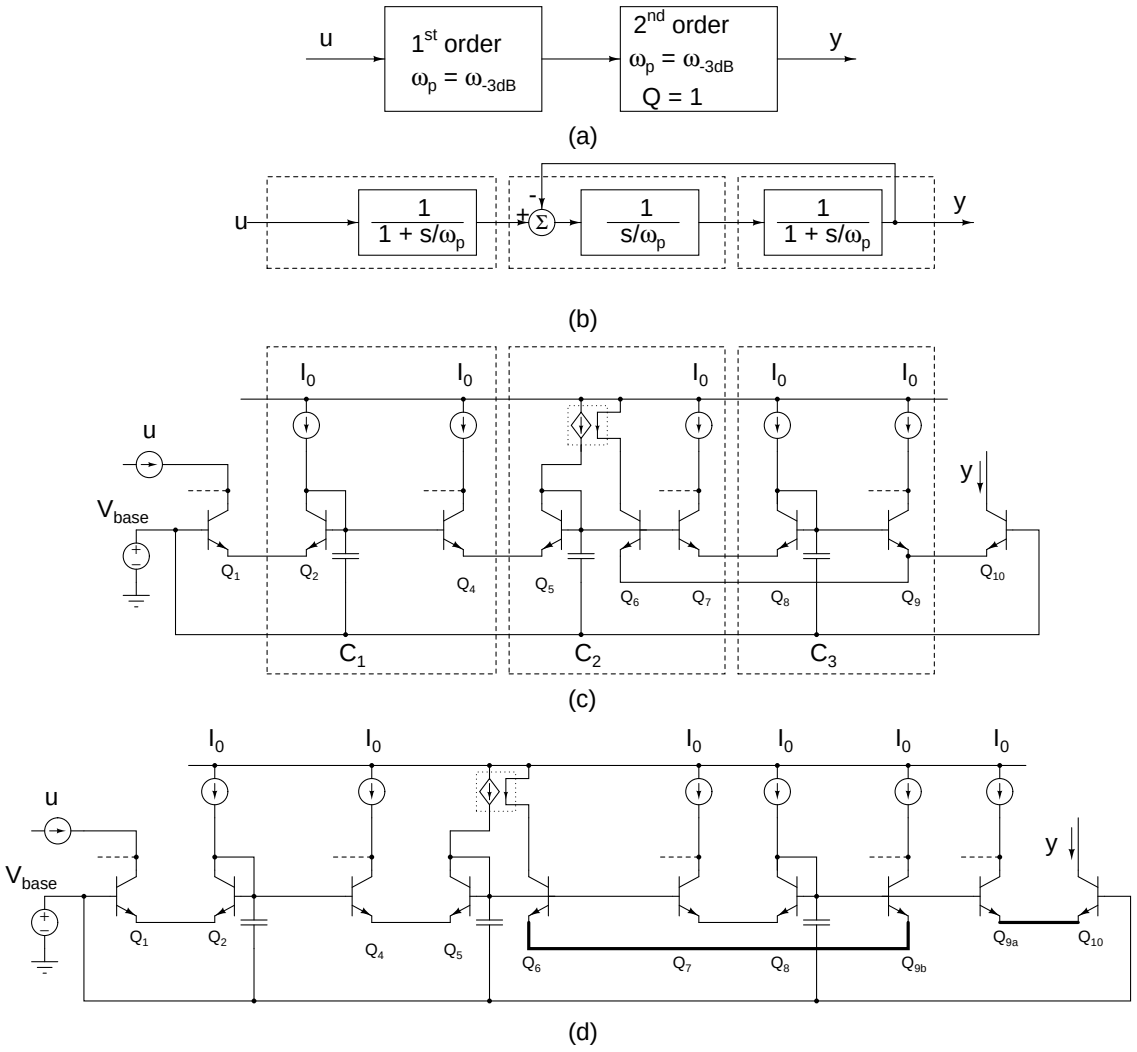

Figure 7.22: Third-order Butterworth filter: (a) Cascade prototype, (b) Block diagram using first order stages, (c) Log-domain realization, (d) Separating the feedback and feed-forward paths.

the filter to verify this assertion. A total of 180 pF of capacitance (in the pseudo differential filter) occupies 0.048 mm^2 as opposed to the 0.47 mm^2 occupied by 240 pF of capacitance used in the ladder filter described in section 7.2. The area saved by using MOS capacitors is indeed significant.

7.3.2 Feedback circuit used to establish the collector currents

The issues discussed in section 7.2.2 regarding the design of the feedback circuit used to establish the collector currents on the transistors of a log-domain filter apply to this filter

as well. The feedback circuit in Fig. 7.6 was used with the cascade realization of the log-domain Butterworth filter.

7.3.3 Transistor sizing at the input and the output

Since dynamic biasing is intended to be used with this filter, the transistor sizing given in Fig. 7.7 is used with this filter.

7.3.4 Feedback and feed-forward paths in the filter

The separation of feedback and feed-forward paths discussed in section 7.2.4 is used in the cascade realization of the filter. The emitter follower Q_9 in Fig. 7.22(c) is split into Q_{9a} and Q_{9b} as shown in Fig. 7.22(d).

7.3.5 Effect of nonidealities

The same considerations apply to the nonidealities in the cascade realization (section 7.2.5) as in the ladder realization. The current sources and current mirrors shown in Fig. 7.10 are used with this filter.

7.3.6 Bias generation

The bias generation circuit used with the cascade realization of the filter is the same as that used with the ladder realization and is shown in Fig. 7.11.

7.3.7 Pseudo differential version of the filter

For pseudo-differential operation, the filter in Fig. 7.22(d) is duplicated and laid out symmetrically.

7.3.8 Automatic tuning of the filter

If desired, the methods mentioned in section 7.3.8 can be applied to the cascade realization as well. No automatic tuning was incorporated into the current design.

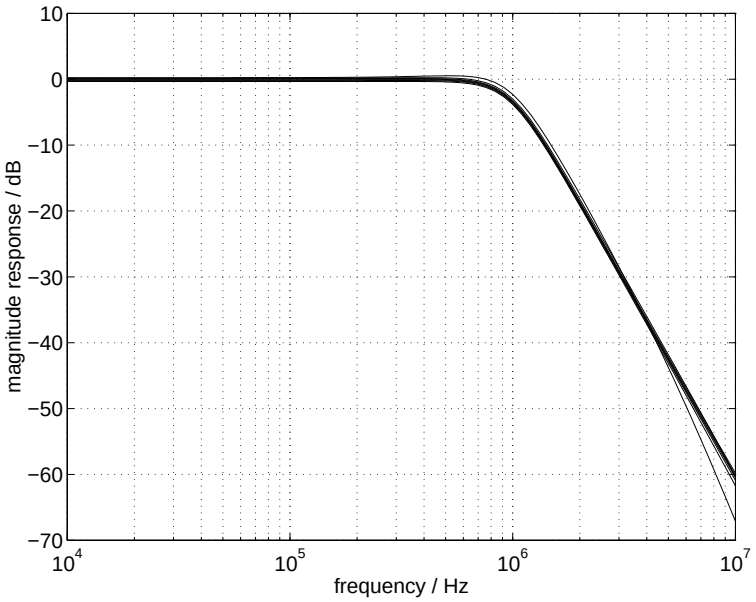

Figure 7.23: Frequency response with the bias current varied from 500 nA to $500\ \mu$A.

7.3.9 Simulation results

Simulations of the cascade realization of the pseudo differential third-order Butterworth filter were carried out in order to assess its performance. The non-linearity of the capacitors was not modeled. Since the voltage swing across the capacitors is a few tens of millivolts in a log-domain filter, it was assumed that the distortion due to them would be negligible [57]. By comparing measured results to those from simulation, the distortion added (if any) by the non-linearity of the MOS capacitors, can be determined. As before, bias currents in the range 50 nA to 500 μA were used in the simulations although the latter is not a hard limit on the bias current that can be applied to the circuit.

Frequency response

Fig. 7.23 shows the magnitude response of the filter with different bias currents. Peaky response was observed with bias currents of 0.5 μA and smaller. Fig. 7.24 shows the pass-band detail. The variation in dc gain with bias is about 0.6 dB.

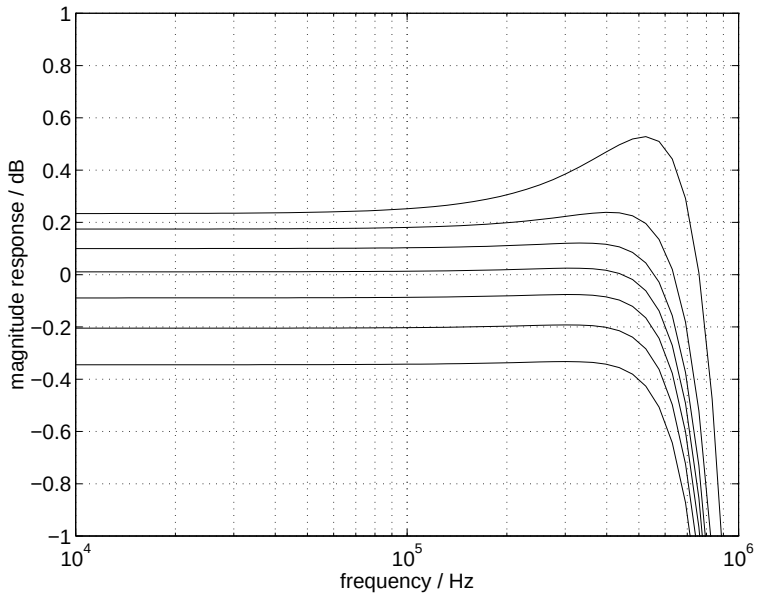

Figure 7.24: Passband detail of the frequency response with the bias current varied from 500 nA to 500 μA.

Output noise and signal to noise ratio

Fig. 7.25 shows the output noise integrated in the (100 kHz, 3 MHz) band. Fig. 7.26 shows the signal to noise ratio assuming a modulation index of 50% (i.e. differential input peak is equal to I_{bias}). As expected, S/N increases by 0.5 dB per dB of bias for small bias and saturates for large bias. If, as before, 0.5 μA is considered to be the lower limit of bias current, the filter maintains a minimum S/N of 48 dB over a dynamic range of 60 dB (up to 500 μA).

Distortion

Fig. 7.27 shows the third harmonic distortion with an input frequency of 312.5 kHz as a function of modulation index for various bias currents. The distortion is higher than in the previous case (Fig. 7.16). Fig. 7.28 shows the distortion vs. frequency for a few values of bias current.

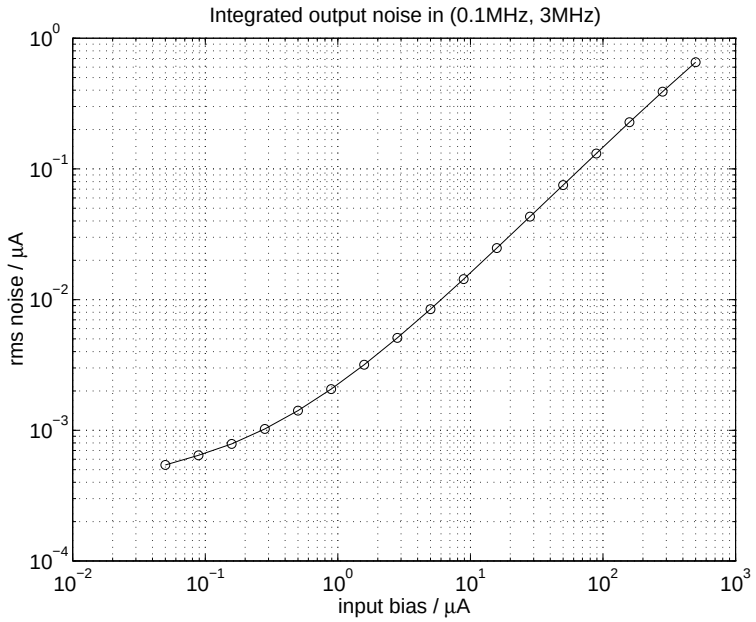

Figure 7.25: Output noise integrated in the 100 kHz to 3 MHz band.

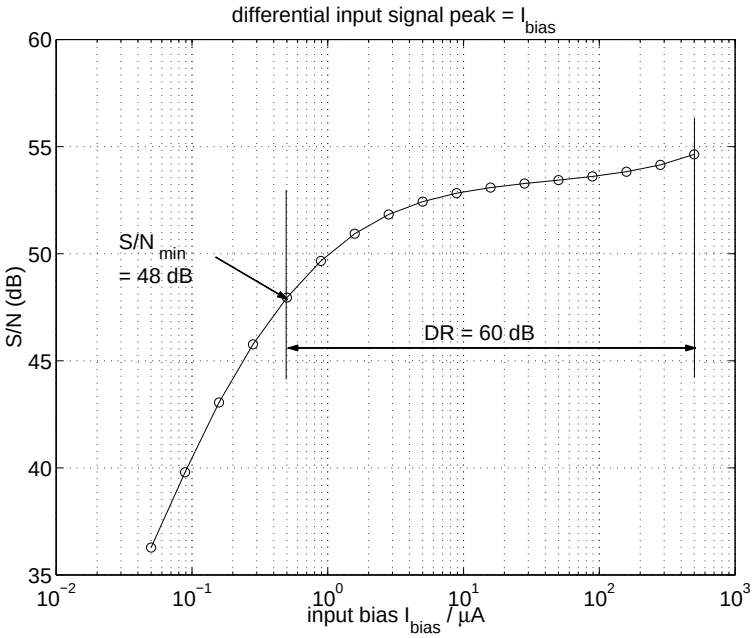

Figure 7.26: Signal to noise ratio with single ended signal peak at 50% of the bias

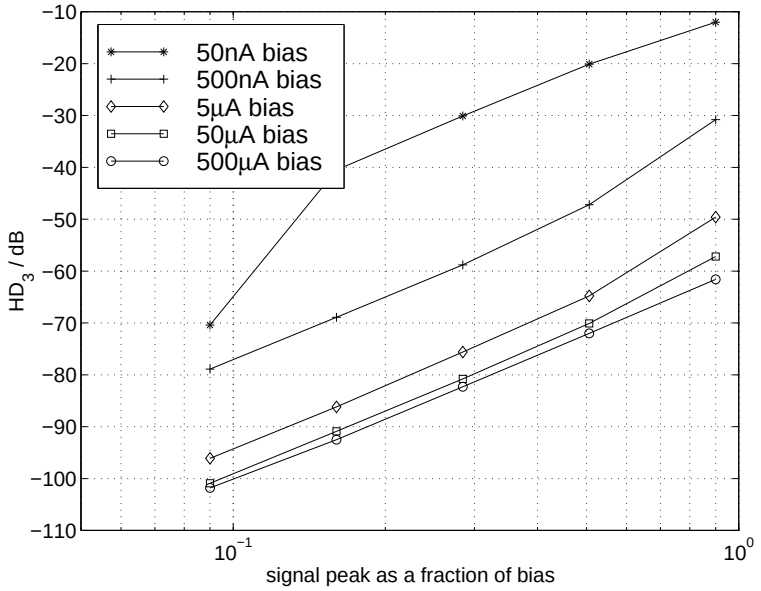

Figure 7.27: Third harmonic distortion as a function of the modulation index for various values of bias currents.

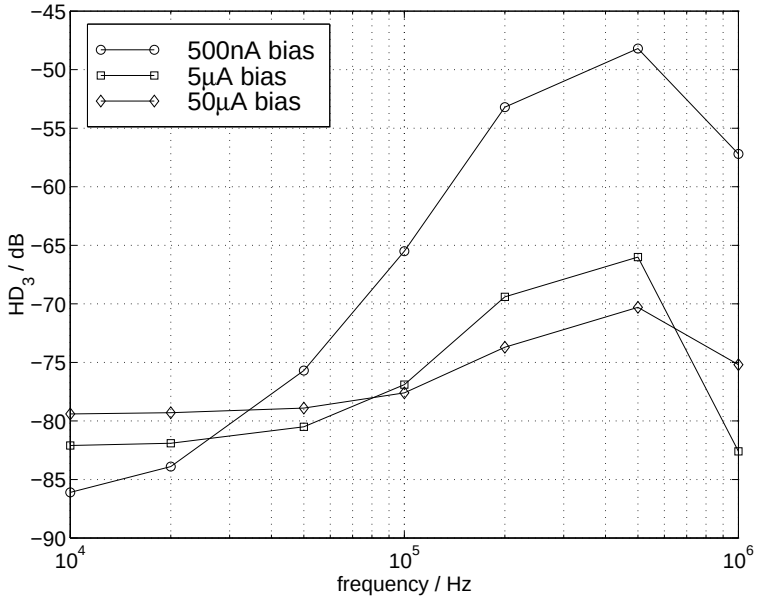

Figure 7.28: Third harmonic distortion as a function of frequency.

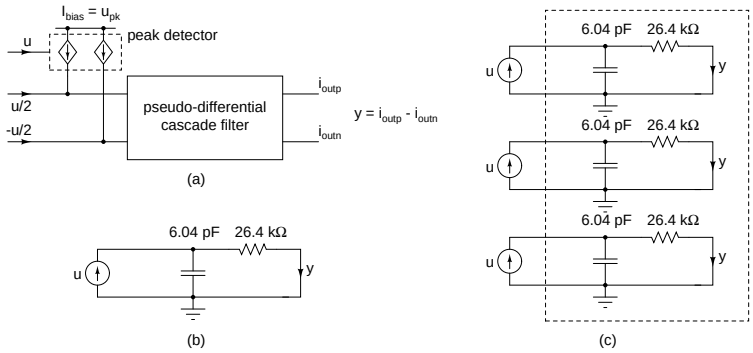

Figure 7.29: (a) Pseudo differential cascade filter biased from an ideal peak detector, (b) Current mode RC filter with $S/N = 48$ dB for a 0.5 μA peak input, (c) Three first-order RC filters.

Power consumption

The filter, excepting the input and output branches (a total of four in the pseudo-differential circuit) draws 190 μA from the power supply. When the input bias is 500 μA, a total of 2.19 mA is consumed in the filter core and the input and output branches. The power consumption is slightly smaller than that of the ladder filter due to the absence of the feedback path from the second stage to the first stage (Fig. 7.22 vs. Fig. 7.3).

Power consumption and dynamic range: comparison to passive filters

Fig. 7.29(a) shows the pseudo differential cascade filter biased from an ideal peak detector. The bias current I_{bias} at the input is set to twice the value of the single ended peak signals. Fig. 7.29(b) shows a passive RC filter designed for a bandwidth of 1 MHz and an output signal to noise ratio of 48 dB for an input current of 0.5 μA peak.

The left and right Y-axes of Fig. 7.30 respectively show the current and power drawn by the cascade realization of the Butterworth filter from a 2.5 V power supply. As with Fig. 7.19, the X-axes shows both the bias current I_{bias} used at the input and the peak value of the differential input signal.

The simulated signal to noise ratio of the pseudo differential filter and the RC filter shown in Fig. 7.29 is shown in Fig. 7.31. Both filters have a signal to noise ratio of 48 dB

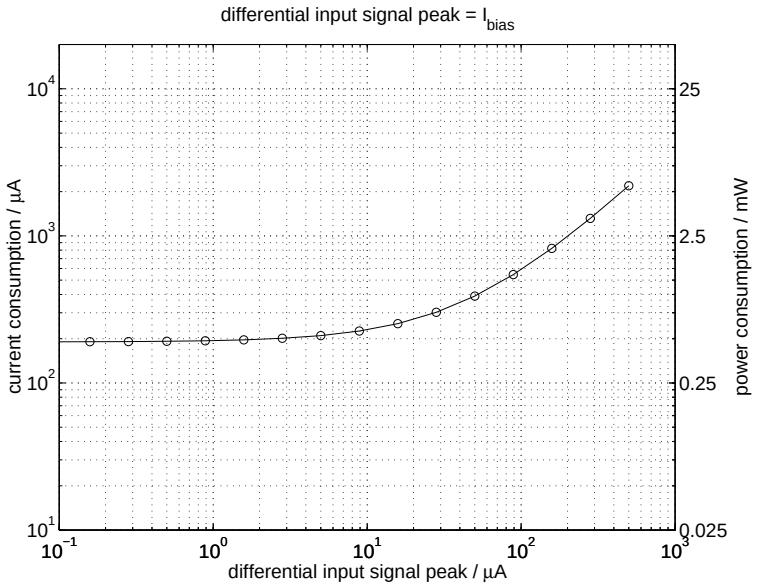

Figure 7.30: Power consumption of the pseudo-differential log-domain filter.

when the input signal is 0.5 μA. The peak signal to noise ratio of the RC filter is 108 dB.

As before, a 3 pole RC filter is formed as shown in Fig. 7.29(c) by simply repeating the first order RC filter in Fig. 7.29(b) three times. Fig. 7.32 shows the variation of the dissipated power with the input signal for the 3 pole filters in Fig. 7.29(a) and Fig. 7.29(c). Again, with the largest input, the two filters dissipate nearly the same power. There is no qualitative difference between the cascade and ladder realizations as far as their power dissipation and dynamic range is concerned. The distortion of the cascade realization is larger than that of the ladder realization, as has been seen.

7.4 Peak detector

7.4.1 General principle

In a dynamically biased log-domain filter, a means of measuring the input signal strength is required. In this design, a peak detector was chosen for this purpose. Since the log-domain filter operates in current mode, it is preferable to have the entire peak detector in the current mode. A transconductor following a voltage mode peak detector is a possi-

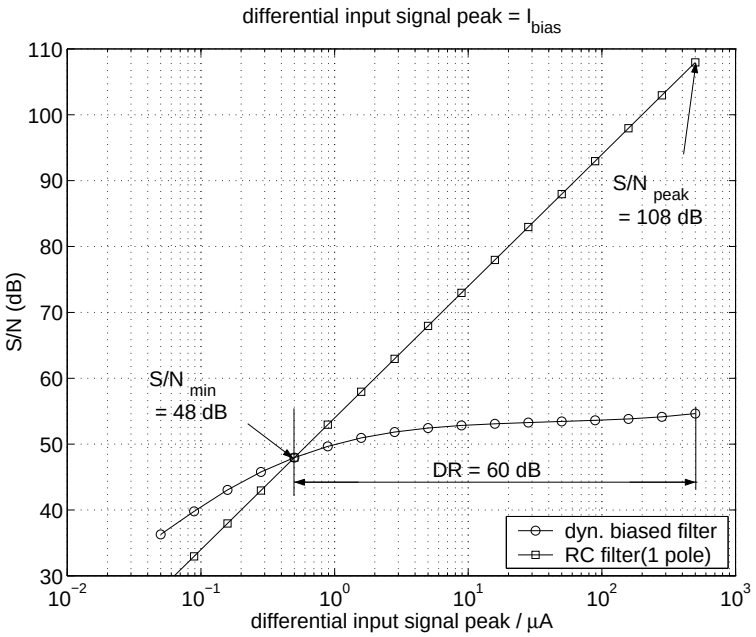

Figure 7.31: Signal to noise ratio of the log-domain and passive RC filters.

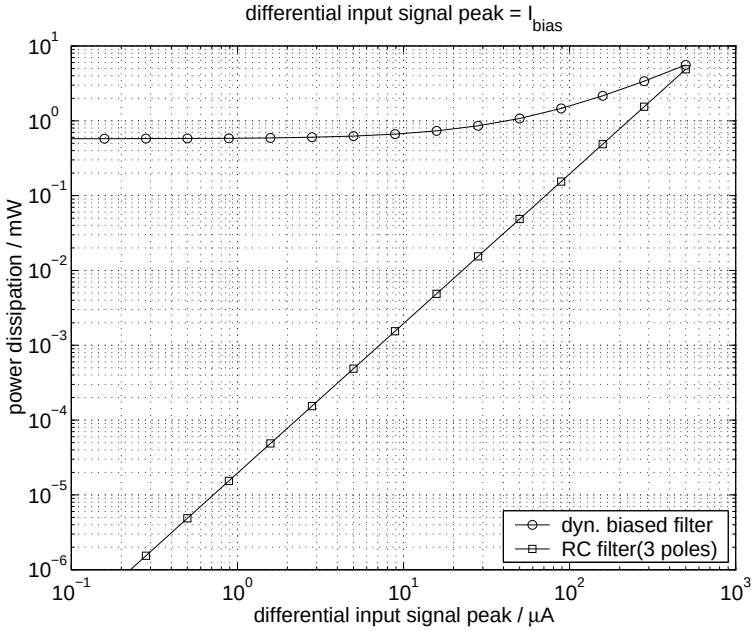

Figure 7.32: Power dissipation versus the input signal: Pseudo-differential cascade filter (Fig. 7.29(a)) and passive RC filter (Fig. 7.29(c)).

bility, but the design of a transconductor operating over the intended wide range of bias currents is non-trivial.

7.4.2 The prototype peak detector

Fig. 7.33(a) shows the classical diode-RC peak detector. This circuit detects the peak value of the input signal if an ideal diode is used. In this context, an ideal diode is one which has a zero voltage drop across it when a forward current is flowing through it and a zero current through it when a reverse voltage is applied across it. Due to the nonzero voltage drop of real diodes the output dc of the peak detector in Fig. 7.33(a) is less than the peak value of the input. This constitutes an error in the peak detector's output and this error increases in a relative sense as the peak value of the input decreases.

The current i_d of a diode is exponentially related to the voltage v_d across it.

$$i_d = I_s \exp(m v_d / V_t) \tag{7.15}$$

where I_s and V_t are the saturation current and the thermal voltage respectively. m is unity for a real diode, but is introduced here to aid the following discussion. The exponential in the equation above is plotted in Fig. 7.34 as a function of the normalized diode voltage v_d/V_t with m as a parameter. It is seen that the diode characteristics become "sharper" with increasing m. $m = \infty$ corresponds to the ideal diode described above.

The peak detector circuit in Fig. 7.33(a) can be redrawn as a generalized block diagram shown in Fig. 7.33(b). The error V_{err} between the input and the output is fed to an exponentiator which drives a low pass filter. V_0 is a normalizing voltage. This block diagram with $V_0 = I_s R_L$, $m = 1$ and $V_{err} = v_d$ corresponds to Fig. 7.33(a).

From the curves in Fig. 7.34, it can be inferred that the accuracy of the peak detector in Fig. 7.33(b) can be improved by increasing the multiplying factor m in the argument of

the exponential. A large value of V_0 helps as well, but not as effectively as a large m that is inside the argument of the exponential.

Fig. 7.33(c) shows another well known circuit used to improve the accuracy of the peak detector for small input signals. The error signal, which is the difference between the input and the output is amplified by an amplifier of gain A before being exponentiated. Initially assume that the offset voltage V_{off} is zero. In this case, the diode current i_d that drives the current mode low pass filter is given by

$$i_d = I_s \exp\left(\frac{A(v_i - v_o) - v_o}{V_t}\right) \tag{7.16}$$

This expression is approximately the same as (7.15) with $m = A$. A large value of A can be used to emulate the characteristic of the ideal diode as closely as possible and maintain the accuracy of peak detection for small signals.

The generalized block diagram in Fig. 7.33(b) is used as the prototype for the current mode peak detector required for this work.

While Fig. 7.34 suggests that the value of m used must be as large as possible, in practice, there may not be much point to increasing m beyond a certain value due to the following reasons.

