Copyright © 1997, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

LOW-POWER DECIMATION FILTER DESIGN FOR MULTI-STANDARD TRANSCEIVER APPLICATIONS

by

Carol J. Barrett

Memorandum No. UCB/ERL M97/88

11 December 1997

COVER IT-6

LOW-POWER DECIMATION FILTER DESIGN FOR MULTI-STANDARD TRANSCEIVER APPLICATIONS

by

Carol J. Barrett

Memorandum No. UCB/ERL M97/88

11 December 1997

ELECTRONICS RESEARCH LABORATORY

College of Engineering
University of California, Berkeley
94720

Low-Power Decimation Filter Design for Multi-Standard Transceiver Applications

by
Carol J. Barrett

Master of Science in Electrical Engineering
University of California, Berkeley

Professor Paul R. Gray, Advisor

Abstract

Recent efforts in the design of wireless RF transceivers focus on high integration and multi-standard operation. Higher integration can be obtained by using receiver architectures, such as wide-band IF with double conversion (WIF), that perform channel select filtering on-chip at baseband. Performing this baseband channel select filtering in the digital domain allows for the programmability necessary to adapt to the different channel bandwidths, sampling rates, and CNR requirements of multiple communication standards. At the back of a wide-dynamic range sigma-delta modulator, a decimation filter can select a desired channel in the presence of both strong adjacent channel interferers and quantization noise from the digitization process.

A low-power decimation filter that performs channel select filtering for the GSM (European cellular) and DECT (European cordless) standards is presented. Automatic gain control is used within the filter to reduce the dynamic range and power consumption. Since the two standards have different blocking profiles and CNR

requirements, the filter can adapt to each standard and reduce its power consumption by powering down unused circuitry. For a 3.3V power supply, the filter consumes 4.4mW in GSM mode and 16.4mW in DECT mode, for input sampling rates of 12.8MHz and 44.8MHz respectively.

Acknowledgments

I would like to thank my advisor, Professor Paul Gray, for his guidance throughout this project. I would also like to thank Professor Bernhard Boser for answering numerous technical questions, especially at the beginning of this project, and for reading my thesis. I am especially grateful to Arnold Feldman for mentoring me throughout this project. He was always available to answer my questions, listen to my ideas, and give advice. Without his help, this project would not have gone as smoothly as is did.

There are many others whom I appreciate for taking the time to discuss research and answer questions. Some of these are Chris Rudell, for helping me understand receiver architectures and for answering dozens of questions, Jeff Weldon, Keith Onodera, Sekhar Narayanaswami, and Andy Abo for many interesting technical and non-technical discussions, Arthur Abnous for giving advice on my project from a digital perspective, Dave Lidsky for help with Power Play, and Herb Huang for being a sounding board for all of my ideas.

Lastly, I would like to thank my mother and the rest of my family for their support, and several others, some mentioned above and some unmentioned, for their friendship.

This work was supported by the National Science Foundation and AT&T Bell Laboratories Fellowships. The standard cells used in this project were provided by SGS-Thompson Microelectronics, with help from Marco Sabatini.

Table of Contents

Chapter 1

Introduction

1.1 Motivation

Current research on radio frequency (RF) communication transceivers emphasizes both higher integration, to meet consumer demand for low-cost, low-power, small-form factor personal communication devices, and the ability to adapt to multiple communication standards. Higher integration can be obtained by using receiver architectures and circuit techniques that eliminate the need for external components. Utilizing a receiver architecture that performs channel select filtering on chip at baseband allows for the programmability necessary to adapt to multiple communication standards. [1]

Wireless telecommunication standards currently used throughout the world have channel bandwidths ranging from 6.25 kHz to 1.728 MHz. [2] A multi-standard receiver that performs baseband channel select filtering in the digital domain must have an analog-to-digital converter (ADC) with a wide dynamic range that can accommodate undesired channels as well as the desired. It must also be able to adapt to the various dynamic range requirements and sampling rates of the standards that are implemented. A wide dynamic range sigma-delta modulator can be used to meet these requirements for multiple standards.

Following the sigma-delta modulator, a lowpass digital decimation filter can select a desired channel in the presence of both strong adjacent channel interferers and quantization noise from the digitalization process. The decimation filter also reduces the sampling rate from the oversampled rate of the sigma-delta to the Nyquist rate of the channel. Including programmability in the filter, allows it to adapt to the different channel bandwidths, worst case interferers, and bit error rate (BER) requirements of the different standards, while keeping the power consumption at a minimum.