1. If the exponentiator has an offset (illustrated in Fig. 7.33(c) with an offset voltage V_{off}), for small inputs, the error may be dominated by the offset and not by the "sharpness" of exponentiation.

2. If a large gain A is realized using complicated circuitry with many poles, the loop may be unstable.

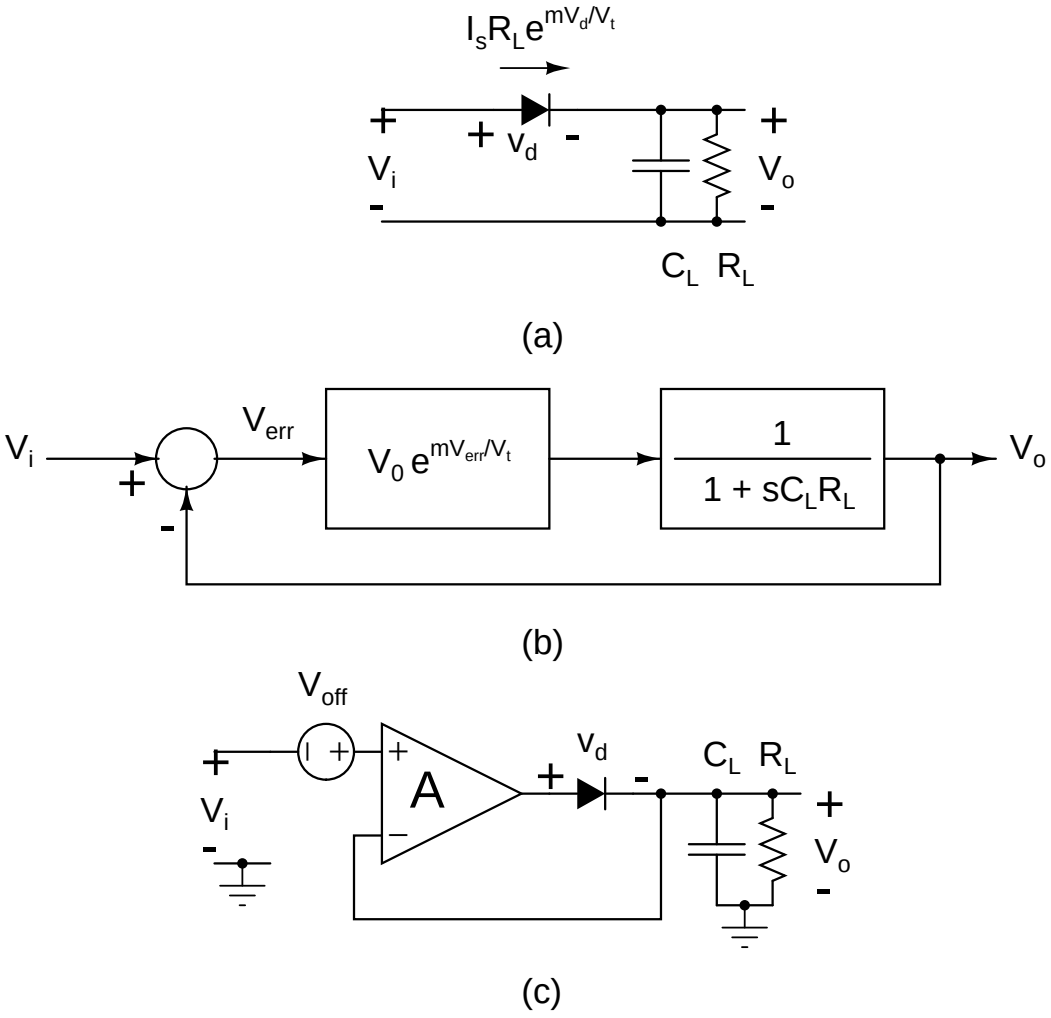

Figure 7.33: Principle of peak detection.

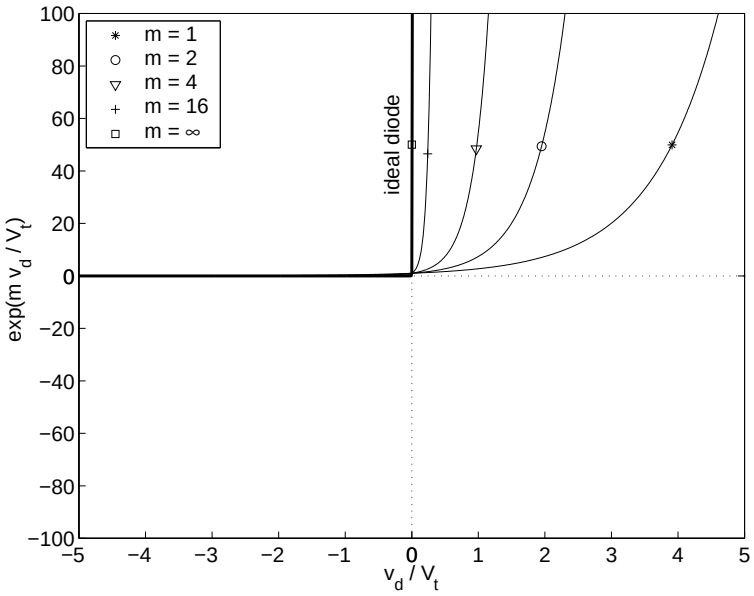

Figure 7.34: Diode characteristics

Attack and decay of the peak detector

When the error V_{err} in Fig. 7.33(b) is negative, i.e. when the input is smaller than the output, the output of the exponentiator is nearly zero. The output V_o is thus the decaying natural response of the low pass filter. For inputs V_i larger than the output V_o, the exponentiator's output shoots up, driving the low pass filter's output up along with it. The increase in the peak detector's output is much faster than its decrease due to the exponentiation of the error signal.

7.4.3 Current mode peak detector

The block diagram in Fig. 7.33(b) should be realized with current input and current output to realize a current mode peak detector. The first-order filter in Fig. 7.1(b) can be used as the low pass filter in Fig. 7.33(b). This filter is shown in Fig. 7.35(a) with u and y substituted by I_1 and I_4 respectively. As discussed in earlier sections, this circuit behaves as a low pass filter between the input I_1 and the output I_4. The emitter voltage V_{e1} of transistor Q_1 in

Fig. 7.35(a) can be written as

$$V_{e1} = V_{base} - V_t \ln\left(\frac{I_1}{I_s}\right) \tag{7.17}$$

Now this circuit is modified as shown in Fig. 7.35(b) where the base of Q_1 has been disconnected from the voltage V_{base} and connected to another voltage $V_{base} - V_b$. The emitter voltage in this case is

$$\begin{aligned} V_{e1a} &= V_{base} - V_b - V_t \ln\left(\frac{I_1}{I_s}\right) && (7.18)\\ &= V_{base} - V_t \ln\left(\frac{I_1 e^{\frac{V_b}{V_t}}}{I_s}\right) && (7.19)\\ & && (7.20) \end{aligned}$$

(7.20) is identical to (7.17) except for the exponential multiplying I_1. Since the circuit in Fig. 7.35(a) filters I_1, it can be concluded from (7.17) and (7.20) that the circuit in Fig. 7.35(b) filters $I_1 e^{\frac{V_b}{V_t}}$.

Thus the behavior of the circuit in Fig. 7.35(b) between the voltage V_b and current I_4 is that of an exponentiator followed by a low-pass filter. The output of the low pass filter I_4 would be the output of the peak detector (Fig. 7.33(b)) in the final configuration. I_4 can be subtracted from the input current I_{in} to generate an error current I_{err} as shown in Fig. 7.35(b). This error current I_{err} must be converted into a voltage V_{err} and fed back as $V_{base} - V_b$ to close the loop and form the peak detector.

The high impedance at the input node (output impedance of Q_4) can itself be used for current to voltage conversion with some clipping circuitry to limit the peak-to-peak swing. The resulting voltage can be amplified to obtain the voltage $V_{base} - V_b$ that drives the base of Q_1. An inverting amplifier should be used to obtain negative feedback around the loop. Fig. 7.36 shows the negative feedback incorporated around the exponentiator-low pass filter combination.

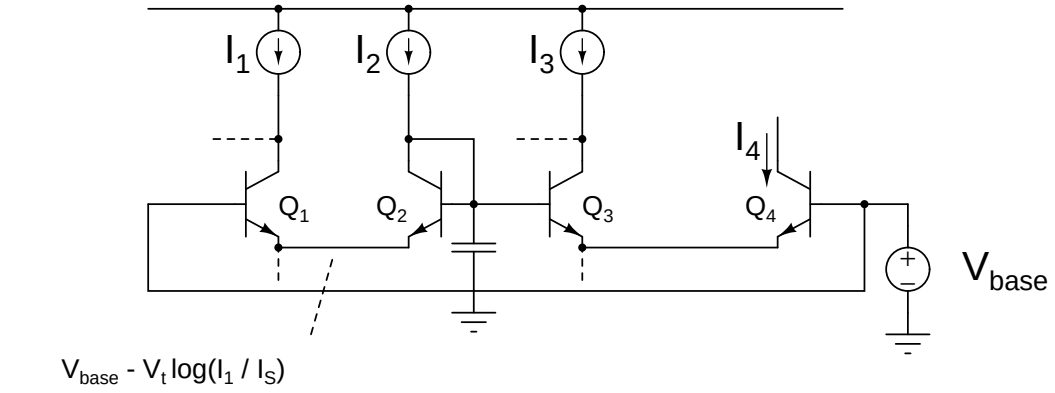

(a)

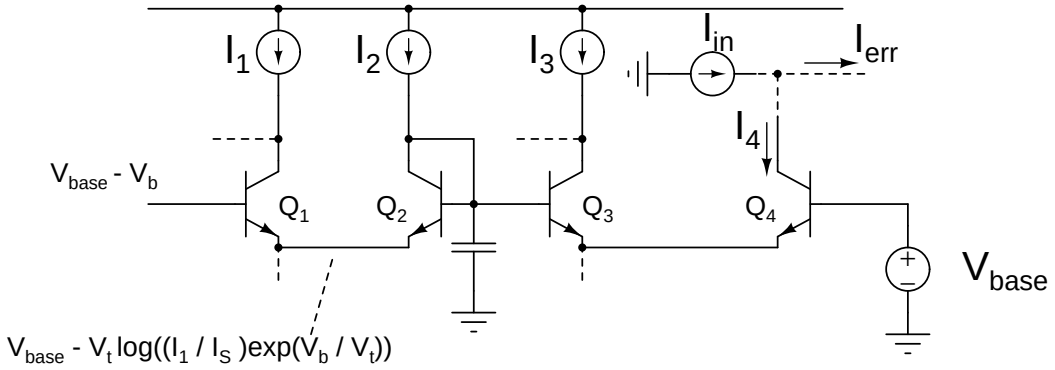

(b)

Figure 7.35: Combination of an exponentiator and a low pass filter.

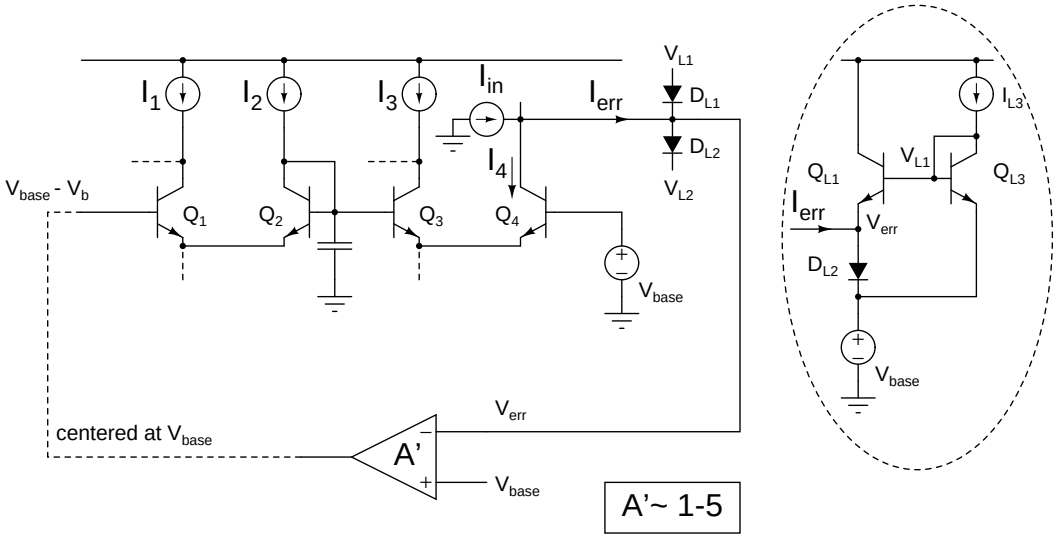

Figure 7.36: Negative feedback around the exponentiator and the low pass filter.

Diodes D_{L1} and D_{L2} with bias voltages V_{L1} and V_{L2} provide limiting of the voltage at the input node. V_{L1} and V_{L2} should be chosen to ensure proper operation of the circuit. We make the following observations:

1. Too small a difference between V_{L1} and V_{L2} results in a large peak-peak swing at the input node. The voltage change at the input node is produced by the error current charging the input capacitance. The error current is of the same order of magnitude as the input current. Therefore, a large peak-peak swing implies a slower charging, especially with small input currents, resulting a slower "attack" of the envelope detector.

2. Too large a difference between V_{L1} and V_{L2} results in a small peak-peak voltage swing at the input node which may not be sufficient to drive the filter to the desired peak level.

3. Too large a value for V_{L1} (for a given difference $V_{L1} - V_{L2}$) causes the input node voltage to rise, which in turn causes a large quiescent output current I_4.

4. Too small a value for V_{L2} (for a given difference $V_{L1} - V_{L2}$) causes the input node voltage to fall, resulting in small or zero output current I_4.

A suitable clipping circuit is shown as an inset in Fig. 7.36. V_{base} is used for V_L. The current I_{L3} and the area of Q_{L3} determine V_{L1}.

A gain A' of 1 to 5 in the inverting amplifier was determined to be sufficient from simulations. As mentioned in section 7.4.1, very large values of gain A' result in instability. Fig. 7.37 shows the inverting amplifier of near unity gain used in this design. A systematic offset is introduced in this amplifier (by using different tail currents and transistors of

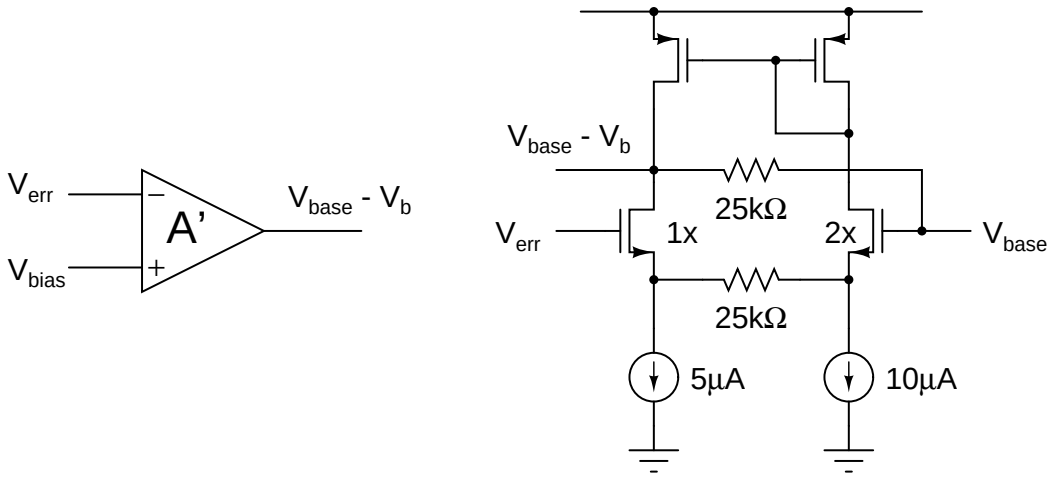

Figure 7.37: Inverting amplifier used in the feedback path of the peak detector.

different aspect ratios in Fig. 7.37) to ensure near zero quiescent output current in the peak detector.

With the feedback loop closed, the circuit in Fig. 7.36 behaves as a peak detector. I_4 is the output of the peak detector. To bias the log-domain filter from this peak detector, another copy of I_4 is required. This can be derived as shown in Fig. 7.38 by adding Q_{3a} and Q_{4a} in parallel with Q_3 and Q_4. Q_{4a} is made 1.5 times larger than Q_4 in order to provide a 50% margin over the peak value of the signal in the added dynamic bias. If this margin is too small, it increases the distortion in the dynamically biased filter. This is not a severe problem, however. Since the inputs are fed separately to the peak detector and the filter in the prototype, a larger input can be fed to the peak detector than to the filter to emulate a larger margin.

The placement of the pole of the low pass filter is determined by the lowest expected signal frequency. This varies with application. In a low pass filter, this can be very low. In this design, an external capacitor was used for filtering to allow for easy experimentation during testing. The bias currents I_2 and I_3 in Fig. 7.38 were 0.6 μA. With a 1 nF capacitor, this results in a pole frequency of roughly 4 kHz.

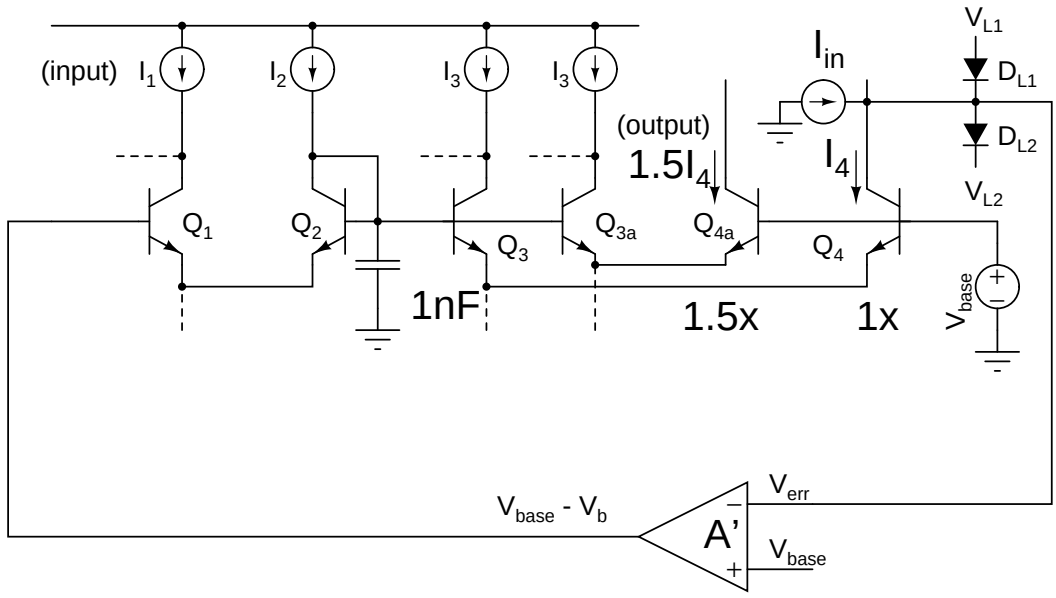

Figure 7.38: Complete circuit of the peak detector.

7.4.4 Simulation results

The single-ended peak detector in Fig. 7.38 (including the padframe) was simulated with a sinusoidal input whose amplitude was varied in logarithmic steps from 50 nA to 5 mA. $V_{dd} = 2.5$ V, $V_{base} = 1.3$ V and $C = 1$ nF were used in the simulations. Fig. 7.39 shows the output dc of the peak detector as a function of the input peak on a logarithmic scale. It can be seen from the curve that the output maintains a near proportionality to the input peak over a wide range.

For a more accurate picture, the ratio of the dc output to the input peak is plotted in Fig. 7.40. A gain of about 1.4 is maintained for input peaks in the range 10 μA to 2 mA. At very large currents, the circuit fails due to limiting of the internal voltage as described in section 7.2. As can be seen in Fig. 7.40, the gain drops abruptly to 1.1 at 5 mA. The gain of the peak detector drops gradually at low currents and is less than unity for currents less than 0.3 μA. This is analogous to the effect of the "diode drop" seen in the circuit of Fig. 7.33(a). Below a certain input amplitude, the diode current is too small to be able to

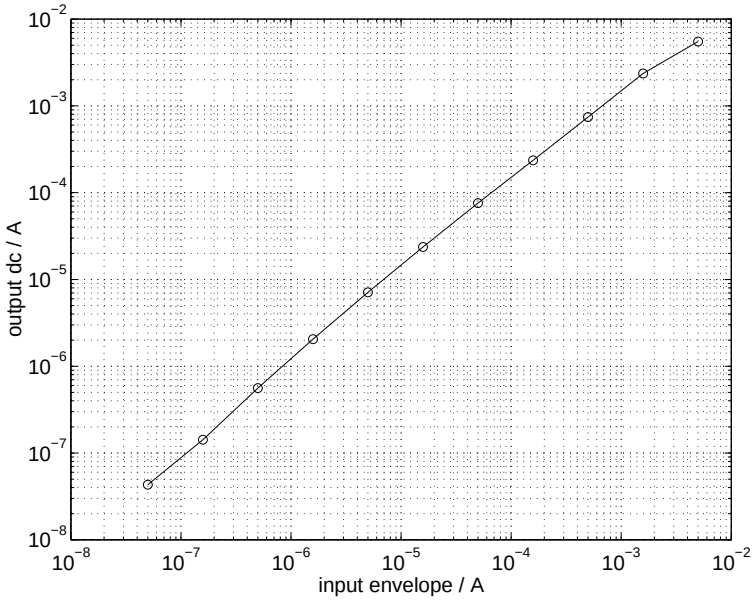

Figure 7.39: Simulated dc output of the peak detector vs. input peak.

rapidly charge the capacitor. In general, a peak detector of the type shown in Fig. 7.33(b) has lower gain at small signals because the error signal is too small for the exponentiator to generate a large drive to the low pass filter. The output ripple as a fraction of the dc output when the input is at a frequency of 1 MHz is shown in Fig. 7.41. The relatively large ripple at the output implies that the output of the envelope detector must be filtered further before being used to dynamically bias the filter.

The dc output of the peak detector was relatively uninfluenced by the input frequency as long as it remained well above the filter's cutoff frequency of 4 kHz.

7.4.5 Peak detector with a differential input.

A signal of wide relative bandwidth (defined here as the ratio of the upper frequency limit to the lower frequency limit) can have a large number of harmonics. This implies exotic shapes—periodic signals can have duty cycles quite different from 50%. A low-pass filter, such as the one described in section 7.2, can be expected to have signals with varying duty cycle at the input.

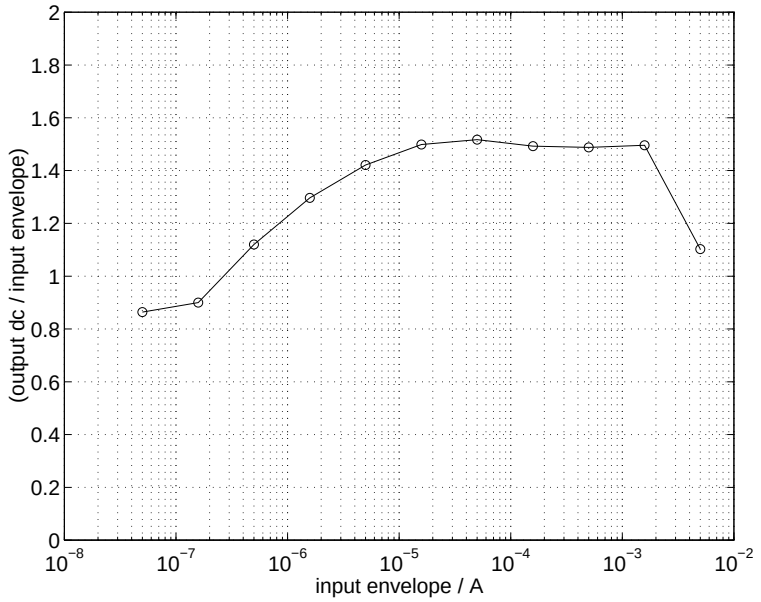

Figure 7.40: Simulated ratio of the dc output to the input peak.

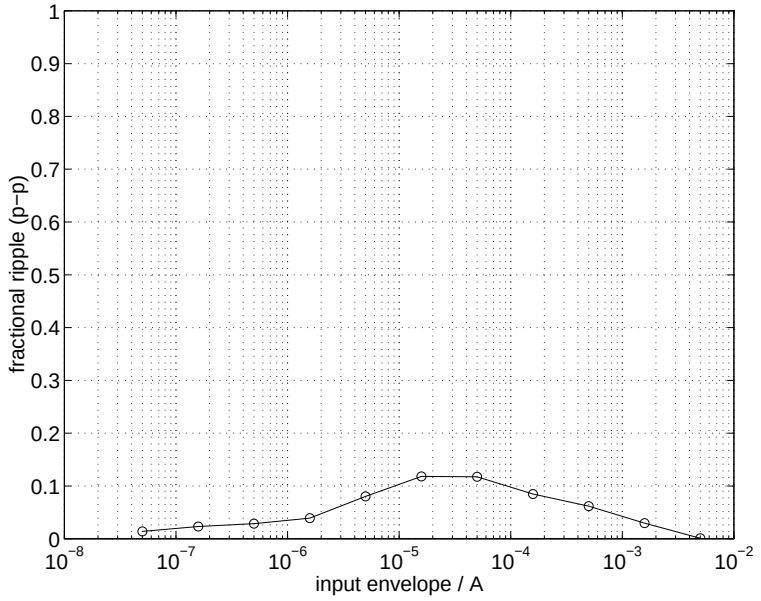

Figure 7.41: Simulated peak-peak ripple as a fraction of the dc output.

Two differential waveform pairs, which could be the inputs to the pseudo-differential log-domain low pass filters described in the earlier sections are shown in Figs. 7.42(a, b). Each of these waveforms has a zero average value. The one in Fig. 7.42(a) has a 50% duty cycle and the one in Fig. 7.42(b) has a duty cycle different from 50%. If a single ended peak detector is used, only one of the waveforms of the differential signal can be fed to the envelope detector. Assume that the waveform shown in dashed lines in Figs. 7.42(a, b) is fed to the single ended peak detector shown in Fig. 7.38. It is evident that the peak detector in Fig. 7.38 detects the *positive* peak of the input waveform I_1. If the peak value so detected is multiplied by 1.2 (to provide a 20% safety margin) and added to the waveforms in Figs. 7.42(a, b) for dynamically biasing the latter, we obtain the waveforms shown in Figs. 7.42(c, d). The signals in Fig. 7.42(c) are always positive, whereas the signal shown in dashed line in Fig. 7.42(d) is not. This is because the negative peak of the signal shown in dashed line in Fig. 7.42(b) is larger than its positive peak (which was used for peak detection and subsequent biasing). Clearly, the waveforms in Fig. 7.42(d) cannot be fed to a log-domain filter.

The solution to this problem is a peak detector operating from differential signals and having the larger of the two peak values as its output. This ensures that sufficient bias is added to both the signals.

Fig. 7.42(a) shows a peak detector with differential inputs. The dc output is the larger of the peak values of V_i and $-V_i$. This is a circuit commonly used as a rectifier with center-tapped transformers in dc power supplies. The same circuit is shown as a block diagram in Fig. 7.43(b). Two exponentiators are used to amplify the error between the output and the negative of the input. The current mode circuit corresponding to this block diagram can be synthesized in a manner similar to that described in section 7.4.3.

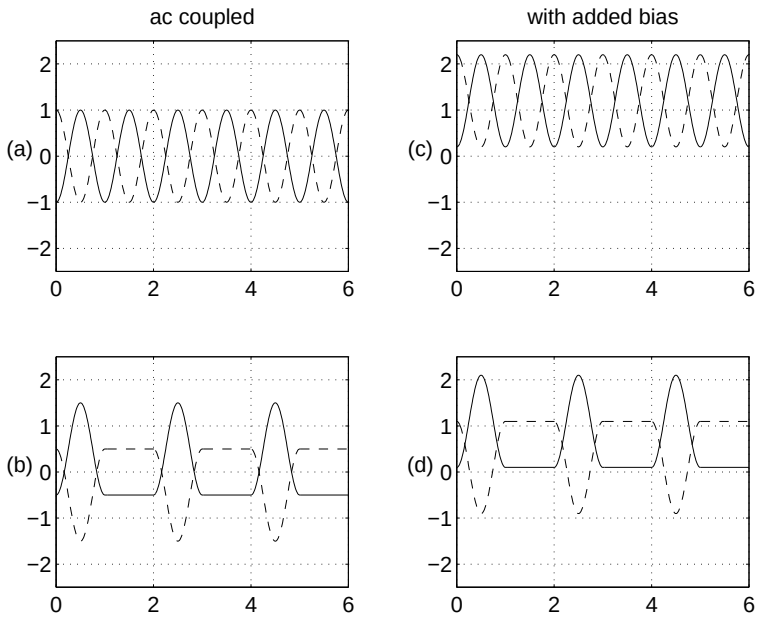

Figure 7.42: (a) Differential signal with equal positive and negative peaks, (b) Differential signal with unequal positive and negative peaks, (c) (a) with added bias, (d) (b) with added bias.

The result is shown in Fig. 7.43(c). This circuit is a combination of two single current mode peak detectors (Fig. 7.38) which have a common hold capacitor.

Due to time constraints, the differential version of the current mode peak detector was not fabricated. The idea can however be tested using the fabricated single ended peak detectors. Two chips can be used with a common hold capacitor as shown in Fig. 7.43(c). The output can be taken from either envelope detector. Compared to an integrated differential peak detector, the only disadvantage of this arrangement is the redundancy of the on-chip bias circuitry. Mismatch between chips can cause asymmetry in its operation. This is akin to using two diodes of different types in Fig. 7.43(a). However, matching between chips from the same fabrication run can be expected to be close enough to avoid such problems.

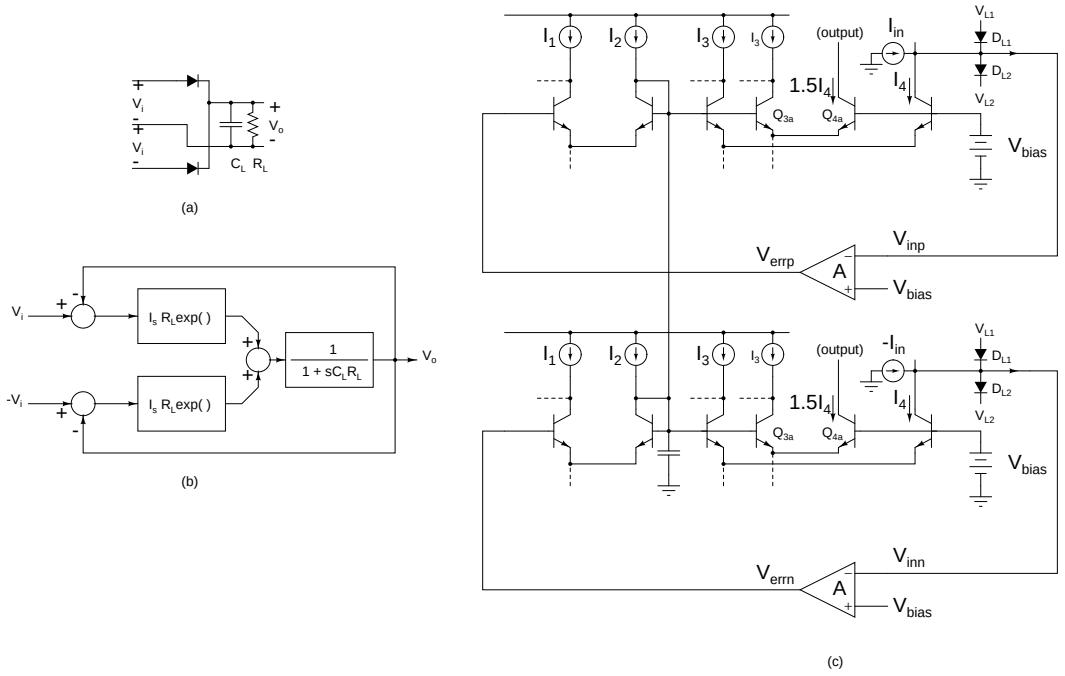

Figure 7.43: Peak detector with differential inputs.

7.5 Combination of the filter and the peak detector

To bias the Butterworth filter from the peak detector, a current mirror is necessary. This current mirror should be able to handle a wide range of currents. Fig. 7.44(a) shows a simple pMOS current mirror. To be able to handle currents as large as a few mA, the transistors need to be very wide. But the resulting large input capacitance makes it unable for the mirror to follow the fast attack of the peak detector at low currents. More importantly, in this technology with a pMOS threshold voltage in excess of 0.9 V, the voltage at the gate of the diode connected device M_1 would be too large to be accommodated at the output of the peak detector.

A solution to this problem is shown in Fig. 7.44(b). An NPN emitter follower is used to drive the gates of the pMOS transistors. Due to the large V_T of M_1, this arrangement results in a drain source voltage of about 0.3 V which is sufficient to keep it in saturation. An I_{bias} of 1.5 μA results in a satisfactory frequency response of the current mirror.

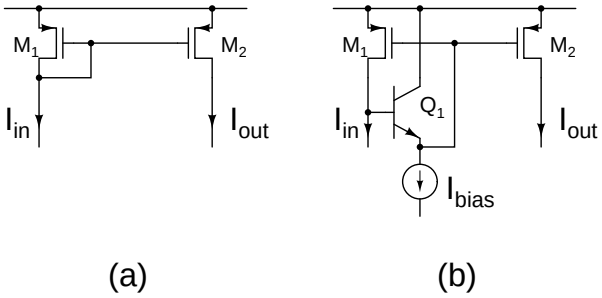

Figure 7.44: Current mirrors.

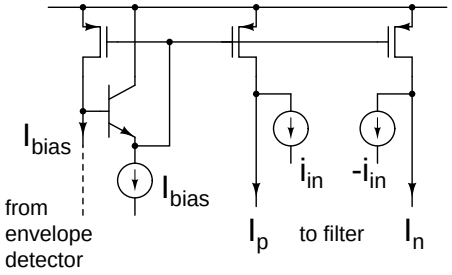

Figure 7.45: Biasing the filter from the peak detector.

Fig. 7.45 shows the filter being biased from the peak detector using this current mirror. The differential input current i_{in} and $-i_{in}$ are fed to the filter as shown in Fig. 7.45.

7.6 Second-order filter using lateral PNP transistors and MOS capacitors.

It was proposed in section 5.2 that log-domain filters could be realized in a CMOS technology using enhanced lateral PNP transistors which have a large dc current gain [54]. To test the viability of such a scheme, a second-order log-domain filter was laid out using only the CMOS part of the technology described in section 7.1. Due to the limited time available to finish the layout of the chip before the fabrication deadline, not much attention was paid to any design details.

Fig. 7.46 shows the block diagram of the filter. A second-order Butterworth filter is formed by placing two lossy integrators in a unity gain feedback loop. The transfer

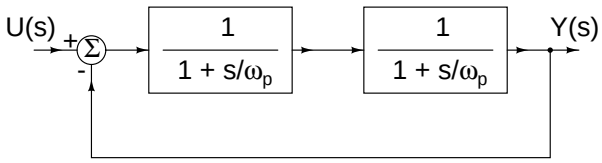

Figure 7.46: Block diagram of a second-order Butterworth filter using two lossy integrators.

function is given by

$$\frac{Y(s)}{U(s)} = \frac{1/2}{1 + (s/\omega_p) + (s/\sqrt{2}\omega_p)^2} \tag{7.21}$$

where ω_p is the pole of the lossy integrator. The resulting Butterworth filter has a -3 dB bandwidth of $\sqrt{2}\omega_p$. At the output of the first integrator, a bandpass characteristic with a non-zero dc gain is obtained. Though this design is slightly more noisy than a second order low-pass filter with one lossless integrator, it was chosen because its modularity enabled a quicker layout.

Fig. 7.47(a) shows the schematic of the filter. Since no dynamic biasing was intended, the simple feedback topology shown in Fig. 7.4(a) [19] (with transistors of opposite polarity) was used to force input and bias currents into Q_1, Q_4 and Q_7. pMOS devices in this technology have a threshold voltage of 0.9 V which keeps both the bipolar and the MOSFET (e.g Q_1 and M_1) comfortably in the proper regions. Since the collector of Q_1 (also Q_4 and Q_7) is at a Volt or so below the top supply rail, it is expected that the filter can operate with a 1.5 V supply. Fig. 7.47(b) shows the schematic of the filter with enhanced lateral PNP transistors (pMOS transistors connected as shown in the figure) and MOS capacitors. The dc bias across the pMOS accumulation capacitors is nearly equal to the potential difference between V_{base} and ground. This can be about 0.7 V while operating from a 1.5 V supply. As before, cascode current mirrors (Fig. 7.47(d)) and current sources (Fig. 7.47(c)) were used in the design. The bias generation circuitry is shown in (Fig. 7.47(e)). pMOS accumulation capacitors of approximately 140 pF each were used for C_1 and C_2. With a $I_0 = 5\,\mu$A,

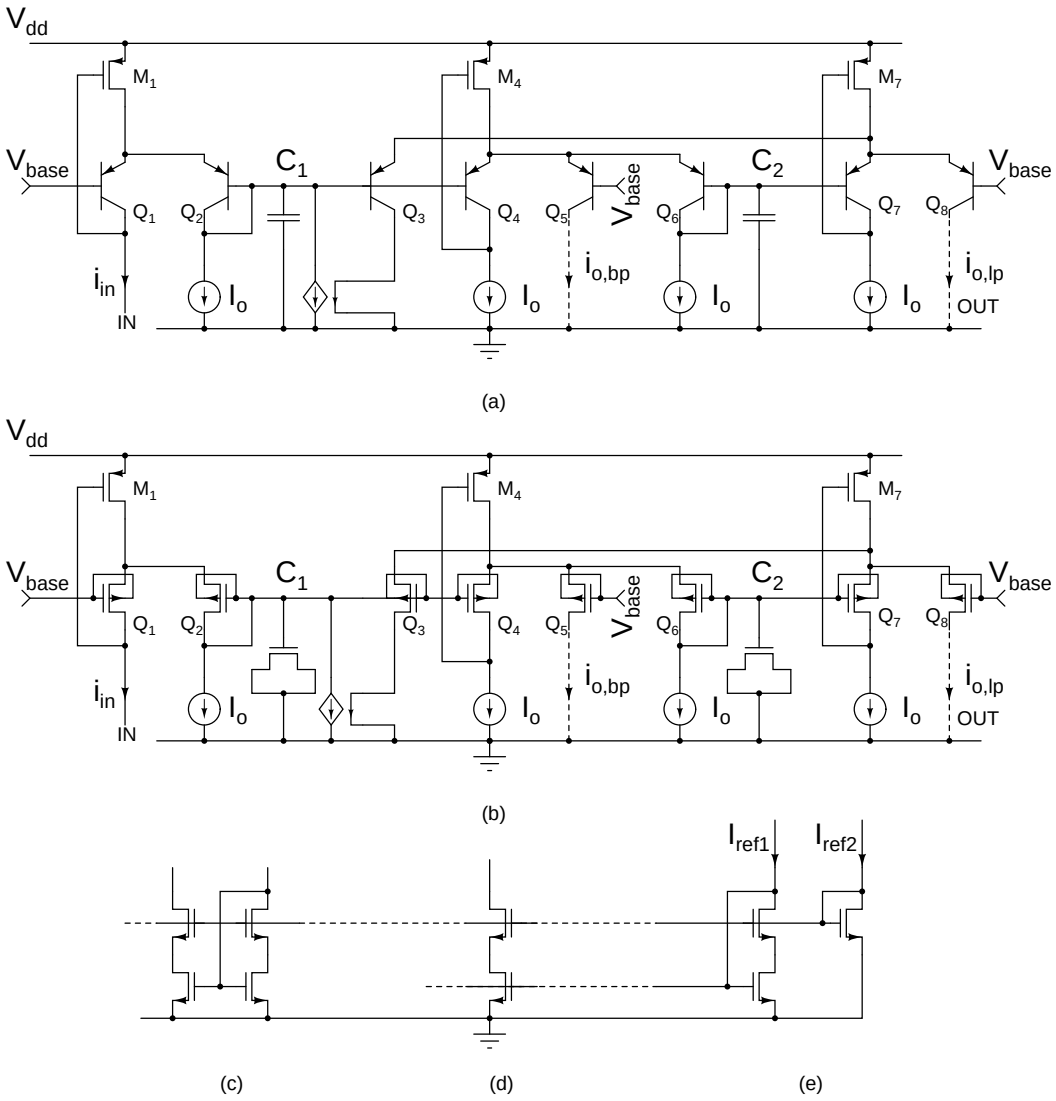

Figure 7.47: CMOS Log-domain realization of Fig. 7.46: (a) Log-domain topology, (b) Circuit in (a) redrawn with lateral PNP transistors (pMOS transistors with gate and bulk tied together) and MOS capacitors, (c) Current mirror, (d) Current source, (e) Bias generator.

the bandwidth of the Butterworth filter would be 31 kHz. Lateral bipolar transistors with 0.36 μm base width and $8 \times 0.4\,\mu\text{m} \times 0.4\,\mu\text{m}$ emitter area[8] were used in the circuit. No simulations were carried out since the models for lateral bipolar transistors were not available at the time. It was ensured to the extent possible that the parasitic capacitances across the two desired capacitances were equal by adding dummy devices. This ensures that the pole frequencies of the two lossy integrators in Fig. 7.46 are equal, resulting in a Butterworth response. Significant parasitics across C_1 and C_2 reduce the bandwidth, but the shape of the magnitude response would be maintained. Some test lateral PNP devices were also laid out so that they could be characterized.

[8]This is the emitter area as seen in the plan view of the transistor (Fig. 5.1(e)). The actual emitter action takes place along the sidewalls of the square emitter.

Chapter 8

Testing Procedures and Measured Results

8.1 Third order Butterworth ladder filter

Fig. 8.1 shows the setup used to measure the pseudo-differential Butterworth ladder filter (section 7.2). The input voltage v_{in} from a signal generator is converted into differential input currents i_{inp} and i_{inn} using a transformer and a pair of resistors R_{in} and fed to the low impedance inputs of the pseudo-differential filter. The differential output currents i_{outp} and i_{outn} of the filter are converted into a single-ended current using a transformer and fed to a transimpedance amplifier to obtain an output voltage v_{out}. The transimpedance amplifier consists of a resistor R_{out} in feedback around a wideband op-amp LM 7121. The 5 μA reference sources are used to set the bandwidth of the filter to 1 MHz. V_{bias} in conjunction with R_{in} is used to vary the input bias current I_{bias} of the filter. To vary I_{bias} over a large range, R_{in} is varied from 2 kΩ to 1 MΩ.

Frequency response

To determine the frequency response of the filter, the frequency response between v_{in} and v_{out} in Fig. 8.1 was measured twice; once with the filter in place and once with the filter replaced by short circuits between its respective inputs and outputs. The difference between the two frequency responses so obtained is the frequency response of the filter between

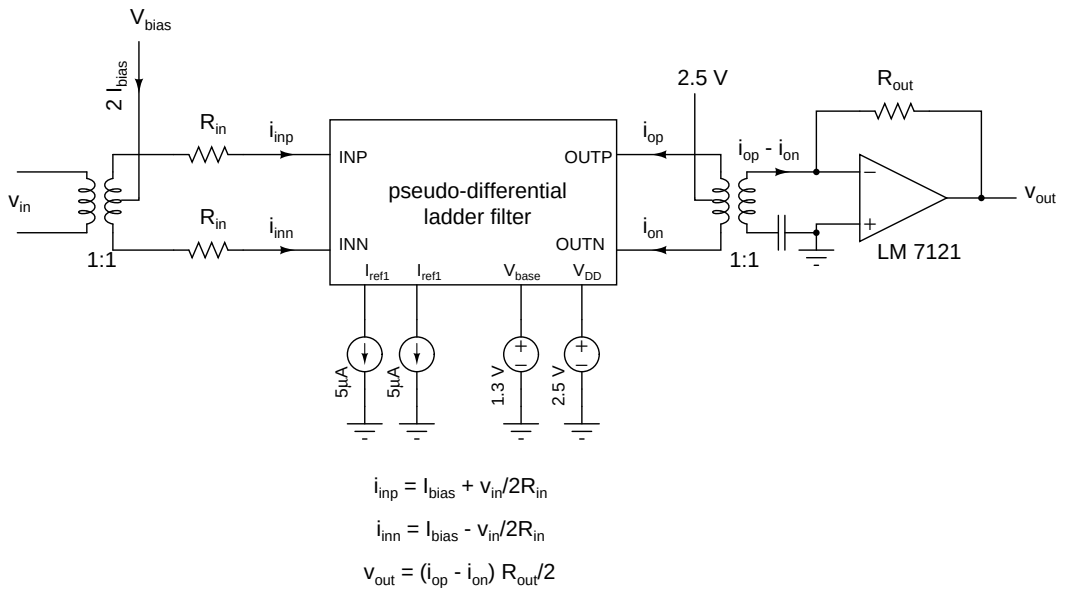

Figure 8.1: Setup used to measure the pseudo-differential ladder filter.

its differential inputs and outputs. The scaling factors due to R_{in} and R_{out}, the attenuation in the cables and the non-flatness in the frequency response of the transimpedance amplifier—if present—are approximately[1] calibrated out using this method.

The measured magnitude response of the filter with I_{bias} varied from 3 μA to 2.5 mA is shown in Fig. 8.2[2]. It follows the expected shape of the Butterworth response (Fig. 7.12). As was anticipated in Chapter 7, the filter is capable of handling bias currents greater than 500 μA. The inset in Fig. 8.2 shows the passband detail of the magnitude response. The measured bandwidth is close to 930 kHz for all bias currents. There is a small variation ($<$ 1 dB) in the passband gain when the bias is varied.

Fig. 8.3 shows the differential output of the filter when signals of identical phase are fed to the two inputs. For this measurement, the two input resistors R_{in} in Fig. 8.1 are connected to one of the ends of the input transformer. Ideally, the differential output

[1]The loop gain around the op-amp is not the same in the two cases. Therefore, the calibration is only approximate.

[2]For values of I_{bias} below 3 μA, the parasitic capacitances on the input leads of the filter on the printed circuit board interfered with the measurements of the frequency response.

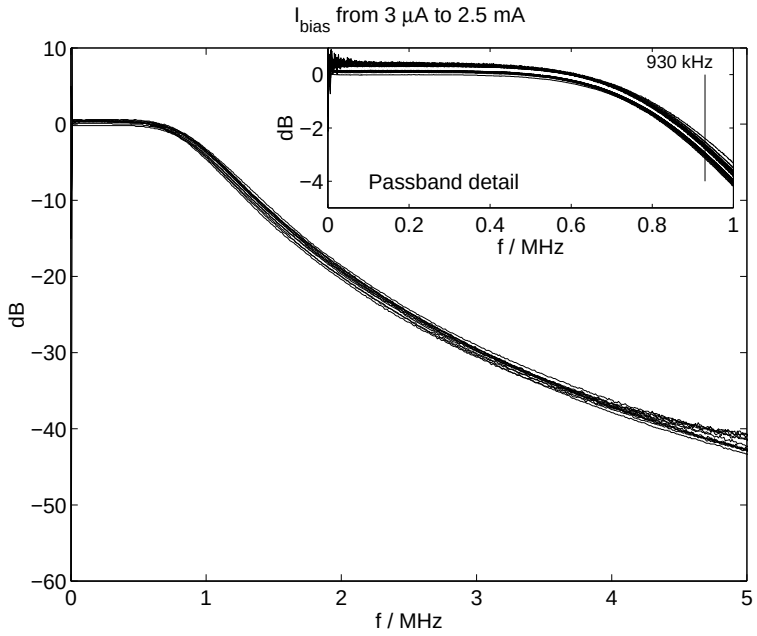

Figure 8.2: Frequency response of the fabricated third order pseudo-differential Butterworth ladder filter.

of the filter should be zero with such common mode inputs. But due to mismatches in the two paths (in the filter as well as in the setup), complete cancellation does not occur. Since the dynamic bias $I_{bias}(t)$ is applied in common mode, such incomplete common mode cancellation manifests itself as leakage of the time-varying bias to the output. The response of the filter to differential input signals is also shown in Fig. 8.3. The measured common mode response in the passband is at least 35 dB below the differential response for all bias currents. The measurement of the common mode response in the stopband is limited by noise.

Output noise

The output noise measured using the test setup in Fig. 8.1 is shown in Fig. 8.4. The input voltage v_{in} is set to zero and I_{bias} is varied. The power spectral density of v_{out} in Fig. 8.1 is multiplied by $(2/R_{out})^2$ to determine the power spectral density of the differential output current noise of the filter. The power spectral density is integrated from 0 to 2 MHz

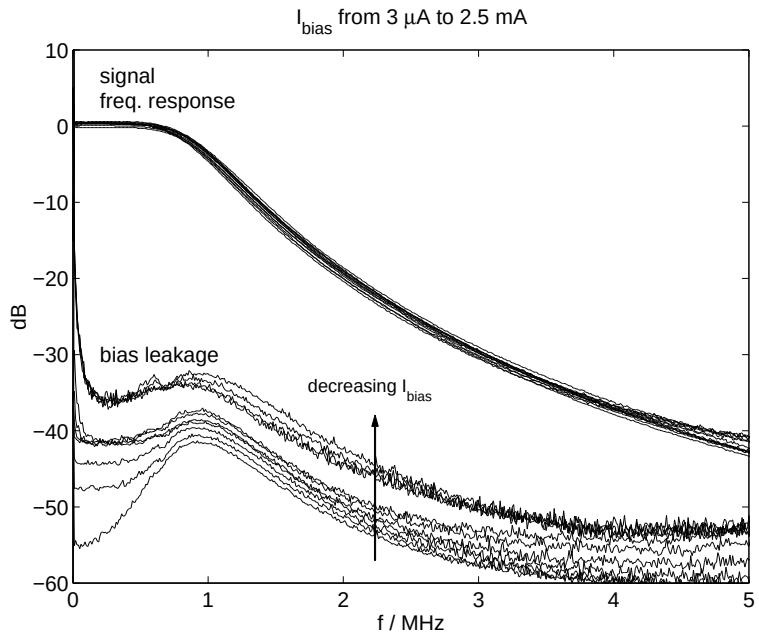

Figure 8.3: Response at the differential output of the filter to common mode and differential inputs.

to obtain the data in Fig. 8.4. The noise from the transimpedance amplifier referred to the differential current output of the filter is also shown on the figure. It is clearly seen that the noise from the setup is dominant for bias currents of 50 μA and less. For larger values of bias currents, the measured noise is reasonably accurate and is seen to be proportional to the bias current. The excessive noise from the setup in Fig. 8.1 is because the low impedance of the output transformer at low frequencies results in a large gain from the input referred voltage noise source of the op-amp to the output. This results in a larger output voltage noise than expected from a transimpedance amplifier that is driven from a current source.

An alternative arrangement to measure the filter's noise is shown in Fig. 8.5. The transformer used at the output of the filter for differential to single-ended conversion in Fig. 8.1 is eliminated. Separate transimpedance amplifiers are used at each of the filter's outputs. The power spectral densities of v_{out1} and v_{out2} are added in a mean square sense

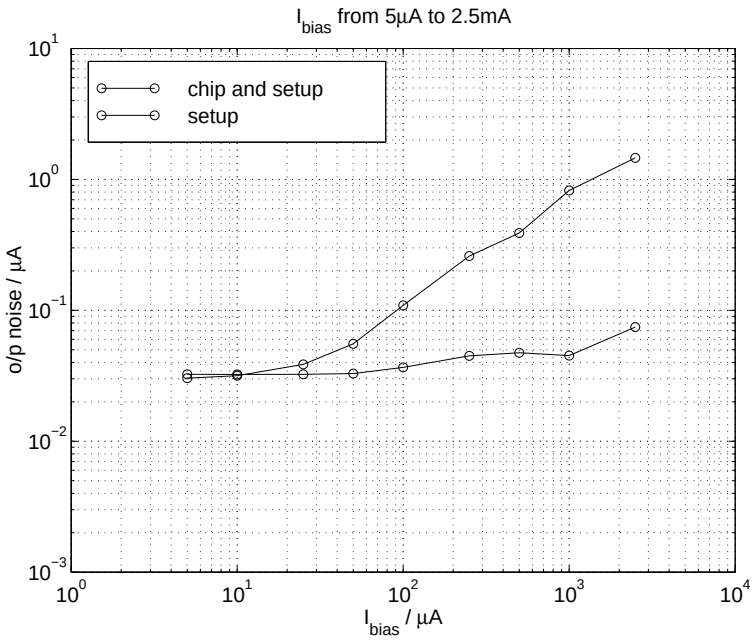

Figure 8.4: Measured output noise.

and divided by ${R_{out}}^2$ to obtain the differential output noise of the filter. The output noise integrated from 0 to 2 MHz is plotted in Fig. 8.6 versus the bias current I_{bias}. The noise from the setup in Fig. 8.5 is also shown. The noise floor of the setup is about[3] 4 nA. The measurement setup in Fig. 8.5 contributes much less noise than the one in Fig. 8.1.

Distortion

As the input frequency is increased, the harmonic distortion of a low-pass filter increases initially and then starts decreasing as the harmonics fall outside the filter's passband. The distortion (both the second- and the third-order) is maximum when the input frequency is somewhere between 1/3 and 1/2 the band edge. In this case, a 400 kHz tone, which results in a large distortion at the output was used for the measurements. Fig. 8.7 shows a family of curves depicting the measured second harmonic distortion versus the bias current at the

[3]This value is consistent with the input referred noise current of $1.9\,\mathrm{pA}/\sqrt{\mathrm{Hz}}$ quoted for LM7121 in its data sheets. The input referred integrated current noise in the 0 to 2 MHz band from two transimpedance amplifiers is $\sqrt{2} \times 1.9 \cdot 10^{-12} \times \sqrt{2 \cdot 10^6} = 3.8\,\mathrm{nA}$. The input referred current noise of the op-amp ultimately limits the lowest value of the current that can be measured using a transimpedance amplifier.