1.2 Research Goals

The focus of this project is to design a power-optimized decimation filter that can perform digital channel selection for the GSM (cellular) and DECT (cordless) communication standards. A secondary goal is to identify possible approaches to decimation filter design for more general multi-standard RF communications receivers. The results of this work are summarized below.

- Designed a decimation filter that meets the GSM and DECT specifications without significantly degrading the receiver noise figure. Estimated power consumption is 4.4mW and 16.4mW for GSM and DECT modes respectively.
- Showed that automatic gain control can be used within a decimation filter that performs channel filtering to reduce power consumption and area.
- Developed a power saving technique for a fixed multi-standard implementation, where filter blocks and sub-blocks can be powered down, enabling the filter to adapt to each standard with reduced power consumption.

1.3 Thesis Organization

The second chapter introduces the application and specifications for this project, including a brief discussion of receiver architectures, emphasizing high-integration and multi-standard capability, the GSM and DECT telecommunication standards as they pertain to the decimation filter design, and a description of the sigma-delta analog-to-digital converter. The

chapter is concluded with a summary of the decimation filter operation and specifications. Chapter 3 describes the method used to design a decimation filter architecture for a low-power, transceiver application. System level issues including automatic gain control, tap encoding, and quantization are included here. The hardware implementation of the decimation filter is discussed in Chapter 4. Two possible multi-standard implementations are described, and power and area estimates are given. Chapter 5 discusses issues related to multi-standard decimation filter design, and proposes future work in this area. Conclusions from this work are given in Chapter 6, and the appendices provide more detailed information on the noise figure calculation, the control signals from the DSP, and the filter taps.

Chapter 2

System Description and Specifications

2.1 Introduction

The performance of a receiver is measured by its ability to select a desired channel in the presence of strong adjacent channels, selectivity, and by its minimum detectable signal, sensitivity. Each wireless communication standard defines the sensitivity and selectivity that a receiver must meet while maintaining a certain bit error rate (BER). To implement more than one wireless standard, the transceiver must be able to meet the performance requirements of each standard and adjust to the different channel bandwidths and carrier frequencies. This chapter discusses receiver architectures and channel select filtering with emphasis on integration, and potential for multi-standard operation. A description of the GSM and DECT standards implemented in this project are next, followed by a brief discussion of the sigma-delta modulator used for analog-to digital conversion. The basic operation of a decimation filter is explained, and the specifications for this project are given.

2.2 Receiver Architecture

A low-power receiver for portable applications should be as highly integrated as possible. The conventional super-heterodyne receiver architecture shown in Figure 2.1 has very good performance due in part to its discreet components. It uses an external image reject (IR) filter and an external, narrow-band, IF channel filter within the receive path. Recent examples

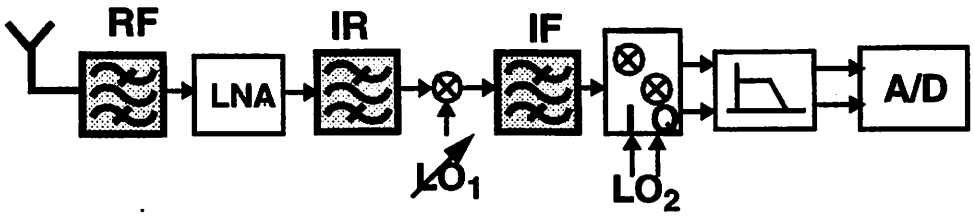

Fig. 2.1: Super-Heterodyne Receiver

of super-heterodyne receivers can be found in [3][4][5][6].