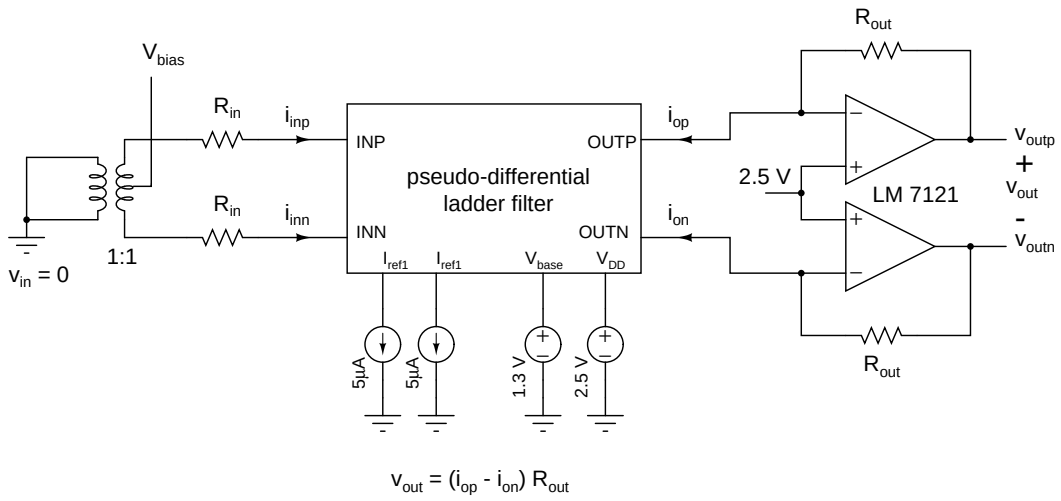

Figure 8.5: Measurement setup used to measure the output noise of the filter.

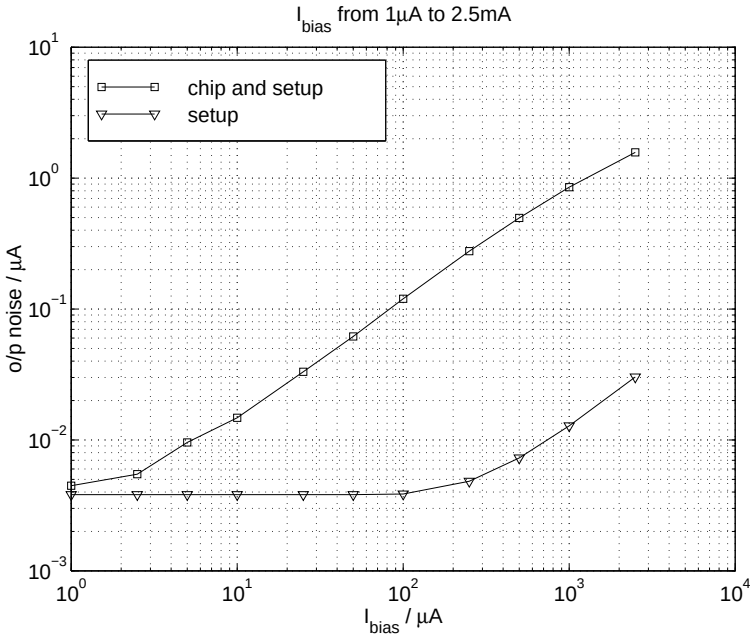

Figure 8.6: Measured output noise.

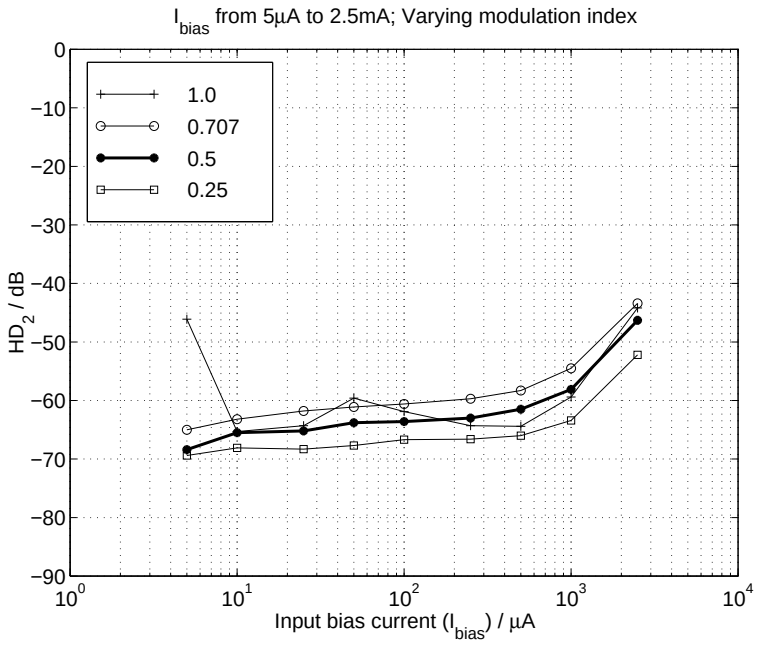

Figure 8.7: Measured second harmonic relative to the fundamental with a 400 kHz input for modulation index values shown in the insert.

input for different values of the modulation index[4]. Similar curves for the third harmonic distortion are shown in Fig. 8.8. The second harmonic is not canceled perfectly at the output due to mismatch and is of a similar magnitude to the third harmonic for the most part. The distortion increases sharply for bias currents more than a mA. This is because of the increasing voltage swing on the collectors of the transistors in the input and the output stages of the filter (Fig. 7.7).

Performance of the filter with suitable biasing

For the following results, the bias current I_{bias} at the input of the pseudo-differential filter is set to twice the value of the single-ended input peak, unless that value is less than 3 μA in which case I_{bias} was maintained at 3 μA. With such a dynamic biasing arrangement, the rms values of the differential output signal and the noise are plotted versus the differential peak input[5] in Fig. 8.9(a). As I_{bias} (= differential peak input) is decreased from

[4]The ratio of the single ended peak input to the bias current I_{bias}.
[5]For values greater than 3 μA the differential peak input is the same as the bias current I_{bias}.

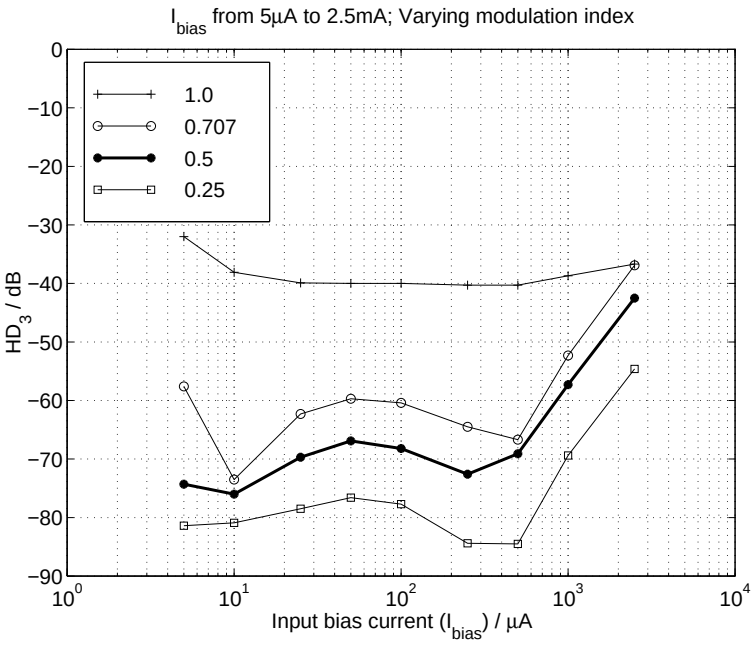

Figure 8.8: Measured third harmonic relative to the fundamental with a 400 kHz input for modulation index values shown in the insert.

its maximum value of a few mA, the output noise decreases down to 4.4 nA rms, corresponding to the smallest I_{bias} of 3 μA. The signal to noise ratio and the signal to total harmonic distortion ratio curves with the biasing arrangement described above are overlaid in Fig. 8.9(b). For differential input peaks greater than 3 μA the latter correspond to the thick lines shown in Figs. 8.7 and 8.8. The signal to noise ratio is lower than the signal to distortion ratio for bias currents below 1 mA. Near constant signal to noise and signal to distortion ratios are maintained for bias (and signal) levels from 1 mA down to 5 μA. The proposed filter maintains $S/N > 0$ dB and THD < -41 dB for total input values ranging over 112 dB (Fig. 8.7(b)). A traditional linear active filter would need a maximum[6] S/N of 112 dB to satisfy the same criterion.

The intermodulation distortion performance of the filter was determined by injecting two tones separated by 40 kHz. Fig. 8.10 shows the ratio of the signal to the third

[6]Measured with an input amplitude that results in an output total harmonic distortion of 41 dB below the fundamental.

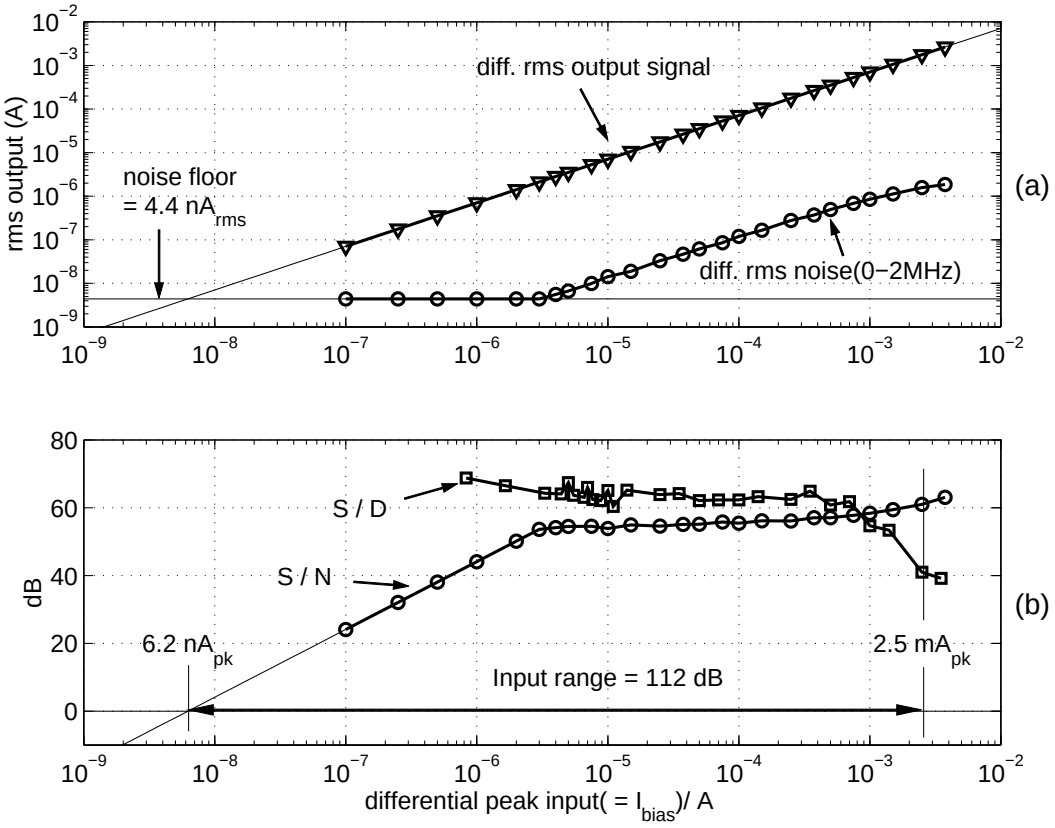

Figure 8.9: Measured noise (in 0-2 MHz band), SNR and THD.

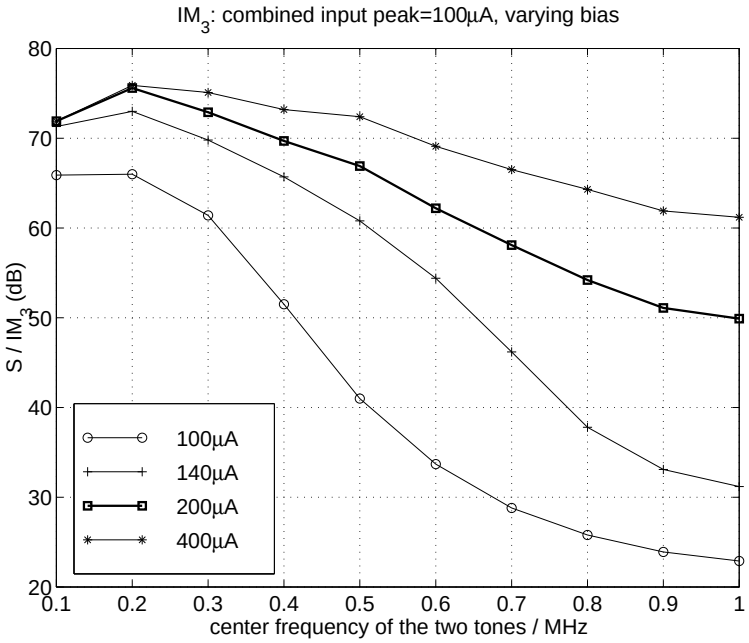

Figure 8.10: Measured third order intermodulation distortion versus frequency with a single ended input peak of 100 μA and values of bias current I_{bias} shown in the insert.

order intermodulation component when the input peak (of the combined two tone signal) is 100 μA and the center frequency of the two tone input signal is varied from 100 kHz to 1 MHz. The bias current I_{bias} is varied to realize different values of the modulation index. The distortion increases with increasing input frequency. For a modulation index of 0.5, i.e. when the single ended input peak is 50% of I_{bias}, the intermodulation distortion is at least 50 dB below the fundamental over the entire band.

Fig. 8.11 shows the intermodulation distortion performance as a function of the differential input signal peak when dynamic biasing is applied as before (i.e. I_{bias} being twice the signal peak subject to a minimum of 3 μA). Three curves corresponding to input frequencies of 100 kHz, 400 kHz and 1 MHz are shown[7]. Below an input amplitude of 1 mA, the ratio of the signal to the intermodulation distortion is more than 40 dB.

[7]The measurement of the intermodulation distortion at 100 kHz and 400 kHz is limited by noise.

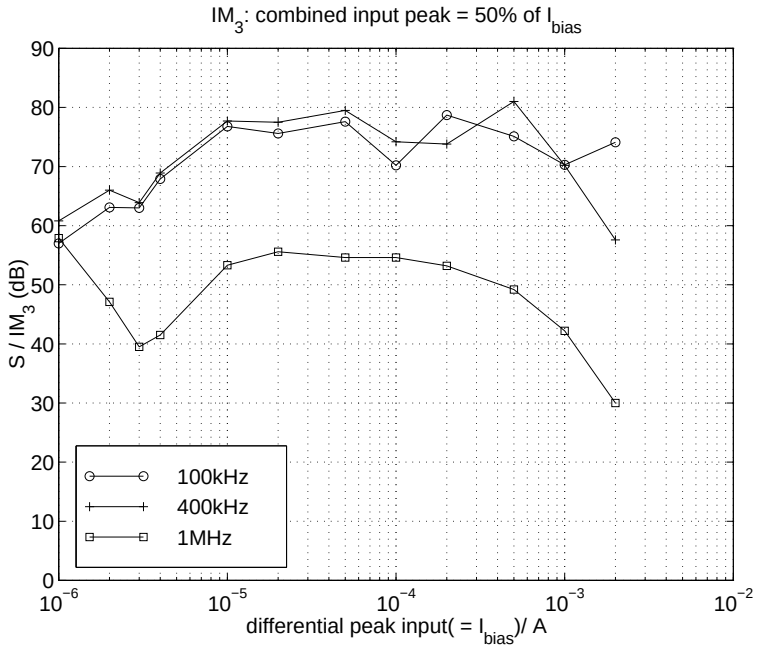

Figure 8.11: Measured third order intermodulation distortion versus input amplitude.

Input-output time invariance

The pseudo-differential log-domain filter is linear and time-invariant between its differential inputs and outputs even in presence of a time-varying bias $I_{bias}(t)$ (chapter 3). To verify this, an input 600 kHz with a single ended peak of 20 μA was fed to the filter and the bias current was switched between 24 μA (which is 20% larger than the *single-ended* peak input) and 114 μA. Fig. 8.12 shows the switched bias current and the differential output. It is seen the output is practically unaffected by transients in I_{bias}. The results in Figs. 8.3 and 8.12 point to the input-output linearity time-invariance of the filter in presence of a varying I_{bias}.

Power consumption

The measured current and the power consumed by the pseudo-differential filter is shown in Fig. 8.13. It follows the expected curve shown in Fig. 7.19. When I_{bias} is at its maximum value of 2.5 mA, the total power consumed is about 26.1 mW. In its quiescent condition,

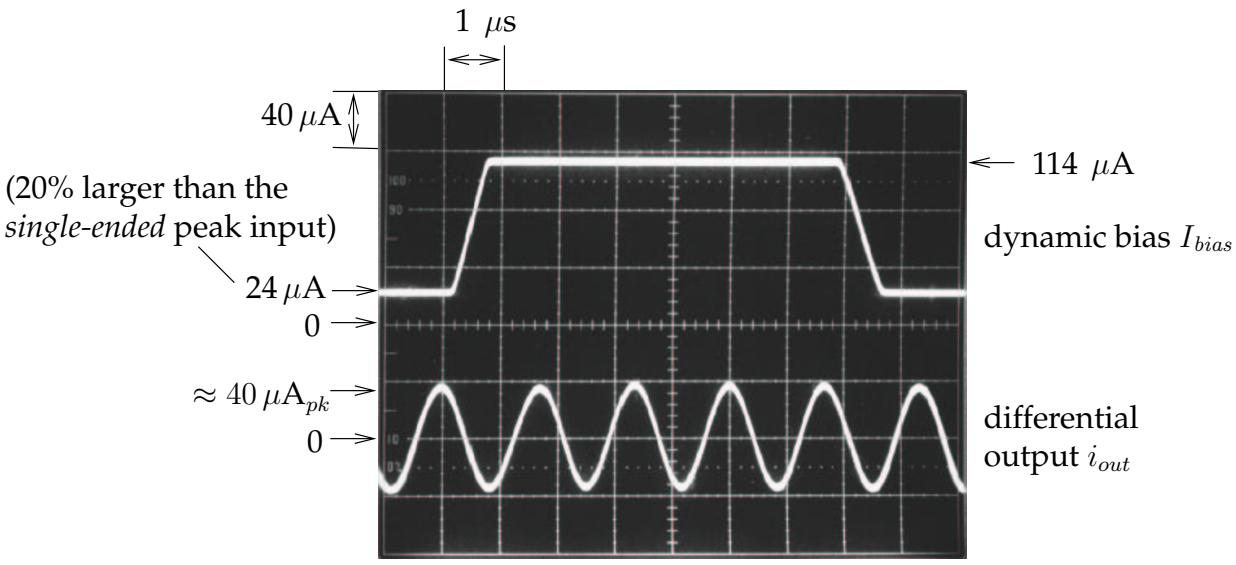

Figure 8.12: Time varying bias current and the resulting differential output of the pseudo-differential filter.

the filter draws 575 μW of power.

Power consumption and dynamic range: comparison to passive filters

The measured performance of the pseudo-differential filter can be compared to that of an equivalent passive RC filter along the lines of section 7.2.9. Fig. 8.14(a) shows a pseudo-differential filter biased using an ideal peak detector such the bias current at the input is twice the peak value of the single ended input currents. Fig. 8.14(b) shows a current mode passive RC filter that has a bandwidth of 930 kHz and a minimum signal to noise ratio of 53.7 dB for an input current of 3 μA peak[8]. As was done in section 7.2.9, the pseudo-differential ladder filter and the RC filter in Fig. 8.14 are compared in the following paragraphs.

Fig. 8.15 shows the variation of signal to noise ratio with the input peak current for the two filters in Fig. 8.14. At 3 μA input current, the two filters have equal signal to noise

[8]The values of R and C in this filter are quite different from those in Fig. 7.18(b). This is because, in the simulations, the lower limit on the bias current, which translates to the lower limit on the signal current, was 0.5 μA whereas in the experiments, it was 3 μA. For a given S/N, the noise floor required depends on the minimum signal current and hence, the different component values.

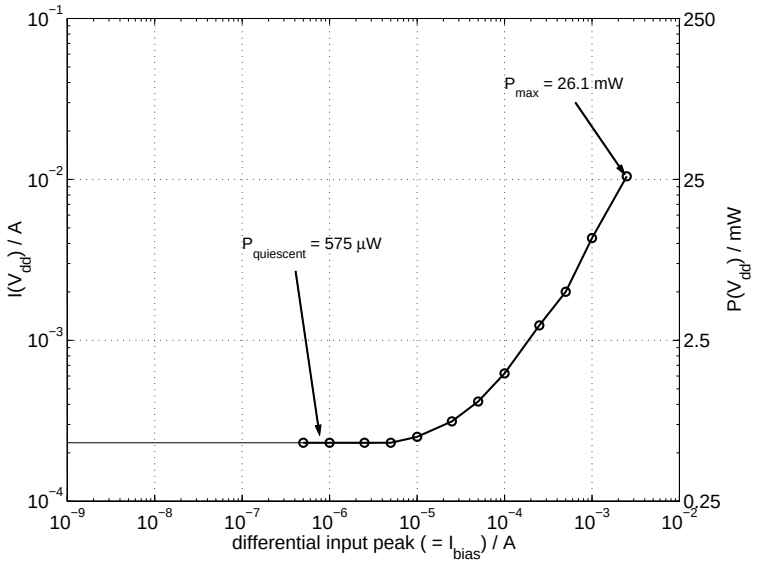

Figure 8.13: Measured current and power consumption of the filter.

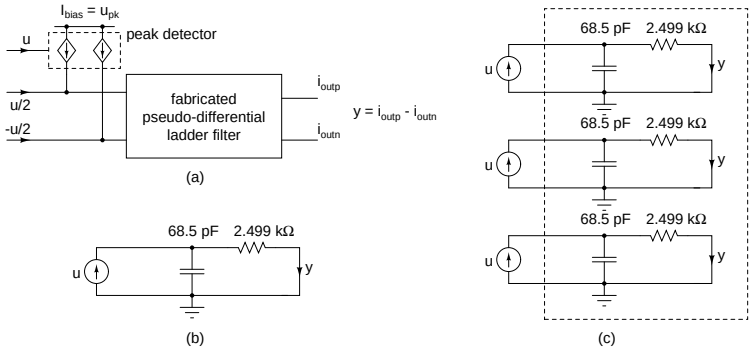

Figure 8.14: (a) Pseudo differential filter biased from an ideal peak detector, (b) Current mode RC filter with $S/N = 53.7$ dB for a 3 μA peak input, (c) Three first-order RC filters.

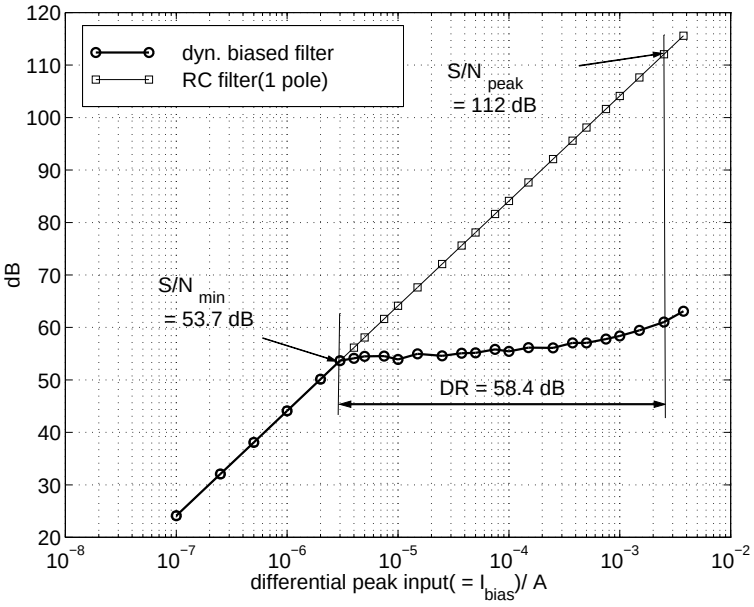

Figure 8.15: Signal to noise ratios of the pseudo-differential filter and the RC filter in Fig. 8.14.

ratios. At the maximum input current of 2.5 mA peak, the signal to noise ratio of the RC filter is 112 dB. The power dissipation of the two filters in Fig. 8.14 is shown in Fig. 8.16 as a function of the peak value of the input u. The qualitative behavior is similar to that in Fig. 7.21. At the highest input current of 2.5 mA, the pseudo-differential filter dissipates about 2.3 times more power than the RC filter.

Comparison to published filters

Table A.2, which lists the power dissipation, the bandwidth, the order, and the dynamic range of several previously published power efficient filters is repeated here as Table 8.1 for convenience. The table also shows these quantities for the pseudo-differential ladder filter presented here. The dynamic range specified is the range of input signals over which THD ≤ 40 dB[9] and $S/N > 0$ dB are maintained. Of the listed filters 3, 5, and 8 are companding (log-domain class-AB) filters and the rest are conventional linear filters. For the

[9]Except for 7 [68] which uses the criterion THD ≤ 49 dB

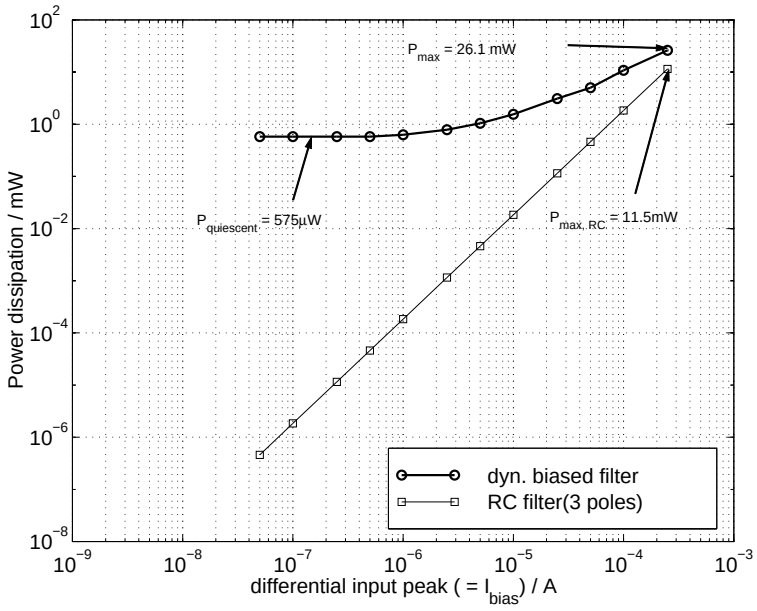

Figure 8.16: Power dissipation versus the input signal: Pseudo-differential ladder filter and passive RC filter.

latter, the dynamic range as defined above is also the S/N when the THD is 40 dB[9] (See Chapter 2).

For the filters listed in Table 8.1[10], the power dissipation normalized to the bandwidth and the order is plotted in Fig. 8.17 versus the dynamic range of the filter. The straight line corresponds to the passive RC filter.

A measure of the efficiency of a filter is the ratio of the power consumed by the first-order RC filter (Appendix A) to the normalized power consumed by the filter in question. The efficiencies and dynamic ranges of the filters in Table 8.1 are shown in Fig. 8.18. The proposed filter represents over an order of magnitude improvement in power efficiency.

Fig. 8.19 shows the photograph of the chip. The entire circuit excluding the pads occupies 0.52 mm^2. The area is dominated by the metal-metal capacitors. Table 8.2. summarizes the measured performance of the chip.

[10]In cases where two values are given for the same filter in Table A.2, the "better" number is plotted.

Table 8.1: Power dissipation per S/N , signal frequency and filter order for published active filters.

	V_{dd} (V)	Power	f$_{-3dB}$	ord.	DR (dB)	power per pole, f$_{-3dB}$ ($\times 10^{-12}$ J)	power per pole, f$_{-3dB}$, DR ($\times 10^{-18}$ J)	Ref.
1	1.0	10.5 μW	100 kHz	5	68 (max.) 57 (min.)	21.0 21.0	3.3 41.9	[69]
2	2.5	40 μW	70 kHz	2	75	285.7	9.03	[70]
3	1.2	65 μW^{q} 170 μW^{m}	320 kHz	3	65	67.7^{q} 177.1^{m}	21.4^{q} 56.0^{m}	[19]
4	1.5	375 μW	525 kHz	5	67	142.9	28.5	[71]
5	1.2	6.5 mW^{q}	30 MHz	3	62.5	72.2	40.6	[22]
6	1.2	23 μW	320 kHz	3	57	24.0	47.9	[20]
7	2.5	13 mW	600 kHz	7	77^{x} 71^{x}	3095.0 3095.0	61.8^{x} 240.3^{x}	[68]
8	1.2	6.5 mW^{q}	100 MHz	3	50	21.7	217.0	[22]
9	5	580 μW	40 MHz	2	41.3	7.25	537.5	[72]
10	2.5	21.6 mW^{m}	930 kHz	3	112	9350	0.059	This work

^{q} In quiescent condition.
^{m} With the maximum input signal.
^{x} [68] quotes a maximum signal of 2 V_{pp}, a noise floor of 196 μV_{rms} and a dynamic range of 77 dB. These numbers are inconsistent. The value corresponding to the quoted maximum signal and noise is 71 dB.

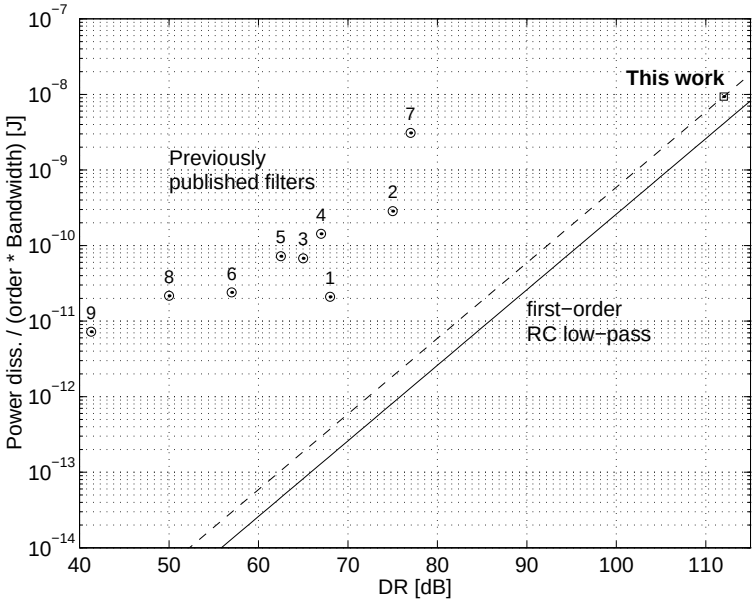

Figure 8.17: Power dissipation per pole and edge frequency vs. dynamic range

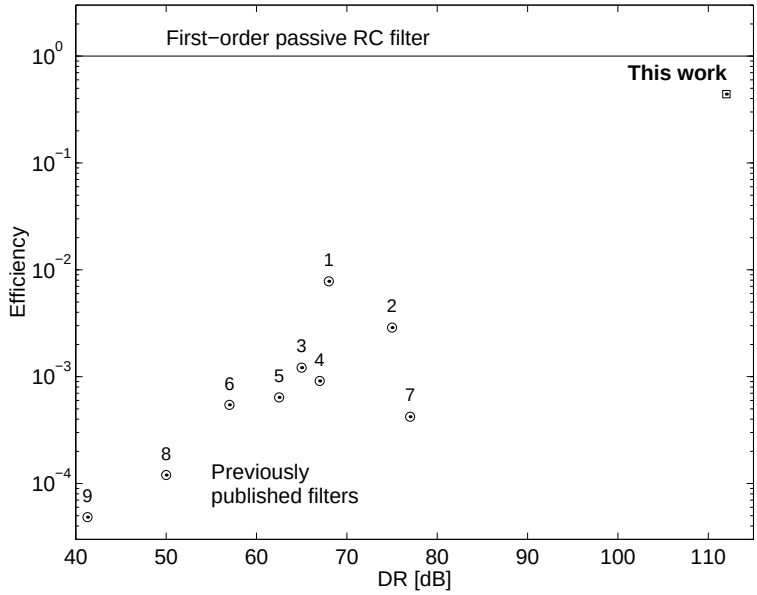

Figure 8.18: Efficiency when compared to a first-order passive RC filter.

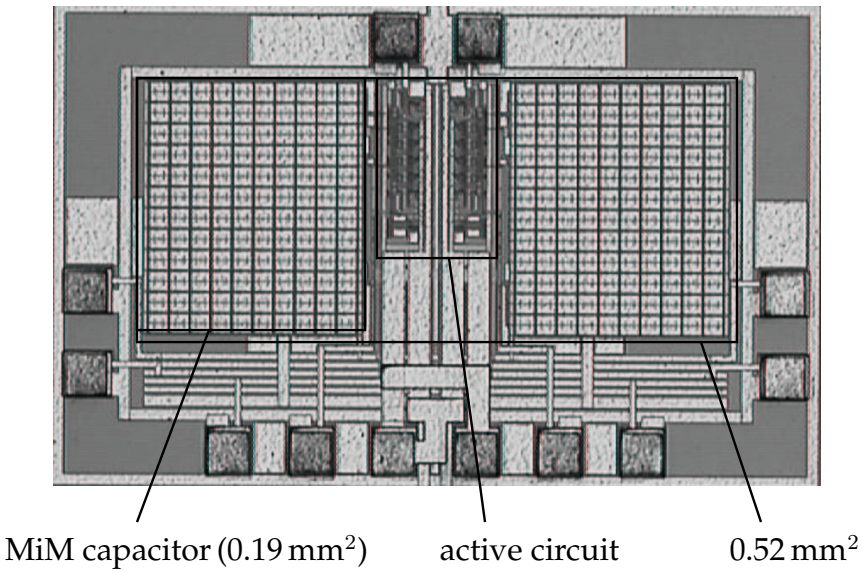

Figure 8.19: Photograph of the third-order pseudo-differential Butterworth ladder filter.

Table 8.2: Performance summary: Third-order Butterworth filter (Fig. 7.8).

Technology	0.25 μm BiCMOS	
Area (excl. pads)	0.52 mm^2	
Supply voltage	2.5 V	
-3 dB BW†	930 kHz	
I_{bias}	3 μA	2.5 mA
Power diss.	575 μW	26.1 mW
Output noise	4.4 nA	1.5 μA
THD	-64.3 dB	-41 dB
Dynamic range‡ DR	112 dB	
$\frac{\text{Max. Power Diss.}}{\text{Order}\cdot\text{BW}}$	9.35 nJ	
$\frac{\text{Max. Power Diss.}}{\text{Order}\cdot\text{BW}\cdot\text{DR}^\ddagger}$	5.9×10^{-20} J	

†BW: Bandwidth
‡ Dynamic range, see text;

8.1.1 Dynamic range of a companding filter

The proposed filter maintains $S/N > 0$ dB and THD < -41 dB for total input values ranging over 112 dB (Fig. 8.9(b)). It is emphasized here that the proposed filter is not equivalent to a conventional filter with a 112 dB dynamic range as the latter would have $S/N = 112$ dB with the largest input. However it would need orders of magnitude larger power dissipation in order to achieve this [1]. The dynamically biased filter presented here is suitable for cases where a modest S/N and a near-optimum power dissipation must be maintained over a large range of input amplitudes.

If the sum of two signals is present at the input, it is the peak value of the *sum* that determines the output noise in the proposed filter. Thus a 1 μA signal appearing simultaneously with a 2.5 mA signal results in an output noise of about 1.5 μA (Fig. 8.9(a)) rms and a signal to noise ratio of about 60 dB (Fig. 8.9(b)) for the *total* signal. Thus, in situations where signals separated by more than 60 dB appear simultaneously at the input, this filter

processes the larger signal with sufficient fidelity and not the smaller one. On the other hand this filter is well-suited for spectral shaping applications in which the total input signal is sometimes large and sometimes small, as in the latter case a large reduction in the power consumption and the output noise can be achieved.

8.2 Peak detector

8.2.1 Single-ended operation

Fig. 8.20 shows the setup used to characterize the peak detector (Fig. 7.38). The bias currents on chip are mirrored from two external 5 μA sources (The bias generation circuit is shown in Fig. 7.11). A 2.5 V supply is used for the chip. The capacitor used for the filter in the peak detector's loop is 1 nF and is connected as shown in the figure. The ac coupled input voltage V_{in} is converted to an input current using a series resistor R_{in}. As mentioned in section 7.4.3, the input node of the peak detector experiences voltage variations. With a large input amplitude, the voltage across R_{in} is nearly equal to V_{in} and input current equals V_{in}/R_{in}. Values of R_{in} between 5 kΩ and 10 MΩ are used to vary the input current from a few mA down to a fraction of a μA. The output current I_{out} is converted to a voltage V_{out} using an opamp with a resistor R_{out} in feedback. The capacitor C_{out} provides first-order filtering of the output current and is used while measuring its dc component.

Fig. 8.21 shows the gain of the peak detector (i.e. the ratio of the output dc component to the input amplitude) for inputs at dc, 200 kHz and 1 MHz. Although the peak detector would never be used with dc inputs, testing with them helps establish the basic functionality. The input is varied from 1.4 μA to 2.8 mA. The measured gain is close to, but less than the desired value of 1.5. The deviation from 1.5 is conjectured to be due to a combination of mismatch and measurement errors. As can be seen, the ratio is close to

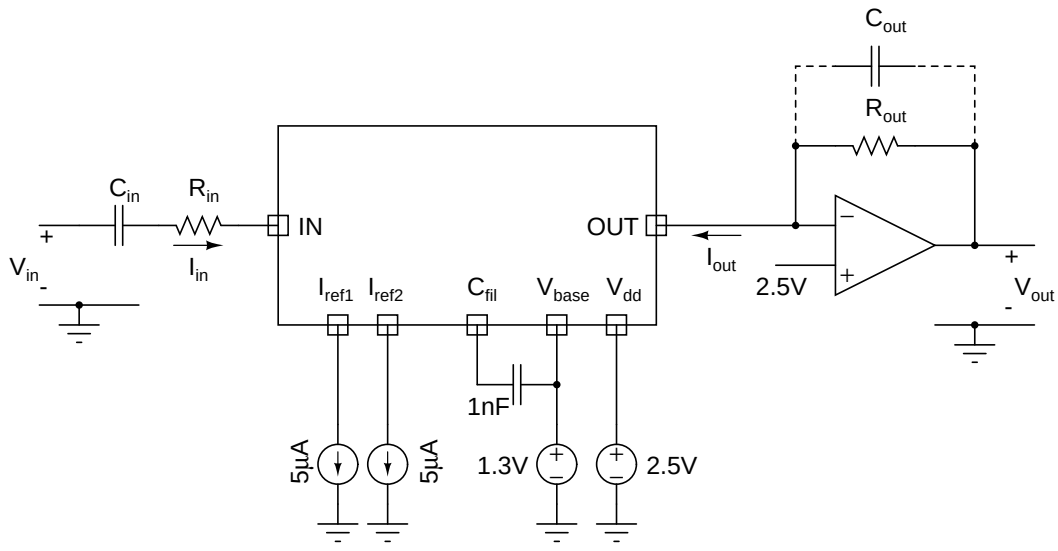

Figure 8.20: Measurement setup for the peak detector.

1.5 for sinusoidal inputs at 200 kHz and 1 MHz as well for input above 10 μA. The measurement setup in Fig. 8.20 is unsuitable for low current ($< 10\,\mu$A) measurements at "high" frequencies (beyond a few tens of kHz). This is because, to supply small currents, the input resistance R_{in} in Fig. 8.20 has to be very large and its shunt capacitance tends to dominate over the resistive part, resulting in inaccurate input voltage to current conversion. Therefore, the smallest current shown in the curves for 200 kHz and 1 MHz inputs in Fig. 8.21 is 14 μA.

Fig. 8.22 shows the peak-peak ripple at the output as a fraction of the output dc component for input frequencies of 200 kHz and 1 MHz. As can be expected, the ripple is lower with a higher frequency input. The peak-peak ripple is very large with an input frequency of 200 kHz. About half of the measured peak-peak is contributed by very narrow spikes which occur each time the capacitor is charged. They can be filtered using a low pass filter. Since the spikes have their energy concentrated at high frequencies, the low pass filter can have a relatively high cutoff frequency so that the dynamic behavior (attack and decay) of the peak detector is not affected. In order to use this peak detector at low frequencies,

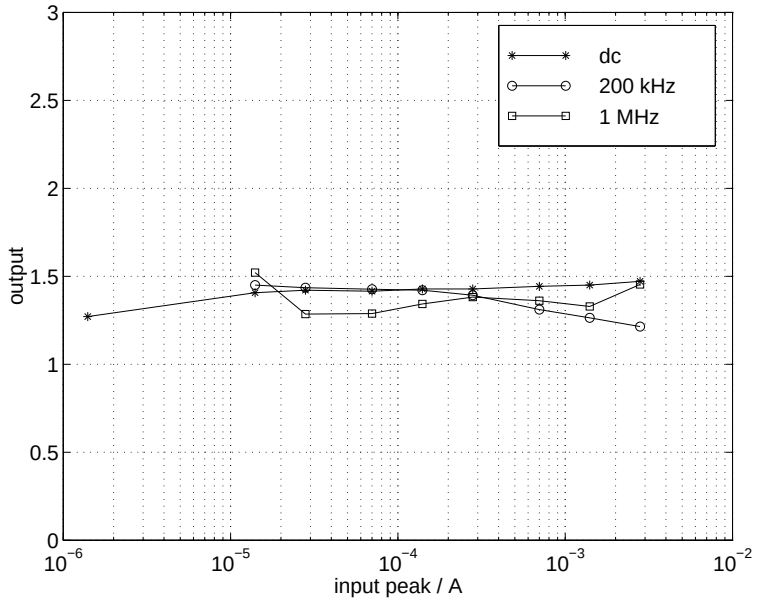

Figure 8.21: Output dc component with dc, 200 kHz and 1 MHz inputs.

either the filtering capacitance of the peak detector must be increased or an additional low pass filter must be used at the output of the peak detector.

Fig. 8.23 shows the frequency response (i.e. variation of the gain with the input frequency) of the peak detector. It can be seen that a relatively flat response is maintained up to 2 MHz. As expected, the gain decreases at low frequencies. This is due to increased ripple. For the lowest input current (10 μA), the gain increases with frequency. This is due to the capacitance of the 1 MΩ input resistor. The gain also increases with frequency for the highest input current of 2 mA. The reason for this behavior is not clear.

The dynamic performance of the peak detector is measured by feeding a 200 kHz sinusoid whose envelope is a square wave at 2 kHz. The amplitude of the sinusoid alternates between I_{pk} and $0.5\,I_{pk}$ as shown in Fig. 8.24. Fig. 8.25 shows the attack and decay times plotted versus I_{pk}. The attack and decay times are the time intervals between between 10% and 90% of the output step[11]. As expected, the attack is much faster than the

[11]Due to the ripple in the output, very precise measurements are not possible. The local average of the

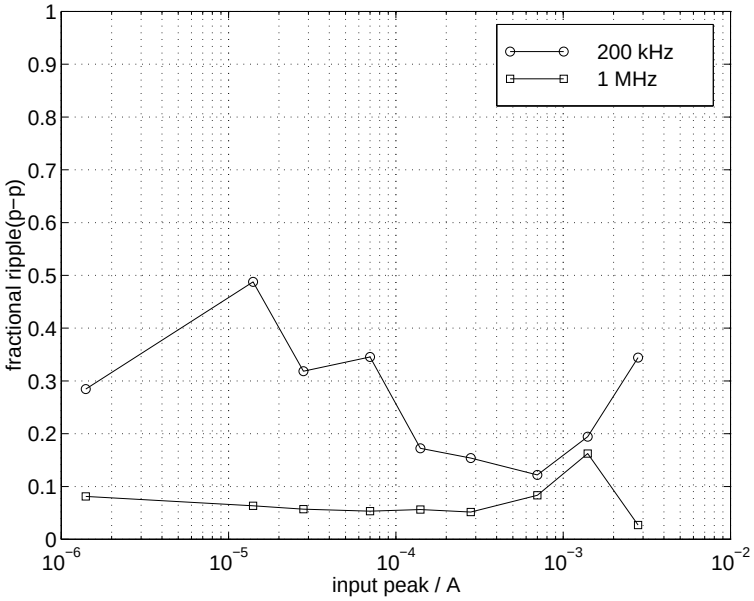

Figure 8.22: Peak-peak ripple at the output with 200 kHz and 1 MHz inputs.

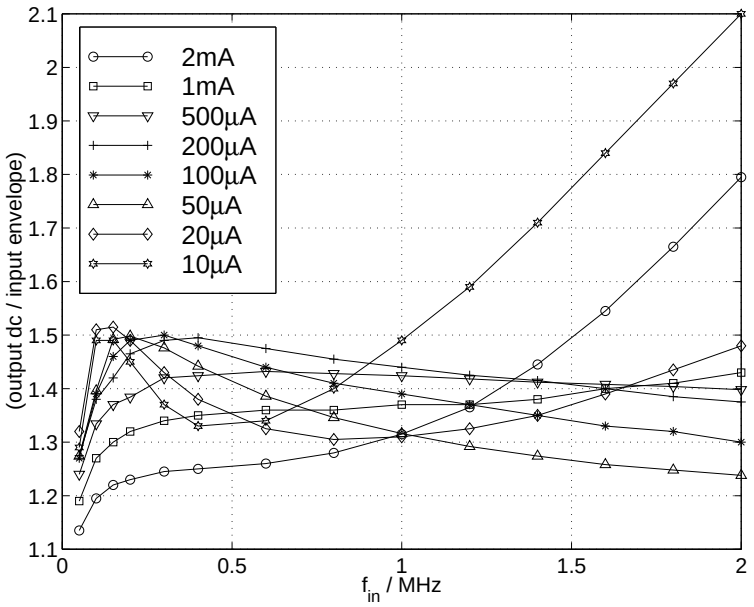

Figure 8.23: Frequency response of the peak detector.

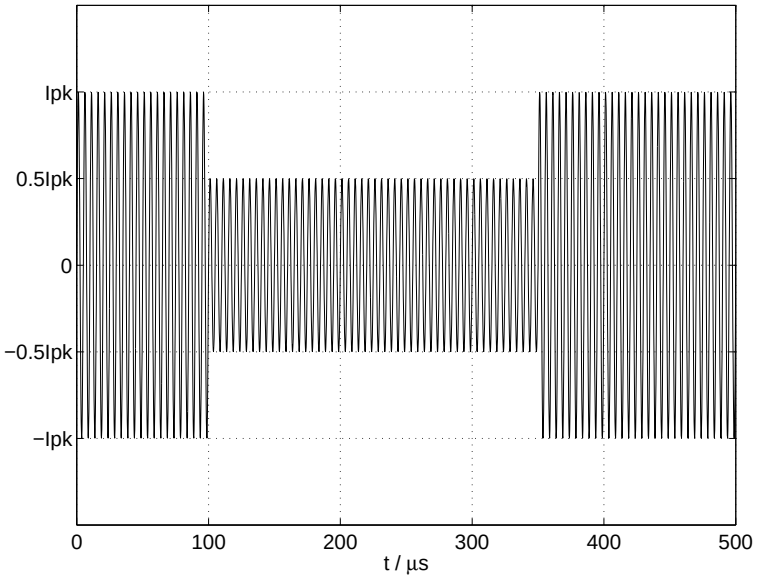

Figure 8.24: Input used to measure the attack and the decay times of the peak detector.

decay. As mentioned in section 7.4, the time constant of the first-order filter in the peak detector is 43.2 μs. With this time constant, the time taken for the output to decay from 90% to 10% when the input amplitude abruptly decreases by a factor of 2 would be 23.6 μs. The measured decay time is larger than this value for very small and very large currents. The possible reasons are:

- With very small input currents, extra band-limiting in the loop may be slowing down the circuit to some extent.
- If the ripple is large (as it is, with large currents), the actual step in the output current is greater than $0.5I_{pk}$ and the time taken for this decay is correspondingly longer.

The current consumed by the peak detector (including the output current I_{out} in Fig. 8.20) is plotted versus the input amplitude in Fig. 8.26. The quiescent current drawn from the power supply is about 65 μA. Out of this, 20 μA is consumed in the bias generation circuit (Fig. 7.11) which is the same as that used for the Butterworth filter. Since

waveform on the oscilloscope screen was established visually.

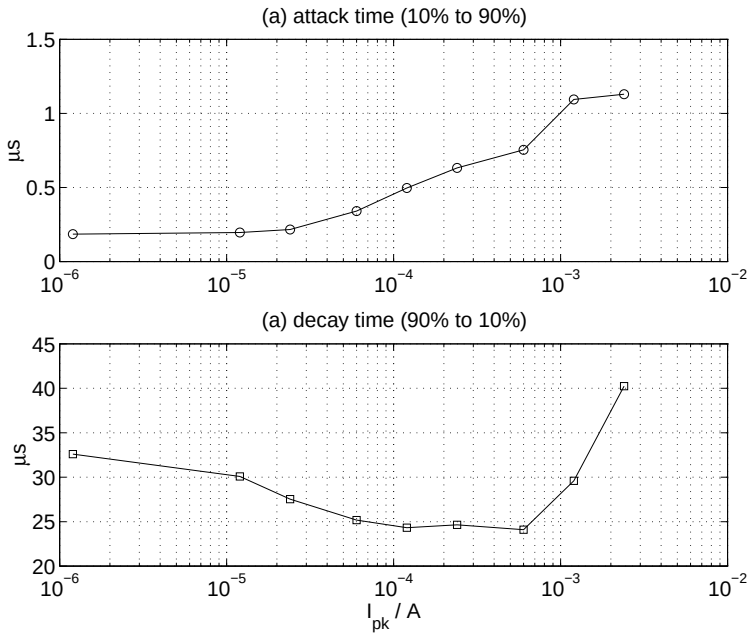

Figure 8.25: (a) Attack time and (b) Decay time of the peak detector as a function of the bias current.

the internal current levels in the peak detector are much lower ($\approx 0.6\,\mu$A), the quiescent current can be reduced by about a fourth by redesigning the bias circuit. The current consumed increases with increasing output current. It does not vary appreciably with input frequency.

Fig. 8.27 shows the chip photograph. Table 8.3 summarizes the performance of the single ended peak detector.

Table 8.3: Performance summary: Single ended peak detector.

Technology	0.25 μm CMOS
Chip area (excl. pads)	0.12 mm^2
Supply voltage	2.5 V
Quiescent power dissipation	162.5 μW
Input range	$\approx$ 1.4 μA to 2.8 mA
Attack time	$\leq$ 1.2 μs
Decay time	$\leq$ 40 μs
Bandwidth	$>$ 2 MHz
Envelope gain	1.3–1.5

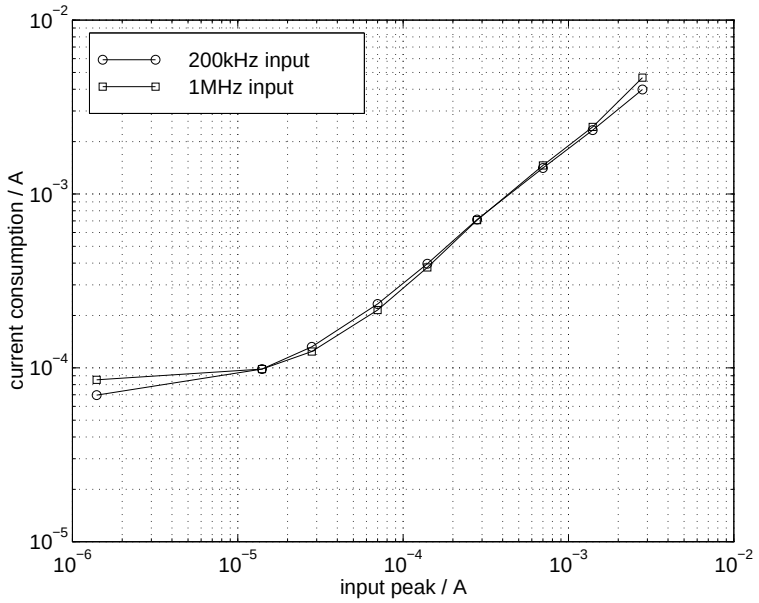

Figure 8.26: Current consumption of the peak detector.

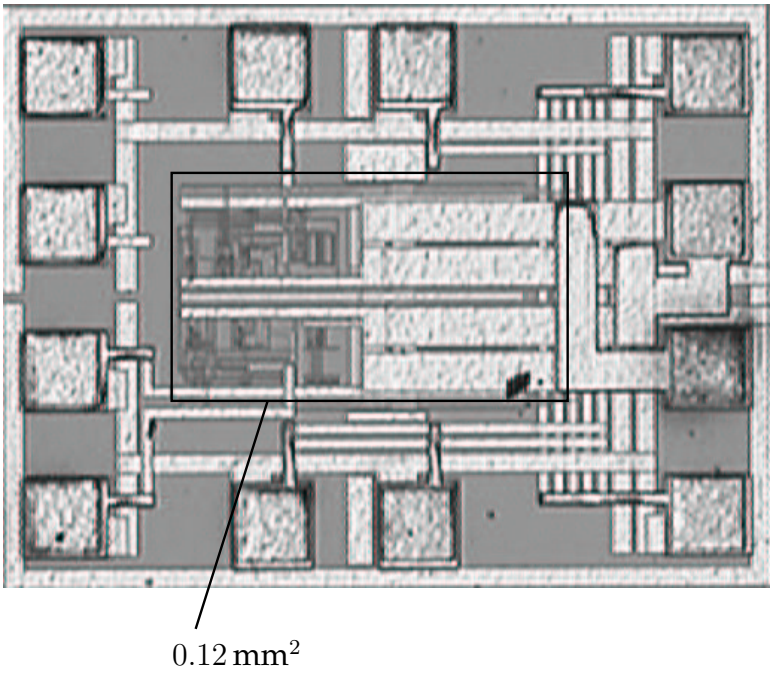

Figure 8.27: Photograph of the single ended peak detector chip.

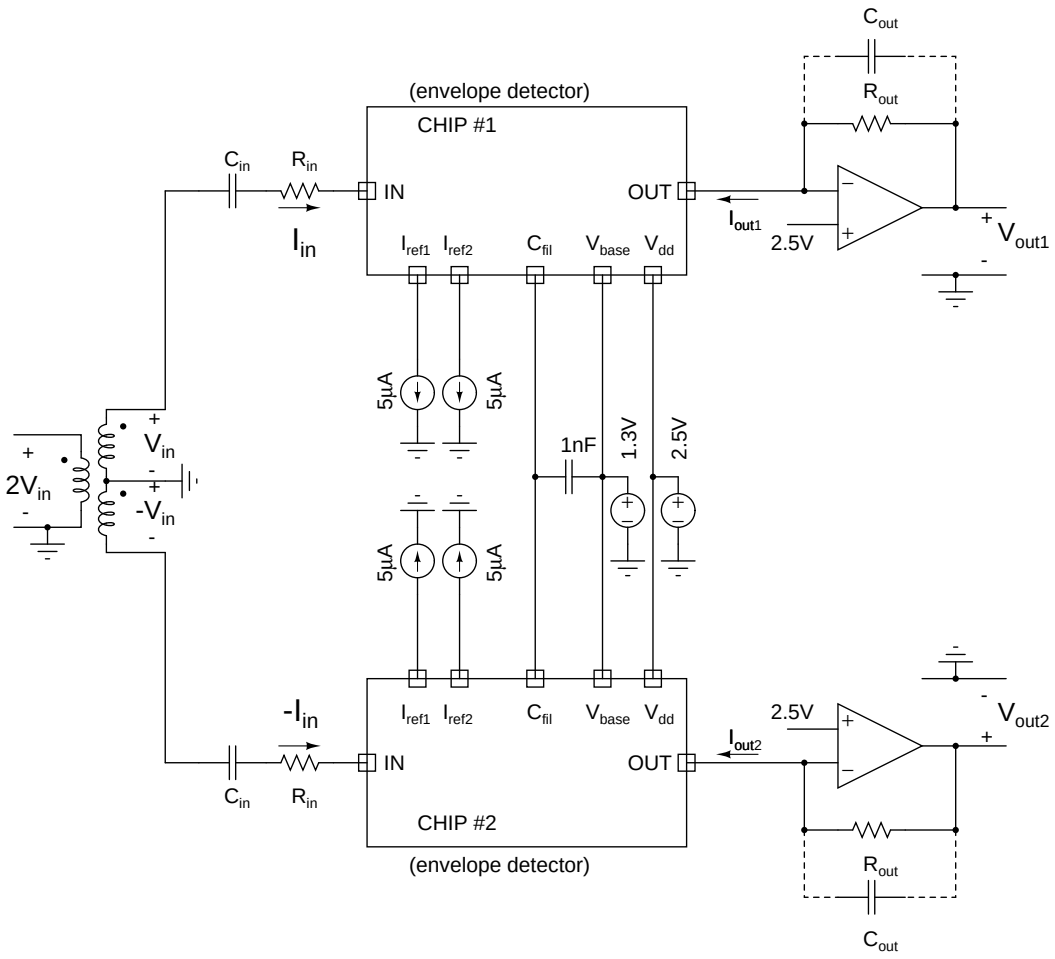

Figure 8.28: Setup to measure a differentially responding peak detector.

8.2.2 Differential operation

The arrangement in Fig. 8.28 is used to characterize the differentially responding peak detector proposed in section 7.4.5. Two chips are connected to a single 1 nF filtering capacitor and are fed with differential inputs. The outputs I_{out1} and I_{out2} of the two chips are identical under perfectly matched conditions and either of them can be taken as the output for dynamically biasing the log-domain filter (chapter 3).

The input voltage V_{in} and the input resistor R_{in} are adjusted such that when ac coupled, the peak value of I_{in} (and $-I_{in}$) is 120 μA for any duty cycle. Fig. 8.29 shows such a signal when its duty cycle changes from 0.2 to 0.6. The signals have unequal positive

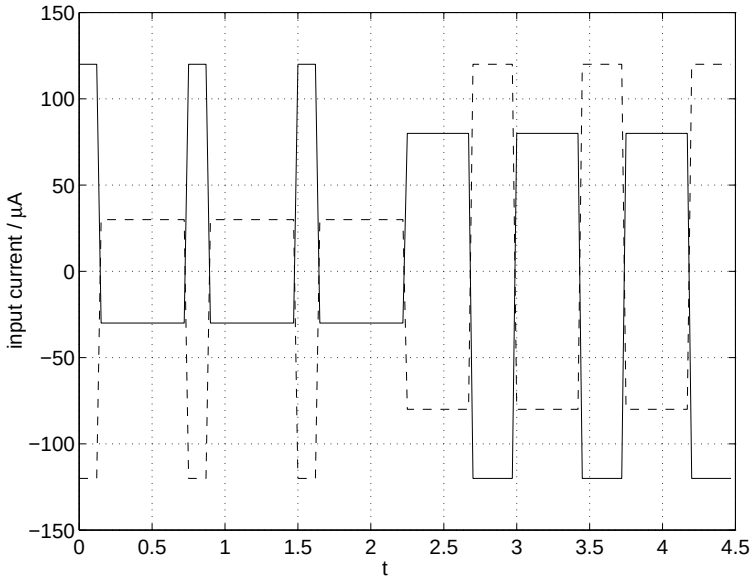

Figure 8.29: Inputs with a constant amplitude and a varying duty cycle (arbitrary X-axis).

and negative peaks, but the larger of the two is 120 μA. For duty cycles less than 0.5, the positive peak is larger than the negative peak. The input frequency is 200 kHz.

Fig. 8.30 shows the gain of the peak detector (ratio of output dc to 120 μA) as the input duty cycle is varied from 0.1 to 0.9. The measured gain is nearly a constant. Since the two peak detectors are coupled, the output is the larger of the positive and negative input peaks.

For comparison, differential inputs with a variable duty cycle are fed to uncoupled peak detectors, each with a 1 nF filtering capacitor. In this condition, the peak detectors measure the positive peak of the respective input currents. The measured gain is plotted versus the duty cycle in Fig. 8.31. As expected, the gain of "chip #1" drops off for duty cycles larger than 0.5. "chip #2" shows the opposite behavior.

These results prove the feasibility of the differential responding peak detector. Detection of the larger of the positive and negative peaks of the input makes this circuit suitable for generating an appropriate dynamic bias for filters with differential inputs when widely

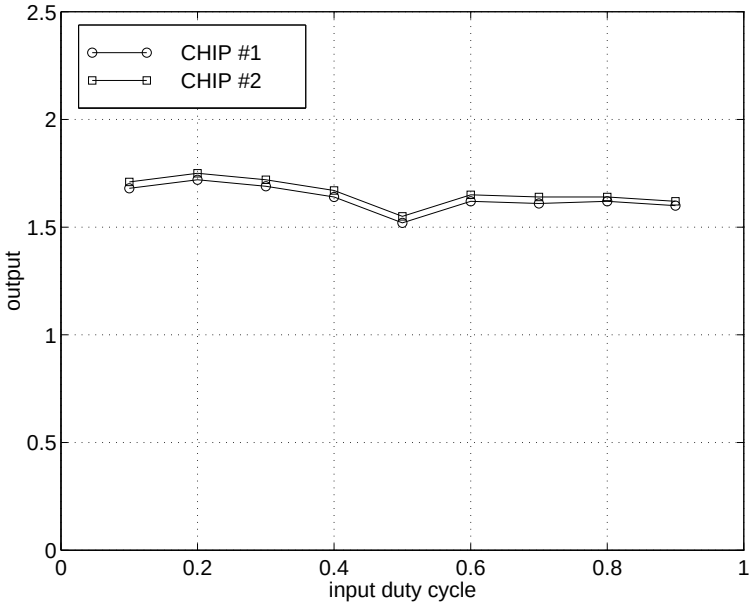

Figure 8.30: Outputs of two coupled peak detectors operating in differential mode.

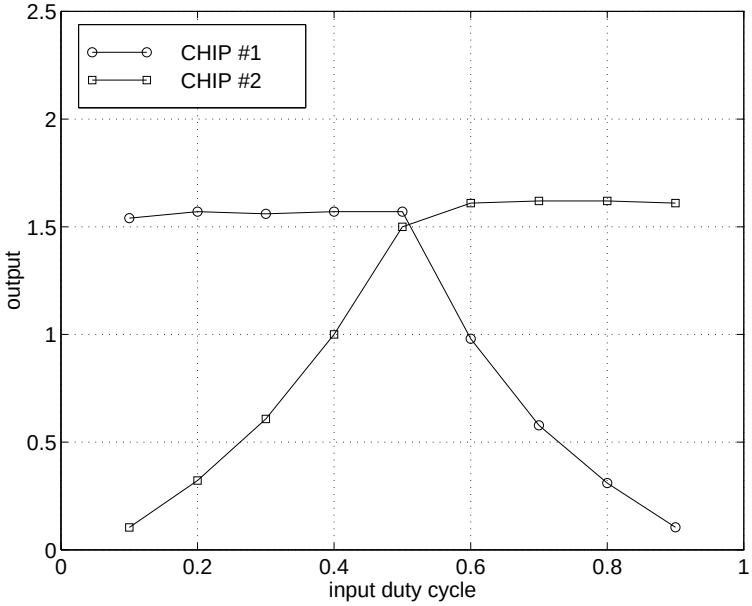

Figure 8.31: Outputs of the two peak detectors when they are uncoupled.

varying duty cycles are expected at the input.

Fig. 8.30 shows an asymmetry in the output around a duty cycle of 0.5. Also the outputs from the two chips are unequal. This is due to mismatch between the chips. Since the two chips are in separate packages, differences in temperature between accentuate the mismatch. Better results can be expected when both halves of the circuit are integrated on the same die.

8.3 Second-order filter using lateral PNPs and MOS capacitors

Fig. 8.32 shows the DC characteristics of the lateral bipolar transistors with the gate tied to the base as described in section 5.2. The variation of the collector current I_C vs. base emitter voltage V_{BE} is plotted on a log scale for two values of collector-emitter voltage and two different base widths (gate lengths). As expected the transistor with a narrower base has a higher saturation current, and a lower Early voltage (greater separation between the curves for the same change in collector-emitter voltage). Log conformity is good up to a collector current of about 0.5 μA for both transistors. The deviation from the exponential seems to be greater for the transistor with the narrower base. The measured slope factors[12] of the transistors with base-widths of 0.32 μm and 0.36 μm are 1.08 and 1.04 respectively. The transistor with the wider base (0.36 μA) was used in the fabricated filter described in section 7.6. Its Early voltage[13] V_A is about 2.3 V.

Fig. 8.33 shows the experimental setup used to characterize the filter in Fig. 7.47(b). The internal bias current I_{tune} that determines the filter's time constants and the input bias current I_{bias} (see Fig. 8.33) can be varied using V_{tune} and V_{bias} respectively. In all of the following measurements, I_{tune} and I_{bias} are set to equal values. Note that this filter is not

[12] The collector current of the transistor is given by $I_c = I_s \exp(V_{BE}/\eta V_t)$ where η is the slope factor.
[13] The small signal output conductance of a transistor is given by $I_c/(V_A + V_{CE})$ in the active region.

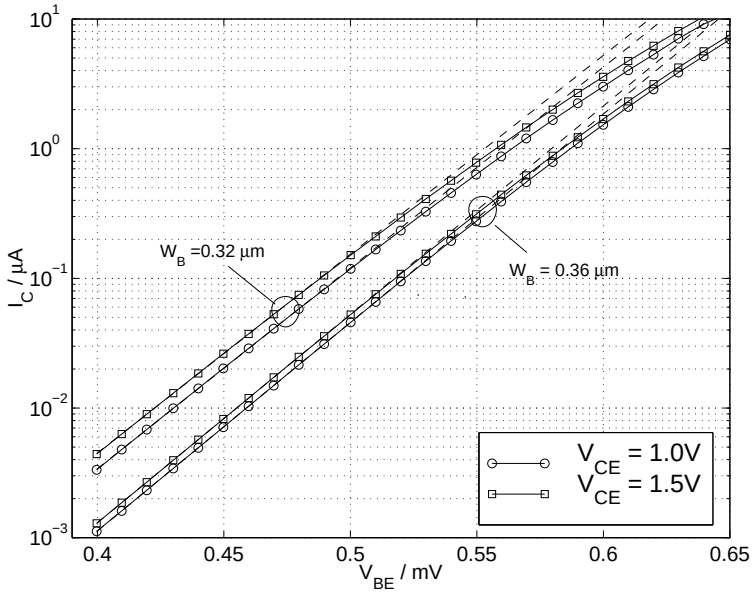

Figure 8.32: Characteristics of the lateral pnp transistor with its base and gate tied together.

tested as a dynamically biased filter, but as a tunable filter.

Fig. 8.34 shows the magnitude response of the filter with various bias currents. As before, the magnitude response of the measurement setup without the filter is subtracted from the magnitude response of the setup with the filter in place to determine the magnitude response of the filter. The dc gain is more than the expected -6 dB (Equation (7.21)) because of the Early effect in the output transistor Q_8 (Figs. 7.47(a, b)). Since the collector of the output transistor is at 0 V due to the measurement setup (Fig. 8.33), its V_{CE} is larger than that of other transistors in Figs. 7.47(a, b) by about 500 mV. Due to the low Early voltage of the transistors, this difference is enough to change the gain by about 1.3 dB. At larger values of bias currents, the proportionality of bandwidth to the bias current is not seen. This due to the transistors deviating from the exponential behavior (Fig. 8.32). With 0.5 μA bias current, the bandwidth is about 23 kHz. This is 25% less than the expected value of 31 kHz (section 7.6). Taking into account the slope factor η of 1.04, the capacitance corresponding to a 23 kHz bandwidth can be calculated to be 181 pF. This would imply an

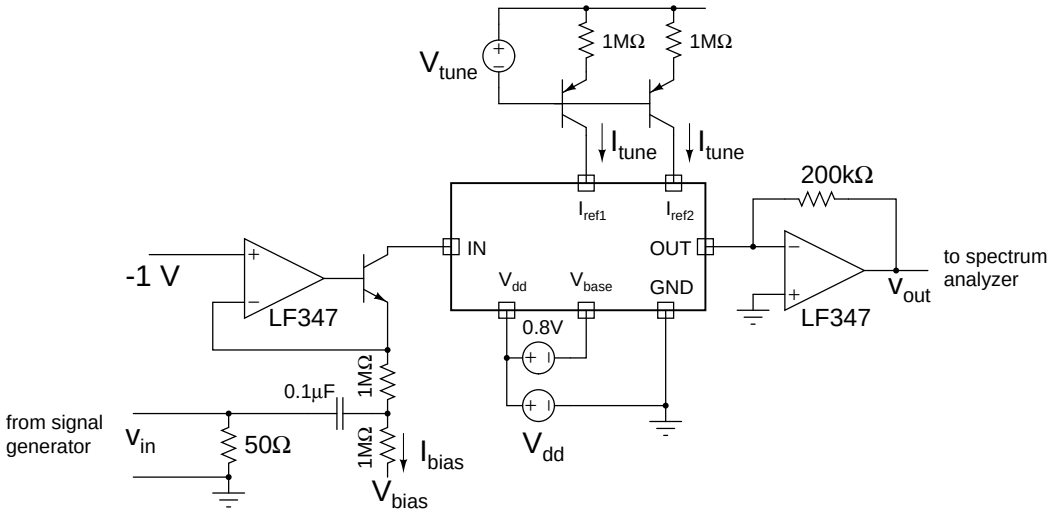

Figure 8.33: Test setup for the second-order filter using lateral bipolar transistors.

additional 41 pF across the intended 140 pF capacitors in Fig. 7.47. This seems too high to be entirely due to parasitics of the transistors connected to that node. The PMOS accumulation capacitor could itself have a larger value than expected due to process variations. Such errors in the bandwidth can be corrected by the use of an automatic tuning circuit.

Fig. 8.35 and Fig. 8.36 show the second- and third-harmonic distortion with bias currents set to 0.5 μA and 1 μA respectively. The larger second-order distortion in this single-ended filter is predominantly due to the small Early voltage. Still, in Fig. 8.35, the second harmonic is more than 35 dB below the fundamental even when the signal input is equal to the bias current. For a given modulation index[14], the distortion is significantly lower with 0.5 μA bias. This is due to increased deviation from the exponential behavior (Fig. 8.32) of the transistor at 1 μA. The distortion increases slightly with the input frequency.

The output fundamental component, noise and harmonic distortion products are plotted in Fig. 8.37 and Fig. 8.38 respectively for bias currents of 0.5 μA and 1 μA. The output noise is integrated up to 50 kHz when the bias is 0.5 μA and 100 kHz when the bias

[14]The ratio of the input signal peak to the bias current I_{bias}.

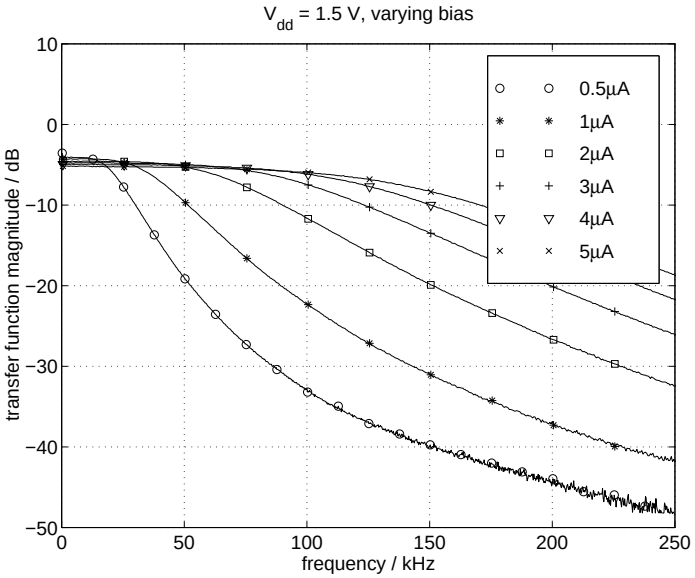

Figure 8.34: Frequency response of the second-order filter with a supply voltage of 1.5 V for various values of I_{tune} ($= I_{bias}$) shown in the insert.

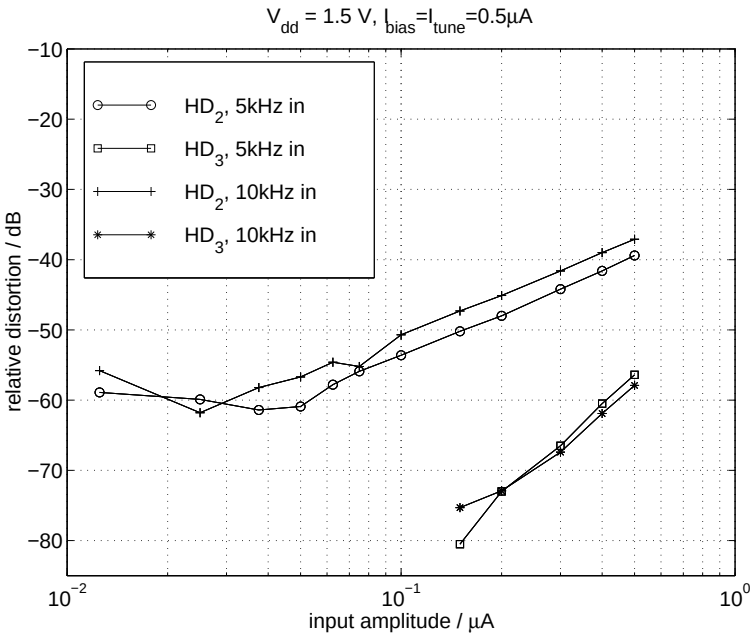

Figure 8.35: Distortion performance of the filter with $0.5\mu A$ bias current.

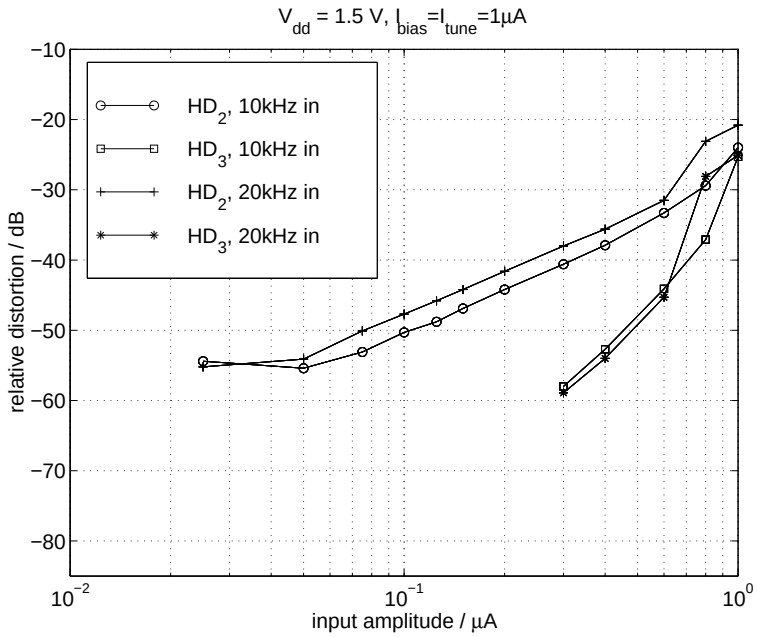

Figure 8.36: Distortion performance of the filter with $1\mu A$ bias current.

current is 1 μA. The input amplitude at which the total harmonic distortion is 40 dB and the signal to noise ratio at this input amplitude is marked on the respective figures. The values of S/N when the second harmonic or the third harmonic equals the noise are also shown. The signal to noise ratio with THD=40 dB are 56.1 dB and 47.0 dB respectively with bias currents of 0.5 μA and 1.0 μA. The maximum value of $S/N + HD_2$[15] for the same two bias currents are 44.9 dB and 40.5 dB respectively. A log-domain filter should ideally maintain the same performance when all the bias currents are changed in the same proportion. But the maximum value of $S/N + HD_2$ is smaller with 1 μA bias current than with 0.5 μA bias current. This again reflects the increased deviation from the exponential at 1 μA. Since the third harmonic is smaller than the second harmonic, performance can be improved by using pseudo differential operation to cancel the former.

As expected, the filter operates satisfactorily over the entire range of bias currents shown in Fig. 8.34 with a supply voltage of 1.5 V. The supply voltage can be increased

[15]This is a measure of the maximum signal to garbage ratio.

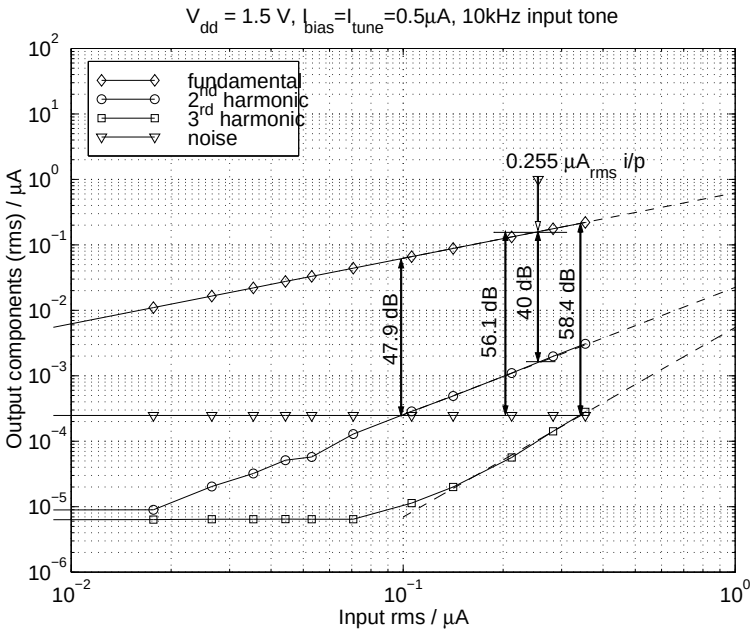

Figure 8.37: Signal, noise and distortion in the second order filter with $0.5\mu A$ bias current.

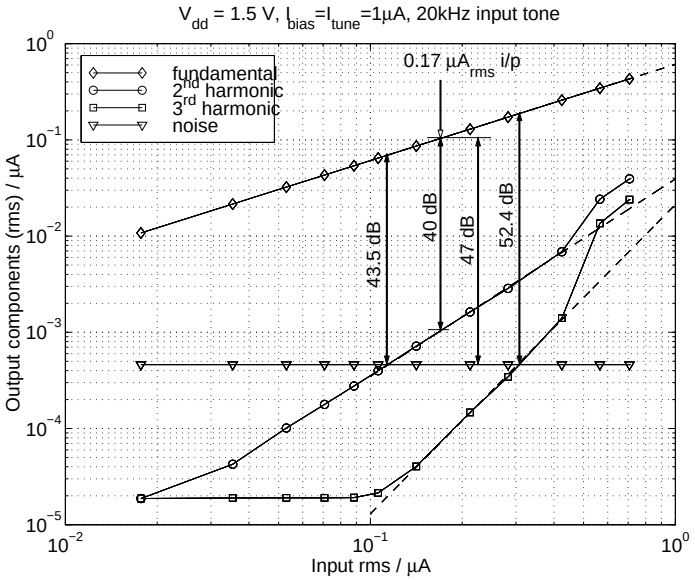

Figure 8.38: Signal, noise and distortion in the second order filter with $1\mu A$ bias current.

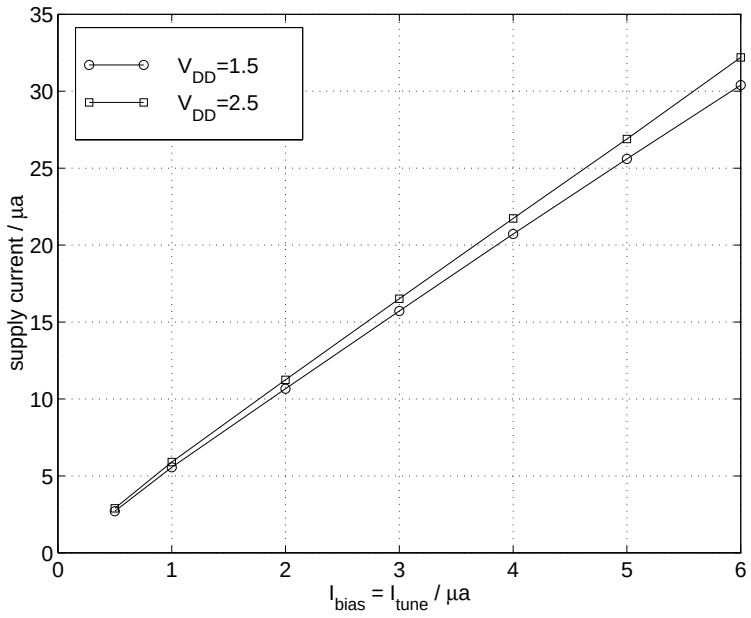

Figure 8.39: Current consumption of the second-order filter.

to 2.5 V (upper limit for this technology) without any change in performance except for a slightly increased dc gain (Due to the increased V_{CE} of the output transistor Q_8 in Fig. 7.47). The current drawn from the power supply is plotted as a function of bias current in Fig. 8.39.

Fig. 8.40 shows the chip photograph. Table 8.4 summarizes the performance of the chip.

With a bias current of 0.5 μA the power dissipation per pole, bandwidth and dynamic range (assuming 1% THD) of the fabricated filter is 0.229 fJ. This is about five times larger than the corresponding value for the single ended class-A log-domain filter using a BiCMOS technology described in [19]. It is however a respectable value considering that this filter was not "designed" in the proper sense of the term. The results do demonstrate the feasibility of log-domain filters using lateral transistors available in standard CMOS processes. Suggestions to improve the performance of filters using lateral PNP transistors

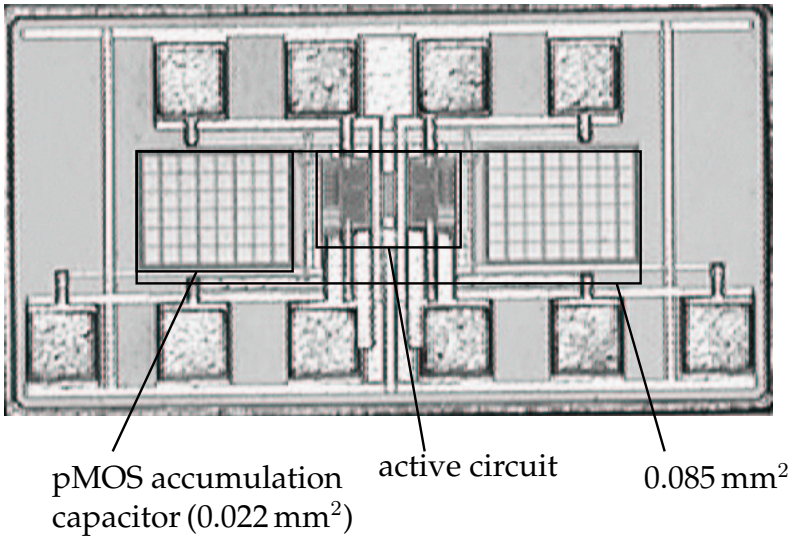

Figure 8.40: Photograph of the 2nd order filter chip using lateral PNP transistors and pMOS accumulation capacitors.

are given in section 8.4.4.

8.4 Conclusions and possible improvements

8.4.1 Third-order ladder filter

Minimizing gain variations with bias current

It was mentioned in section 7.2.9 that Early effect in the transistors in the input and the output stages (Fig. 7.7) was responsible for the variation in the dc gain with the bias current. Fig. 8.41(a) shows the feedback circuit used to establish the collector currents. The gate-source voltage of M_f and the base-emitter voltage of Q_f are denoted by V_{GSf} and V_{BEf} respectively. The collector-emitter voltage V_{CE1} of Q_1 increases with increasing bias current due to an increase in V_{GSf} and V_{BEf}. A cascode transistor Q_x can be used as shown in Fig. 8.41(b) in order to solve this problem. In this case, V_{CE1} is held constant. The increasing V_{GSf} and V_{BEf} have only a second order effect.

The disadvantages of this technique are (a) Increased supply voltage requirements and (b) Introduction of another pole in the feedback loop.

Table 8.4: Performance summary: Second order filter with lateral PNPs and MOS capacitors (Fig. 7.47).

Technology	0.25 μm CMOS	
Chip area (excl. pads)	0.085 mm^2	
Supply voltage	1.5 V	
Bias current ($I_{bias} = I_{tune}$)	$1\mu A$	$2\mu A$
-3 dB bandwidth (kHz)	22	41
Power dissipation (μW)	4.1	8.3
Output noise (rms nA)	0.25 (0 to 50 kHz)	0.46 (0 to 100 kHz)
Dynamic range (THD $<$ 40 dB)†	56.1 dB	47.0 dB
Dynamic range (max(S / N+THD))‡	44.9 dB	40.5 dB
$\frac{\text{Power dissipation}}{\text{order}\cdot\text{Bandwidth}}$	93.2 pJ	101.2 pJ
$\frac{\text{Power dissipation}}{\text{order}\cdot\text{Bandwidth}\cdot\text{DR}^\dagger}$	2.29×10^{-16} J	2.02×10^{-15} J
$\frac{\text{Power dissipation}}{\text{order}\cdot\text{Bandwidth}\cdot\text{DR}^\ddagger}$	3.01×10^{-15} J	9.02×10^{-15} J

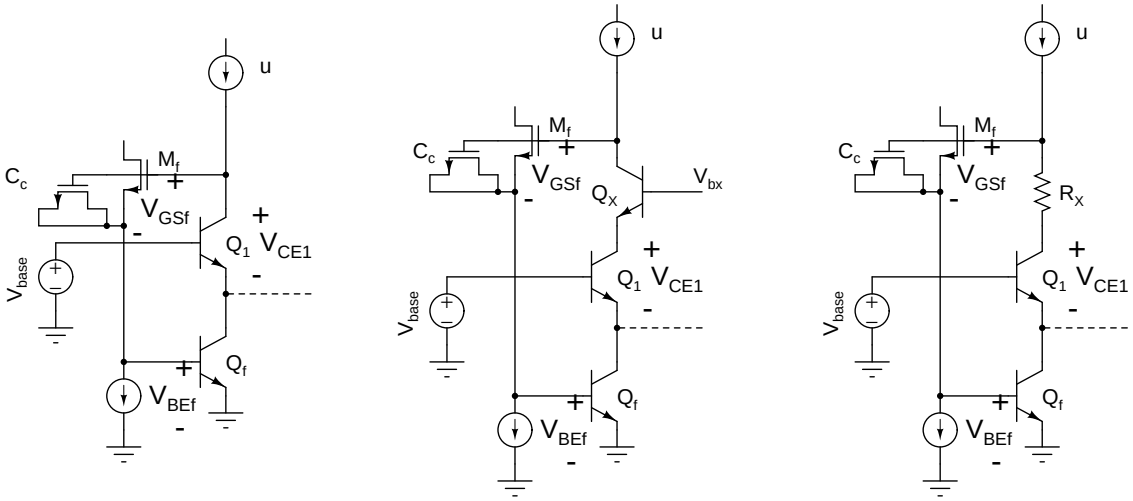

Figure 8.41: (a) Feedback used to establish the collector currents, (b) Cascoding to maintain a constant collector-emitter voltage of Q_1, (c) Using a resistor to minimize the variations in the collector-emitter voltage of Q_1.

Fig. 8.41(c) shows an alternative technique in which a resistor R_x is placed in series with the collector. The variations in V_{CE1} can be minimized by choosing the correct value of R_x. For a given R_x, as the input current increases V_{CE1} increases initially due to increasing V_{GSf} and V_{BEf}. As the current is increased further, the voltage drop across R_x counters the increase in V_{GSf} and V_{BEf}. Eventually, the increasing voltage drop across R_x dominates and V_{CE1} decreases with increasing bias current. A small R_x can be chosen such that the maximum of V_{CE1} versus the bias current occurs near the largest intended bias current. This technique provides a partial remedy to the variation of the gain with the bias current and can only be used at the input (Fig. 7.7(a)) and not at the output[16] (Fig. 7.7(b)).

In addition to these, techniques that maintain identical collector-emitter voltages of opposing transistors in a translinear loop can also be used [12].

Dynamic biasing in discrete steps

Fig. 8.42(a) shows the schematic of the log-domain filter in a general form. The core of the filter is embedded between logarithmic input and exponential output stages. The simple feedback stage shown in Fig. 7.4 is shown here to avoid cluttering the figure. But the following techniques are applicable to to other feedback circuits (e.g. Fig. 7.5, Fig. 7.6). The bias I_{bias} is varied continuously as a function of the input signal u.

Discrete variation of I_{bias} as shown in Fig. 8.42(b) can be used instead. Current sources are switched in and out in order to increase or decrease the bias current fed to the filter. The chief advantage of this technique is that the ripple that exists in I_{bias} when it is derived from one of the detectors described in Chapter 4 is avoided.

A further variation of discrete dynamic biasing is shown in Fig. 8.43. Fig. 8.43(a) shows the general form of the log-domain filter. The emitter follower Q_E in the filter core

[16] Because the current in the corresponding branch in Fig. 7.7(b) is a constant and a variable voltage drop cannot be produced by connecting a resistor.

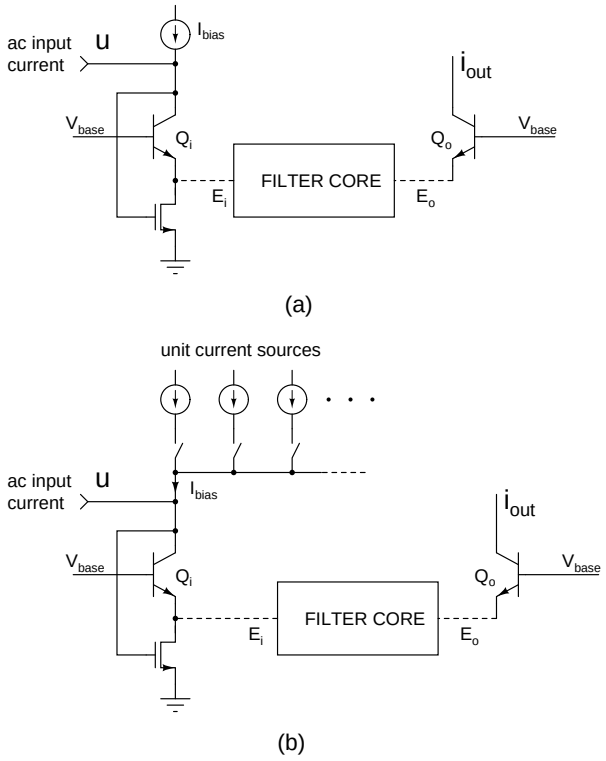

Figure 8.42: (a) Log-domain filter drawn in a general form, (b) Dynamic biasing in discrete steps.

that is used to drive the output stage is shown explicitly. The input and output cells are marked. The feedback transistor in the output stage is split into two parts M_{fE} and M_{fo}. The aspect ratios and hence the drain currents of M_{fE} and M_{fo} are in proportion to the quiescent current in Q_E (the emitter follower) and Q_o (the output transistor) respectively.

It was mentioned in section 7.2.3 that increasing the size of both the input and the output transistors in the same proportion leaves the transfer function of the filter unaffected. Thus instead of switching the bias currents at the input as shown in Fig. 8.42(b), entire input and output stages can be switched.

Fig. 8.43(b) shows an array of N input cells connected in parallel. Each of the cells is activated by turning on the appropriate switches. For example, cell-1 is active when S_1 is closed and inactive when S_{1B} is closed. Each of the cells is designed for the minimum

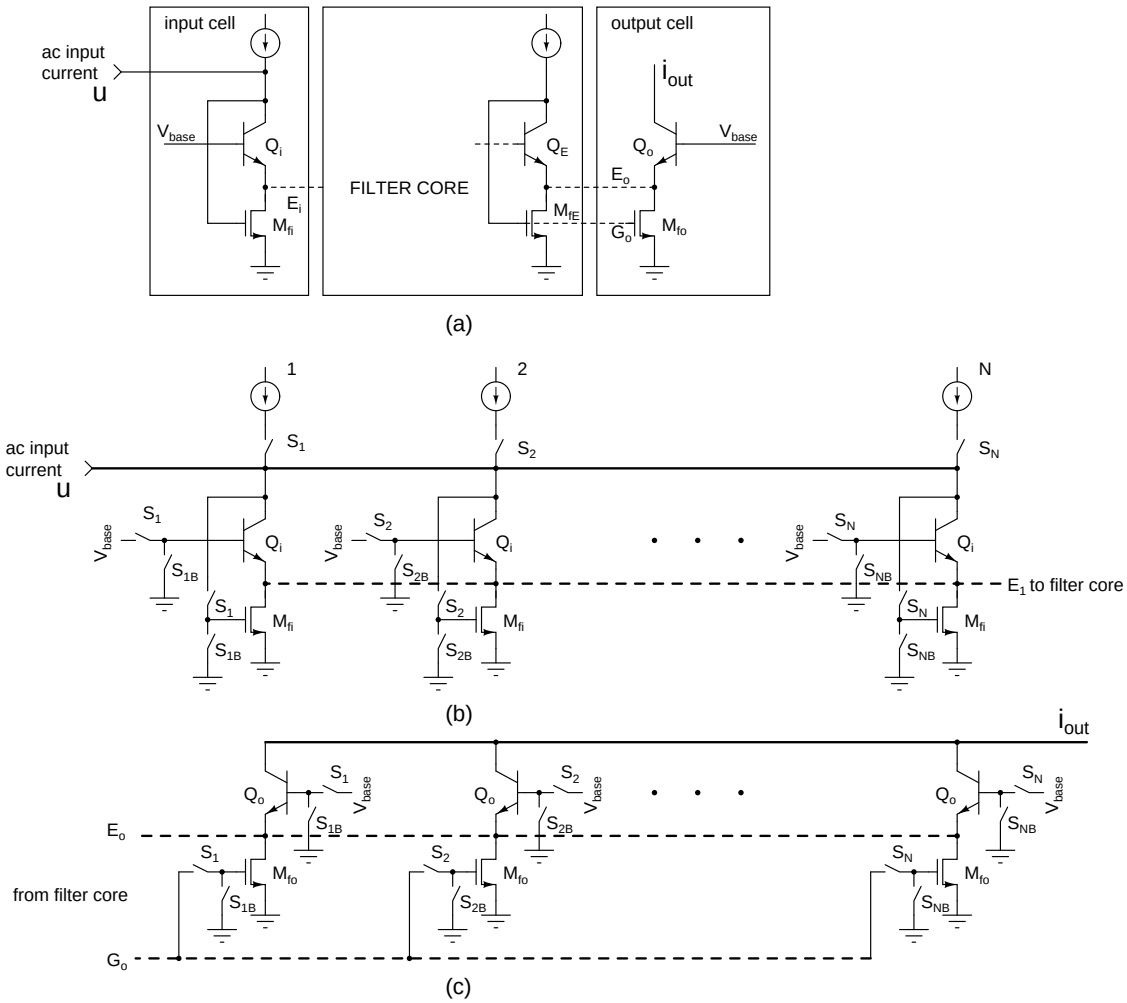

Figure 8.43: (a) Log-domain filter drawn in a general form, (b) Array of switched input cells, (c) Array of switched output cells.

required bias current. The number of cells turned on is varied according to the strength of the input signal u. Fig. 8.43(c) shows N output cells which can be turned on and off similarly. The two advantages of this scheme over using continuous variation of the input bias current are:

1. The collector voltage of the input transistor Q_1 and the final emitter follower Q_E do not vary with the bias current used at the input. Since each unit feedback transistor M_{fi} and M_{fo} used at the input and the output carries a constant current, there is no variation in their gate source voltages. Thus the variation in the dc gain of the filter

with the bias current is eliminated. Since the gate source voltage of the feedback transistor is no longer varying, the simpler feedback structure in Fig. 7.4 can be used instead of the feedback circuit with high transconductance shown in Fig. 7.6 even when the bias current is varied over a large range[17].

2. The ripple in the output of the envelope detector used to measure the signal strength does not appear at the input of the filter.

When the input bias current is intended to be varied over a large range, the scheme in Fig. 8.43 requires a large number of cells. The parasitics resulting from a large array of cells can limit the maximum usable frequency of such a scheme. In such cases, a combination of discrete and continuous variation of bias currents can be used to strike a compromise between the high frequency capabilities of the circuit and the constancy of its dc gain over a wide range. Also when a large variation in I_{bias} is intended, the number of cells that are turned on in consecutive steps can be increased; e.g. in a binary weighted fashion.

If input cells are switched and a single fixed cell is used at the output, the system in Fig. 8.43 behaves like as a combination of a filter and a variable gain amplifier that maintains a near constant output amplitude[18]. This provides a simple way of adding AGC (automatic gain control) to the filter.

Measurement of the filter

Difficulties were encountered in the measurement of the filter at small bias currents. These were mainly due to the parasitics of the printed circuit board at the input node. Fig. 8.44(a) shows the schematic of the input stage including the parasitics on the board. Typically the board parasitic C_{BOARD} ($\approx 3 - 10$ pF) is quite a bit larger than the on-chip parasitic

[17] Assuming of course that the threshold voltage of M_{fi} and M_{fo} are such that Q_1 and Q_E are operating in their active regions.

[18] The extent of variation of the output amplitude depends on the manner in which input cells are switched as a function of the input amplitude.

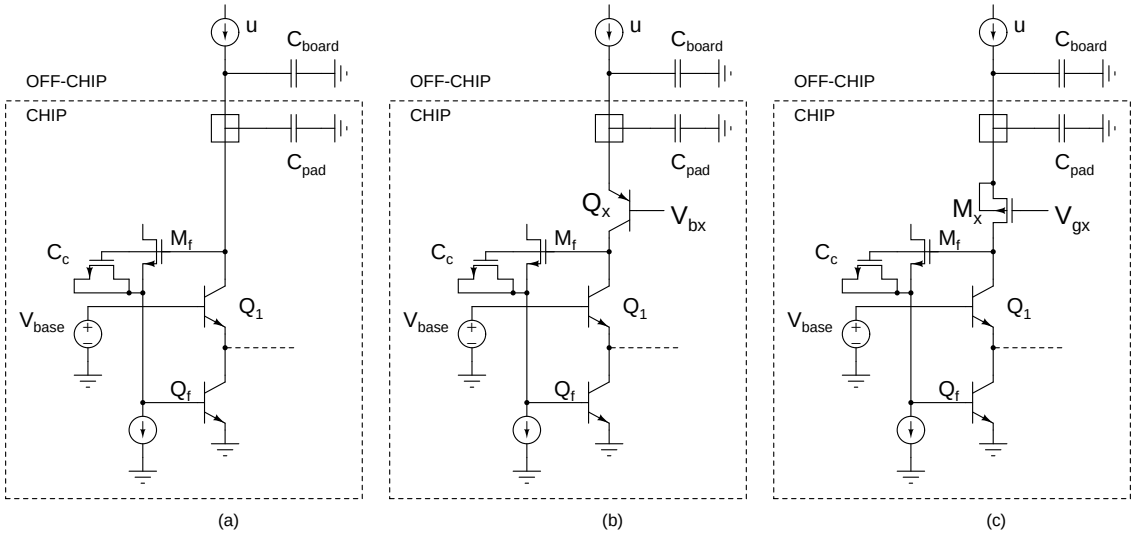

Figure 8.44: (a) Existing input stage with pad and board parasitics, (b) Using a PNP transistor at the input to isolate the board parasitics, (c) Using a pMOS transistor at the input to isolate the board parasitics.

C_{PAD} ($\approx 0.5 - 1$ pF). The effect of the parasitic C_{BOARD} (and C_{PAD}) is to cause peaking in the closed loop frequency response of the feedback loop used to establish the input current. The measured frequency response at low bias currents is thus inferior to the situation in which the filter is driven from an on-chip source (C_{BOARD} is absent).

The integrated circuit can be designed to isolate the board parasitics from the input node of the filter. If PNP devices are available in the technology, a cascode transistor Q_x can be used at the input to isolate C_{BOARD} from the input node of the filter as shown in Fig. 8.44(b). A small Q_x can be used for wide bandwidth of cascoding. Fig. 8.44(c) shows an alternative using a pMOS cascode device M_x in case suitable PNP devices are unavailable. The base (gate) bias V_{bx} (V_{gx}) must be chosen as large as possible to avoid limiting the voltage swing at the collector of Q_1 while keeping Q_x in the active region. If V_{bx} or V_{gx} is chosen to be close to V_{dd}, the voltage on the input pad swings above the power supply. It must be ensured that no junctions break down under this condition. Also, the diode that is customarily connected between the pad and V_{dd} in order to protect the circuit

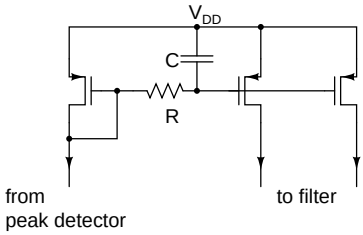

Figure 8.45: Filtering the dynamic bias current.

from electrostatic discharge must be omitted as there is a risk of it being forward biased and shunting input current away from Q_x.

In our setup C_{BOARD} was minimized using a single-sided printed circuit board and scraping off the metal around the input leads.

8.4.2 Third-order cascade filter

The improvements suggested above for the ladder filter also apply to the cascade filter.

8.4.3 Dynamically biased filter

The output of the envelope detector has a fairly large ripple at low frequencies. To avoid added distortion in the filter, the output of the envelope detector must be filtered before being used as the bias for the dynamically biased filter. A possible way of including filtering in the current mirror is shown in Fig. 8.45. C can be a nonlinear capacitor.

8.4.4 Second-order filter using lateral PNP transistors

Minimizing distortion due to Early effect

Comparison to existing log-domain filters in Bipolar/BiCMOS technologies suggests that there is room for improvement in the distortion performance of log-domain filters using lateral PNP transistors. Although no simulations of the filter in Fig. 7.47(b) were undertaken, the curves in Fig. 7.9 suggest that the awfully low Early voltage of the lateral PNP transistors (Fig. 8.32) may be responsible for the distortion in the filter. Larger base widths

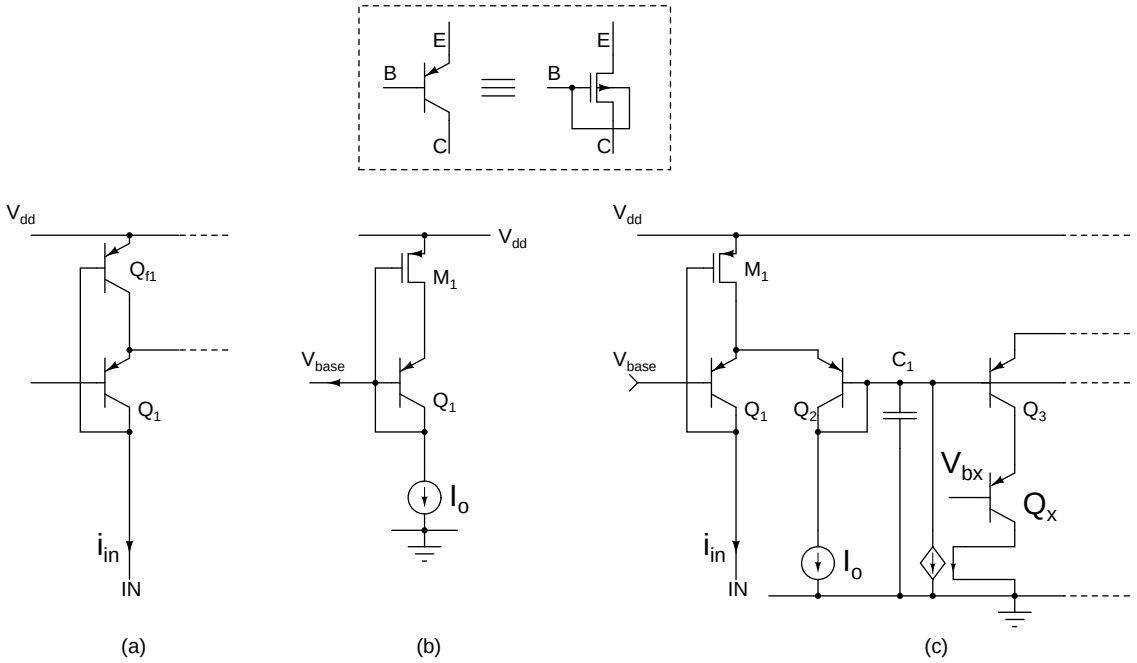

Figure 8.46: (a) Using a lateral PNP transistor instead of a pMOS transistor to increase the g_m of the feedback path, (b) Generating V_{base}, (c) Using a cascode transistor equalize collector voltages of all the transistors.

can be used to increase the Early voltage. A transistor with a wider base has a lower current gain β, but some degradation of β is acceptable since the current gain of the enhanced lateral bipolar transistor is large enough[54, 55]. As suggested earlier, cascoding of transistors and increasing the transconductance of the feedback loop used to establish the collector currents can be used as well. The latter is shown in Fig. 8.46(a). If Q_{f1} is an enhanced lateral bipolar with a current gain of 1000, the distortion introduced by its base current can be insignificant especially if dynamic biasing is not intended to be used at the input and the output. The feedback transconductance is higher than if a pMOS were used in place of Q_{f1}.

Sensitivity to variations in the power supply

For the circuit in Fig. 7.47(b) to be insensitive to the power supply, V_{base} should be a constant with respect to V_{dd}. If V_{base} is a constant with respect to ground, a variation in V_{dd}

produces corresponding variations in the collector emitter voltages of $Q_{1,4,7}$. These variations act as spurious inputs to the filter. To remedy this situation, V_{base} should be held constant with respect to V_{dd} in Fig. 7.47(b). Fig. 8.46(b) shows a technique for generating a constant voltage with respect to V_{dd}. If the current densities in Q_1 and M_1 are identical to the corresponding bipolar and pMOS transistors used in the filter (i.e. Q_1 & M_1, Q_7 & M_7, Q_7 & M_7 in Fig. 7.47(b)), the quiescent V_{CB} s of the transistors in the filter will be zero. The current density in Q_1 and M_1 can be altered to obtain the desired quiescent V_{CB} s for the transistors in the filter. With this arrangement, a change in V_{dd} causes an equal change in the voltages at all the nodes except the collectors of Q_3 and Q_8 (Fig. 7.47(b)). The collector voltage of Q_3 in Fig. 7.47(b) is defined with respect to ground (by the diode connected transistor of the current mirror). The arrangement in Fig. 8.46(c) can be used to bootstrap the collector voltage of Q_3 to V_{dd}. Q_3 is cascoded by Q_x whose base voltage V_{bx} is referenced to V_{dd} using a technique similar the one in Fig. 8.46(b). Analogous techniques can be used with Q_8. The resulting circuit is insensitive to variations in V_{dd}.

Chapter 9

Conclusions and Suggestions for Future Work

9.1 Conclusions

The use of companding for improving the dynamic range of log-domain filters for a given power consumption was investigated in this thesis. The limitations of traditional linear active filters and existing methods of companding used to overcome these were discussed in Chapter 2.

A simplified technique for dynamic biasing in log-domain filters was given in Chapter 3. It was shown that owing to its simplicity, the proposed technique enjoys various advantages over the presently known methods. The issues involved in the generation of a dynamic biasing signal based on the input signal strength were discussed in Chapter 4.

Based on the ideas presented in Chapter 3, the design of pseudo-differential filters that maintain external linearity in presence of dynamic biasing was presented in Chapter 7. The measured results presented in Chapter 8 confirm the ability of these filters to handle large dynamic range input signals. The design and the measured performance of a peak detector that can be used for dynamic biasing was also presented in these chapters.

Possible techniques for the realization of log-domain filters in using a standard CMOS process and experimental results from a prototype second-order filter were presented in

Chapters 5 and 8. These filters are attractive for use at low frequencies in view of CMOS being the dominant technology of the day.

9.2 Future work

Companding filters are relatively new entrants into the arena of analog filters. The following are few areas that need further investigation.

Companding in presence of blockers

The degradation of signal to noise ratio in presence of large signals in the stopband is inherent to companding filters. This point was alluded to in section 8.1.1. The problem of processing a small signal with a sufficient signal to noise ratio in presence of a large signal needs to be investigated. In presence of a large signal in the stopband, the input amplifier (Fig. 1.1) is constrained to have a small gain[3]. This situation can perhaps be remedied by applying companding to the individual stages instead of the filter as a whole. Since the blocker in the stopband is progressively attenuated by the filter, the amount of compression applied can be reduced. This ensures that the desired small signal is not buried in noise.

Applying companding filters in practical situations

Companding filters have to be tailored for use in practical situations. Companding is most beneficial in applications that require modest signal to noise ratios to be maintained over a large range of input signals. Examples of this are hearing aids. The use of log-domain filters in the former is presented in [65, 66]. A translinear gain control circuit with quiescent control for hearing aids is presented in [67]. The dynamic biasing techniques presented in this dissertation could perhaps be used to combine the gain control and the filtering functions in a hearing aid.

Another place where large dynamic range signals are encountered and the signal to noise ratio demanded is modest is in radio receivers used in digital communication systems. However, such systems also encounter large out of band blocking signals. Companding filters in the present form offer no advantages over traditional active filters in such situations. The problem of degradation of signal to noise ratio in presence of blocking signals has to be tackled before companding filters can be used in radio receivers.

Bibliography

[1] E. Vittoz, "Low power low-voltage limitations and prospects in analog design", in R. J. v. d. Plassche, W. M. C. Sansen, and J. H. Huijsing, eds., *Analog Circuit Design, Low-Power, Low-Voltage, Integrated Filters and Smart-Power*, Kluwer, Boston, 1995.

[2] Y. P. Tsividis, V. Gopinathan, and L. Tóth, "Companding in signal processing", *Electronics Letters*, vol. 26, pp. 1331-1332, Aug. 1990.

[3] Y. Tsividis, "Externally linear time-invariant systems and their applications to companding signal processors", *IEEE TCAS-II*, vol. 44, no. 2, pp. 65-85, Feb. 1997.

[4] S. Haykin, *Communication Systems*, 3rd edition, John Wiley, 1994.

[5] E. Blumenkrantz, "Analog floating point technique", *Proc. 1995 IEEE Symposium on Low Power Electronics*, San Jose, CA, Oct. 1995, pp. 72-73.

[6] E. Seevinck, "Companding current mode integrator: a new circuit principle for continuous time monolithic filters", *Electronics Letters*, vol. 26, no. 24, pp. 2046-2047, Nov. 1990.

[7] D. R. Frey, "Log-domain filtering: an approach to current mode filtering", *IEE Proc. G 1993*, vol. 140, no. 6, pp. 406-416, Dec. 1993.

[8] Y. Tsividis, "General approach to signal processors employing companding", *Electronics Letters*, vol. 31, no. 18, pp. 1549-1550, Aug. 1995.

[9] Y. Tsividis and D. Li, "Current-mode filters using syllabic companding", *Proc. IEEE ISCAS*, pp. 121-124, May 1996.

[10] R. W. Adams, "Filtering in the log domain," Preprint 1470, presented at 63rd Audio Eng. Soc. Conf., NY, May, 1997.

[11] B. Gilbert, "Translinear circuits: a proposed classification", *Electronics Letters*, vol. 11, no. 1, pp. 14-16, Jan. 1975.

[12] E. Seevinck, *Analysis and Synthesis of Translinear Circuits*, Elsevier, Amsterdam, 1988.

[13] E. M. Drakakis, A. J. Payne, and C. Tomazou, "Log-domain filters, translinear circuits and the Bernoulli cell", *Proc. IEEE ISCAS*, vol. 1, pp. 501-504, Hong Kong, 1997.

[14] D. Perry and G. W. Roberts, "Log-domain filters based on LC ladder synthesis", *Proc. 1995 IEEE ISCAS*, vol. 1, pp. 311-314, 1995.

[15] J. Mulder, A. C. van der Woerd, W. A. Serdijn, and A. H. M. van Roermund, "General current mode analysis of translinear filters", *IEEE TCAS-I*, vol. 44, no. 3, pp. 193-197, Mar. 1997.

[16] D. Frey and A. Tola, "A state space formulation for externally linear class AB dynamical circuits", *IEEE TCAS-II*, vol. 46, no. 3, March 1999.

[17] J. Mulder, W. A. Serdijn, A. C. van der Woerd, and A. H. M. van Roermund, "A syllabic companding translinear filter", *Proc. 1997 ISCAS*, vol. 1, pp. 101-104, Hong Kong, 1997.

[18] D. R. Frey and Y. P. Tsividis, "Syllabically companding log domain filter using dynamic biasing", *Electronics Letters*, vol. 33, no. 18, pp. 1506-1507, 28 Aug. 1997.

[19] M. Punzenberger and C. Enz, "A 1.2-V low power BiCMOS class AB log-domain filter", *IEEE Journal of Solid State Circuits*, vol. 32, pp. 1968-1978, Dec. 1997.

[20] M. Punzenberger and C. Enz, "A compact low-power BiCMOS log-domain filter", *IEEE Journal of Solid State Circuits*, vol. 33, pp. 1123-1129, Jul. 1998.

[21] D. R. Frey, "Log-domain filtering for RF applications", *IEEE Journal of Solid State Circuits*, vol. 31, no. 10, pp. 1468-1475, Oct. 1996.

[22] M. El-Gamal, R. A. Baki, and A. Bar-Dor, "30-100 MHz npn-only variable gain class AB companding based filters for 1.2 V applications", *Proceedings 2000 ISSCC*, vol. 43, pp. 150-151, San Fransisco, Feb. 2000.

[23] G. van Ruymbeke, C. C. Enz, F. Krummenacher, and M. Declercq, "A BiCMOS programmable continuous-time filter using image-parameter synthesis and voltage-companding technique", *IEEE JSSC*, vol. 33, no. 3, Mar. 1997.

[24] Y. Tsividis, "Minimising power dissipation in analogue signal processors through syllabic companding", *Electronics Letters*, vol. 35, no. 31, pp. 1805-1807, 14. Oct. 1999.

[25] E. A. Vittoz and Y. Tsividis, "Frequency-Dynamic Range-Power", *to be published*, 2000.

[26] Y. Tsividis, "On linear integrators and differentiators using instantaneous companding", *IEEE TCAS-II*, vol. 42, pp. 561-564, Aug. 1995.

[27] D. R. Frey, "Exponential state space filters: A generic current mode design strategy," *IEEE Trans. Circuits Syst. I*, vol. 43, pp. 34-42, Jan. 1996.

[28] B. P. Lathi, *Linear Systems and Signals*, Berkeley-Cambridge Press, Carmichael, California, 1992.

[29] R. M. Fox, "Design oriented analysis of log-domain filters", *IEEE TCAS-II*, vol. 45, no. 7, Jul. 1998.

[30] D. Frey, Y. Tsividis, G. Efthivoulidis, and N. Krishnapura, "Syllabic companding log domain filters", *to be published in IEEE TCAS-II*, 2000.

[31] R. Fox, M. Nagarajan, and J. Harris, "Practical design of single ended log-domain filter circuits", *Proc. 1997 ISCAS*, Hong Kong, 1997.

[32] J. Mulder, M. H. L. Kouwenhoven, W. A. Serdijn, A. C. van der Woerd and A. H. M. van Roermund, "Analysis of noise in translinear filters", *Proc. 1998 ISCAS*, vol. 1, pp. 337-340, Monterey, 1998.

[33] L. Tóth, Y. Tsividis, N. Krishnapura, "Analysis of noise and interference in companding signal processors", *Proc. 1998 ISCAS*, vol. 1, pp. 143-146, Jun 1-3 1998, Monterey, California.

[34] D. Frey, "Synthesis of distortion compensated log-domain filters using state space techniques", *Proc. 1998 ISCAS*, vol. 1, pp. 321-324, Monterey, 1998.

[35] Y. P. Tsividis and J. O. Voorman (editors), *Integrated Continuous-Time Filters—Principles, Design and Applications*, IEEE Press, Piscataway, 1992.

[36] M. Punzenberger and C. Enz, "Noise in instantaneous companding filters", *Proc. 1997 ISCAS*, vol. 2, pp. 337-340, Hong Kong, 1997.

[37] N. Krishnapura, Y. Tsividis, and D. R. Frey, "Simplified technique for syllabic companding in log-domain filters", *Electronics Letters*, vol. 36, no. 15, pp. 1257-1259, 20 Jul. 2000.

[38] D. R. Frey, "On instantaneous vs. syllabic companding in log domain filters", *Proc. 1999 ISCAS*, vol. 2, pp. 672-676, Orlando, 1999.

[39] R. M. Fox, "Enhancing dynamic range in differential log-domain filters based on the two-filters approach", *Proc. 2000 ISCAS*, vol. 2, pp. 617-620, Geneva, 2000.

[40] C. C. Todd, "A Monolithic Analog Compandor", *IEEE Journal of Solid State Circuits*, vol. 11, no. 6, Dec. 1976.

[41] S. Armstrong, "The dynamics of compression: Some key elements explored", *The Hearing Journal*, vol. 46, no. 11, pp. 1-4, Nov. 1993.

[42] D. Sheingold, *Nonlinear Circuits Handbook*, Analog Devices Inc., Norwood, MA, 1996.

[43] D. R. Frey, "Explicit log-domain square root detector", United States Patent no. 5585757, Dec. 1996.

[44] J. Mulder, A. C. van der Woerd, W. A. Serdijn, and A. H. M. van Roermund, "An RMS-DC converter based on the dynamic translinear principle", *IEEE Journal of Solid State Circuits*, vol. 32, no. 7, Jul. 1997.

[45] Y. Tsividis, *Operation and Modeling of the MOS Transistor*, Second edition, WCB/McGraw-Hill, 1999.

[46] C. Toumazou, J. Ngarmnil, and T. S. Lande, "Micropower log-domain filter for electronic cochlea", *Electronics Letters*, vol. 30, pp. 1839-1841, 27 Oct. 1994.

[47] C. C. Enz and M. Punzenberger, "1-V log-domain filters", *Proc. Advances in Analog Circuit Design 1998*, Apr. 1998.

[48] Dominique Python, Manfred Punzenberger and Christian Enz, "A 1-V CMOS log-domain integrator", Proc. 1999 ISCAS, vol. 2, pp. 685-688, 1999.

[49] J. Mulder, *Static and Dynamic Translinear Circuits*, PhD dissertation, Delft University Press, 1998.

[50] Y. P. Tsividis and R. W. Ulmer, "A CMOS voltage reference", *IEEE Journal of Solid State Circuits*, vol. 13, no. 6, pp. 774-778, Dec. 1978.

[51] E. A. Vittoz, "MOS transistors operated in the lateral bipolar mode and their applications to bipolar technology", *IEEE Journal of Solid State Circuits*, vol. 18, no. 3, pp. 273-279, Jun. 1983.

[52] T. Pan and A. Abidi, "A 50-dB variable gain amplifier using parasitic bipolar transistors in CMOS", *IEEE Journal of Solid State Circuits*, vol. 24, no. 4, pp. 951-961, Jul. 1989.

[53] X. Arreguit, *Compatible Lateral Bipolar Transistors In CMOS Technology: Model And Applications*, PhD thesis, EPFL, Lausanne 1989.

[54] S. Verdonckt-Vandebroek, S. S. Wong, J. C. S. Woo, and P. K. Ko, "High-gain lateral bipolar action in a MOSFET structure", *IEEE Transactions on Electron Devices*, vol. 38, no. 11, Nov. 1991.

[55] S. Verdonckt-Vandebroek, J. You, J. C. S. Woo, and S. S. Wong, "High-gain lateral p-n-p bipolar action in a p-MOSFET structure", *IEEE Electron Device Letters*, vol. 13, no. 6, Nov. 1992.

[56] A. Annema, "Low-power bandgap references featuring DTMOST's", *IEEE Journal of Solid State Circuits*, vol. 34, no. 7, Jul. 1999.

[57] S. Pavan and Y. Tsividis, *High Frequency Continuous Time Filters in Digital CMOS Processes*, Kluwer, 1999.

[58] S. O. Rice, "Response of periodically varying systems to shot noise–Application to switched RC circuits", *Bell System Technical Journal*, vol. 49, no. 9, pp. 2221–2247, Nov. 1970.

[59] T. Ström and S. Signell, "Analysis of periodically switched linear circuits", *IEEE TCAS* vol. CAS-24, pp. 531-541, Oct 1977.

[60] A. Papoulis, *Probability, Random Variables and Stochastic Processes*, McGraw Hill, 1965.

[61] E. Kreyszig, *Advanced Engineering Mathematics*, John Wiley, 1998.

[62] A. V. Oppenheim and R. W. Schafer, *Discrete-Time Signal Processing*, Prentice Hall, 1989.

[63] A. Hajimiri and T. H. Lee, "A general theory of phase noise in electrical oscillators", *IEEE Journal of Solid State Circuits*, vol. 33, no. 16, pp. 1547-1548, Aug. 1998.

[64] C. D. Hull and R. G. Meyer, "A systematic approach to the analysis of noise in mixers", *IEEE Transactions on Circuits and Systems I*, vol.40, no. 12, pp. 909-19, Dec. 1993.

[65] W. A. Serdijn, M. Broest, J. Mulder, A. C. van der Woerd, A. H. M. van Roermund, "A low-voltage ultra-low power translinear integrator for audio filter applications", *IEEE Journal of Solid State Circuits*, vol. 32, no. 4, Apr. 1997.

[66] L. P. L. van Dijk, A. C. van der Woerd, J. Mulder, and A. H. M. van Roermund, "An ultra-low-power, low-voltage electronic audio delay line for use in hearing aids", *IEEE Journal of Solid State Circuits*, vol. 33, no. 2, Feb. 1998.

[67] A. C. van der Woerd, W. A. Serdijn, J. Davidse, and A. H. M. van Roermund, "Fully integrable class-AB rear-end with smart quiescent current control for a general purpose hearing aid chip", *Proc. 1996 European Solid State Circuits Conference*, pp. 162-165, 1996.

[68] F. Yang and C. C. Enz, "A low distortion BiCMOS seventh-order Bessel filter operating at 2.5 V supply", *IEEE Journal of Solid State Circuits*, vol. 31, pp. 321-330, Mar. 1996.

[69] D. Python, A. S. Porret and C. Enz, "A 1 V 5^{th}-order Bessel filter dedicated to digital standard processes", *Proc. 1999 CICC*, pp. 505-508, 1999.

[70] D. G. Python and C. C. Enz, "A 40 μW, 75 dB dynamic range, 70 kHz bandwidth biquad filter based on complementary MOS transconductors", *Proc. 1999 European Solid State Circuits Conference*, pp. 38-41, Duisburg, Germany, 1999.

[71] R. H. Zele and D. J. Allstot, "Low-power CMOS continuous-time filters", *IEEE Journal of Solid State Circuits*, pp. 157-168, vol 31, no. 2, Feb. 1996.

[72] Y. P. Tsividis, "Minimal Transistor-only micropower integrated VHF active filter", *Electronics Letters*, vol. 23, no. 15, pp. 777-778, Jul. 1987. 1.1, 1.1, 2.1.1, 4, 8.1.1, A 1.1, 1.2, 2 1.1, 1.1, 1.2, 2.1.3, 2.2.1, 2.2.2, 5, 3.3.1, 6.1, 6.3.3, 9.2 1.1, 6.2.3 1.1, 1.2, 2.1.3, 5 (document), 1.1, 1.2, 2.2.1, 2.2.1, 2.2.1, 3.2, 3.3.2, 5.1, 7.1 1.1, 1.2, 2.2.1, 2.2.1, 1, 1 1.2, 2.1.3, 2.2.2, 5 1.2, 5, 5 1.2 1.2, 2.2.1, 2.2.1 1.2, 2.2.1, 2.2.1, 52 1.2 1.2, 2.2.1, 2.2.1, 1 1.2 1.2, 2.2.1 1.2, 5 1.2, 1.3, 7, 3.2, 1, 3.3.1 (document), 1.2, 2.2.1, 2.2.1, 2.2.1, 2.2.1, 2.2.1, 3.2, 3.3.2, 2, 5.2.2, 7.1, 7.2, 7.2.1, 7.4, 7.2.2, 7.6, 8.1, 51, A.2 1.2, 2.2.1, 2, 7.2, 7.2.1, 8.1, A.2 1.2 (document), 1.2, 2.2.1, 7.5, 29, 8.1, A.2 1.2 1.2, 1.3 2.1.1, A 2.2.1, 2.2.1, 6.2.1 2.2.1, 2.2.1, 2.2.1, 1 2.2.1 2.2.1 2.2.1, 7, 2, 3.3.1, 4.2.2, 33 2.2.1 2.2.1, 2.2.3, 6.1, 6.3.3 2.2.1, 2.2.3, 6.1, 6.2.1, 6.2.4, 6.3.3 2.2.1 2.2.1, 2 2.2.3, 6.1 3.1 3.2 3.2, 3.3.2 2 4.2.2, 18 4.3.2 4.3.2, 15 4.3.2,

15 5.1, 29, 30, 31, 4 5.1 5.1, 21, 5.2.3 5.1 5.1 22 22 22, 5 22 5.2.2, 2, 7.6, 8.4.4 5.2.2, 8.4.4 5.2.2 5.2.3, 7.3.9 6.2.2, 6.2.2 6.2.2, 6.2.2 6.2.3 1 6.3.1 6.3.1 6.3.2 9.2 9.2 9.2 9, 8.1, 2, A.2 8.1, A.2 8.1, A.2 8.1, A.2 8.1, A.2

Appendix A

Power Dissipation for a Given Signal to Noise Ratio

In this appendix, the analytical expressions [1, 25] for the power consumed in simple first-order circuits for a given signal to noise ratio (S/N) are derived. The conformity of higher order filters to these expressions is examined through simulation.

A.1 First-order RC filter

Fig. A.1(a) shows a first-order filter using a resistor R and a capacitor C. The -3 dB bandwidth of the filter is given by $\omega_p = 1/RC$. The transfer function of the filter is given by

$$\frac{V_c(s)}{V_i(s)} = \frac{1}{1 + s/\omega_p} \tag{A.1}$$

Thermal noise from the resistor results in noise at the filter's output. The integrated noise at the output is given by the well known expression

$$\overline{v_{on}^2} = \frac{kT}{C} \tag{A.2}$$

where k denotes the Boltzmann constant and T the absolute temperature.

The power drawn from the input source and the output S/N depends on the applied input signal. A sinusoid at the -3 dB frequency of the filter is chosen to be the representative input. A voltage $v_i(t) = V_p \cos(\omega_p t)$ is applied to the filter's input. The output voltage v_c is

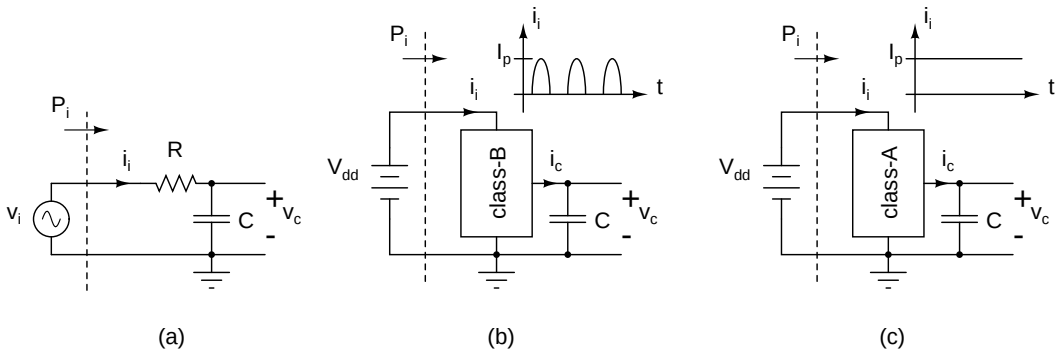

Figure A.1: (a) First-order RC filter, (b) Class-AB circuit driving a capacitor, (c) Class-A circuit driving a capacitor.

given by

$$v_c(t) = \frac{V_p}{\sqrt{2}} \cos(\omega_p t - 45^o) \tag{A.3}$$

The output S/N is the ratio of the mean-square value of signal v_c to that of noise.

$$S/N = \frac{V_p^2 C}{4kT} \tag{A.4}$$

The current drawn by the circuit with the given input v_i is

$$i_i(t) = \frac{V_p}{\sqrt{2}R} \cos(\omega_p t + 45^o) \tag{A.5}$$

The power drawn from the input source is given by

$$P_i = \int_{-T_p/2}^{T_p/2} v_i(t) i_i(t) dt \tag{A.6}$$

$$= \frac{V_p^2}{4R} \tag{A.7}$$

where $T_p = 2\pi/\omega_p$ is the period of the sinusoidal input.

V_p can be eliminated from (A.7) and (A.4) to relate the power dissipated in the circuit P_i to the output signal to noise ratio.

$$P_i = \frac{kT\, S/N}{RC} \tag{A.8}$$

$$= kT\omega_p\, S/N \tag{A.9}$$

$$= 2\pi kT f_p\, S/N \tag{A.10}$$

where f_p is the filter's bandwidth in Hz. The power dissipation is proportional on the bandwidth ω_p and the signal to noise ratio S/N.

A.2 Capacitor driven from an ideal class-B driver

Fig. A.1(b) shows an ideal class-B driver driving a capacitor C. The positive output currents in an ideal class-B driver are drawn from the positive supply (V_{dd}) and negative output currents from ground. The mean squared integrated noise voltage on the capacitor is assumed to be kT/C, the same as if a resistor was shunting the capacitor. The output S/N can be calculated assuming a sinusoidal voltage $v_c(t) = V_p \cos(\omega_0 t)$ across the capacitor. The peak voltage V_p can at most be $V_{dd}/2$ since the output voltage is limited by the supply rails. With this peak swing, the output S/N is given by

$$S/N = \frac{V_{dd}^2 C}{8kT} \tag{A.11}$$

The current through the capacitor when the voltage across it is $v_c(t) = (V_{dd}/2)\cos(\omega_0 t)$ is

$$i_c(t) = -\frac{\omega C V_{dd}}{2} \sin(\omega_0 t) \tag{A.12}$$

Thus, the current i_i drawn from the positive power supply is a train of half sinusoidal pulses of amplitude $\omega_0 C V_{dd}/2$. The power drawn from this supply is[1]

$$P_i = V_{dd} \int_{-T_p/2}^{T_p/2} i_i(t)dt \tag{A.13}$$

$$= \frac{\omega_0 C V_{dd}^2}{2\pi} \tag{A.14}$$

Eliminating V_{dd} from (A.11) and (A.14), the power dissipation can be related to S/N as

$$P_i = \frac{8kT\omega_0\, S/N}{2\pi} \tag{A.15}$$

$$= 8kT f_0\, S/N \tag{A.16}$$

[1]The average value of a half wave rectified sinusoid of unit amplitude over one period is $1/\pi$.

where f_0 is the frequency of the capacitor voltage in Hz.

The proportionality constants in (A.10) and (A.16) are different, but not by a great amount. In both cases, the power dissipation is proportional to S/N and frequency.

A.3 Capacitor driven from an ideal class-A driver

Most filters use class-A output stages to drive capacitors. A constant bias current I_0 then flows through the output branch. The peak value of the current that can be driven into the capacitor is equal to this bias current. Fig. A.1(c) shows a capacitor being driven from a class-A driver. If the voltage across the capacitor has the same amplitude and frequency as considered in the previous section, the minimum permissible value for this bias current is $\omega_0 C V_{dd}/2$. The power drawn from the power supply in this case is

$$P_i = V_{dd} I_0 \tag{A.17}$$

$$= \frac{\omega_0 C V_{dd}^2}{2} \tag{A.18}$$

$$= 8\pi k T f_0 \, S/N \tag{A.19}$$

This is π times larger than the power drawn from a class-AB driver. This is consistent with the well known result that the maximum power efficiencies of class-A and class-B amplifiers (when no inductors are used) are $1/4$ and $\pi/4$ respectively.

A.4 Summary

(A.10), (A.16), and (A.19) are derived using some simplifying assumptions for circuits with a single capacitor. For higher order filters, it is assumed that equal power is dissipated in each pole-forming circuit. The power dissipated in a first-order filter is multiplied by the filter's order n to estimate the overall power dissipation. It is not proven that the power dissipation in the expressions above are the minimum required for a given S/N,

Table A.1: Power dissipation per S/N, signal frequency and filter order ($T = 300$ K).

filter	$P_i/(S/N\ f_p\ n)$ (J)	
	analytical expression	calculated/ simulated value
First-order RC filter	$2\pi kT$	2.6×10^{-20}
Capacitor with class-B driver	$8kT$	3.3×10^{-20}
Capacitor with class-A driver	$8\pi kT$	10.4×10^{-20}
Second-order Butterworth doubly terminated RLC	—	3.62×10^{-20}
Third-order Butterworth doubly terminated RLC	—	5.12×10^{-20}
Third-order Butterworth singly terminated RLC	—	1.69×10^{-20}

although it is certainly hard to imagine a circuit that would consume less power than a passive RC filter. Nevertheless they are often used as the fundamental lower limits for the power dissipation in filters for a given S/N. The power dissipation for a given S/N, signal frequency and filter order for several passive filters, including the ones considered above are listed in Table A.1. The numbers for higher order filters are arrived at using SPICE simulations. In each case, the input to the filter is a sinusoid at the -3 dB frequency. It can be seen that the resulting numbers are quite close to that derived analytically for the first-order RC filter.

Table A.2 lists the power dissipation, the bandwidth, the order, and the dynamic range of several previously published active filters. The list is arranged in the ascending order of the normalized power dissipation (normalized to the order, the bandwidth, and the dynamic range).

The dynamic range specified is the range of input signals over which THD ≤ 40 dB[2] and $S/N > 0$ dB are maintained. Of the listed filters 3, 5, and 8 are companding (log-domain class-AB) filters and the rest are conventional linear filters. For the latter, the

[2]Except for 7 [68] which uses the criterion THD ≤ 49 dB.

Table A.2: Power dissipation per dynamic range DR, -3 dB bandwidth, and filter order for published active filters.

	V_{dd} (V)	Power	f_{-3dB}	ord.	DR (dB)	power per pole, f_{-3dB} ($\times 10^{-12}$ J)	power per pole, f_{-3dB}, DR ($\times 10^{-18}$ J)	Ref.
1	1.0	10.5 μW	100 kHz	5	68 (max.) 57 (min.)	21.0 21.0	3.3 41.9	[69]
2	2.5	40 μW	70 kHz	2	75	285.7	9.03	[70]
3	1.2	65 μW[q] 170 μW[m]	320 kHz	3	65	67.7[q] 177.1[m]	21.4[q] 56.0[m]	[19]
4	1.5	375 μW	525 kHz	5	67	142.9	28.5	[71]
5	1.2	6.5 mW[q]	30 MHz	3	62.5	72.2	40.6	[22]
6	1.2	23 μW	320 kHz	3	57	24.0	47.9	[20]
7	2.5	13 mW	600 kHz	7	77[x] 71[x]	3095.0 3095.0	61.8[x] 240.3[x]	[68]
8	1.2	6.5 mW[q]	100 MHz	3	50	21.7	217.0	[22]
9	5	580 μW	40 MHz	2	41.3	7.25	537.5	[72]

[q] In quiescent condition.
[m] With the maximum input signal.
[x] [68] quotes a maximum signal of 2 V_{pp}, a noise floor of 196 μV_{rms} and a dynamic range of 77 dB. These numbers are inconsistent. The value corresponding to the quoted maximum signal and noise is 71 dB.

dynamic range as defined above is also the S/N when the THD is 40 dB[2] (See Chapter 2).

To the best of the author's knowledge, these are the filters with the highest dynamic range per unit power consumption. It can be seen that even the best of these is more than a hundred times as power hungry as the passive RC filter.