The direct-conversion homodyne approach shown in Figure 2.2, has no external components within the receive path. Since its local oscillator (LO) is at the radio frequency (RF), a DC offset is created at the output of the mixer, reducing the dynamic range of the receiver. Baseband circuitry and the mixer also contribute 1/f noise, which further reduces the sensitivity. Because the LO must be a high-frequency, low-phase-noise, channel select

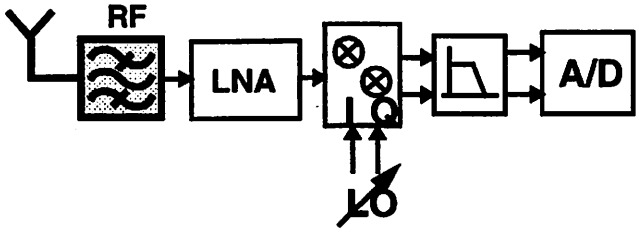

Fig. 2.2: Direct-Conversion Homodyne Receiver

synthesizer, it will be difficult to implement with low-Q elements that are available on chip.[7] Recent examples of homodyne receivers can be found in [8][9]. Because the channel filtering is performed at baseband, unlike the super-heterodyne, it is possible to design baseband circuits

for the homodyne receiver that are multistandard capable, however, the noise and DC offset must be reduced to achieve adequate dynamic range.

A recently proposed low-IF architecture [10] shown in Figure 2.3 has no external components in the receive path, but unlike the homodyne receiver, it mixes the signal to a low-IF, avoiding the DC offset problem. This architecture can be used to effectively implement standards with a relatively small channel bandwidth. For a wide-bandwidth standard such as DECT, it would be extremely difficult to implement a low-power multi-standard baseband section using current technology since the clock rate of the A/D would be very high.[11].

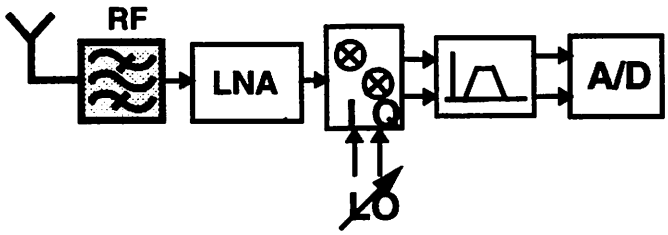

Fig. 2.3: Low-IF Architecture

The Wide-Band Intermediate Frequency With Double Conversion (WIF) architecture [12] shown in Figure 2.4 has no discreet components within the receive path. Because the channel filtering is done at baseband instead of at IF, a wide-band IF filter can be used, eliminating the need for an off-chip IF filter. With this architecture, channel filtering can be done digitally, allowing the receiver to adapt to multiple communication standards. This approach requires a high dynamic range A/D converter that can adapt to the different channel bandwidths and sampling ratios that the different standards require. In contrast to the super

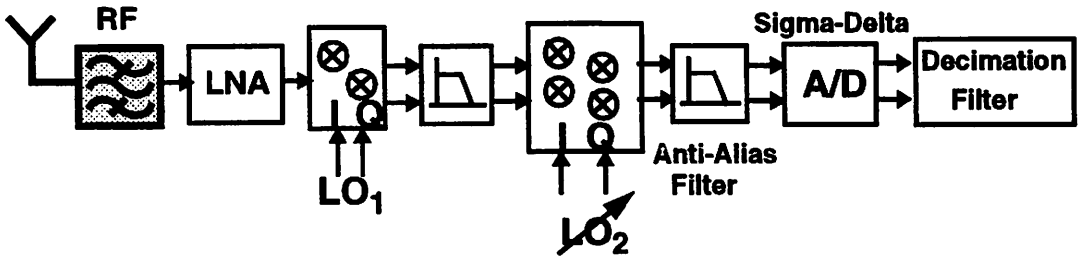

Fig. 2.4: Wide-Band IF Double Conversion Receiver

heterodyne, the WIF has a fixed LO1 and a variable LO2. A fixed LO1 can use a higher reference frequency, and therefore have a higher phase-locked loop (PLL) bandwidth and lower phase noise. This allows the phase noise requirement for LO2 to be relaxed.[7] This architecture uses image-reject mixers, eliminating the off-chip IR filter. The frequency synthesizers use on chip spiral inductors and varactors, so the only other off-chip component needed is the crystal.

The wide-band IF architecture as it appears in Figure 2.4 has 1/f noise in the baseband section that significantly degrades the noise figure of the receiver when it is operating under the GSM standard. To combat this problem, the receiver includes the programmability necessary to translate the spectrum to a very low-IF, one channel away from DC. This spectrum passes through the anti-alias filter and the ADC, and is then filtered by a simple digital low-pass filter to prevent quantization noise aliasing when it is translated to baseband by a digital image-reject mixer. When the receiver is operating in DECT mode, the 1/f noise is tolerable, so the digital mixer section is bypassed as shown in Figure 2.5.

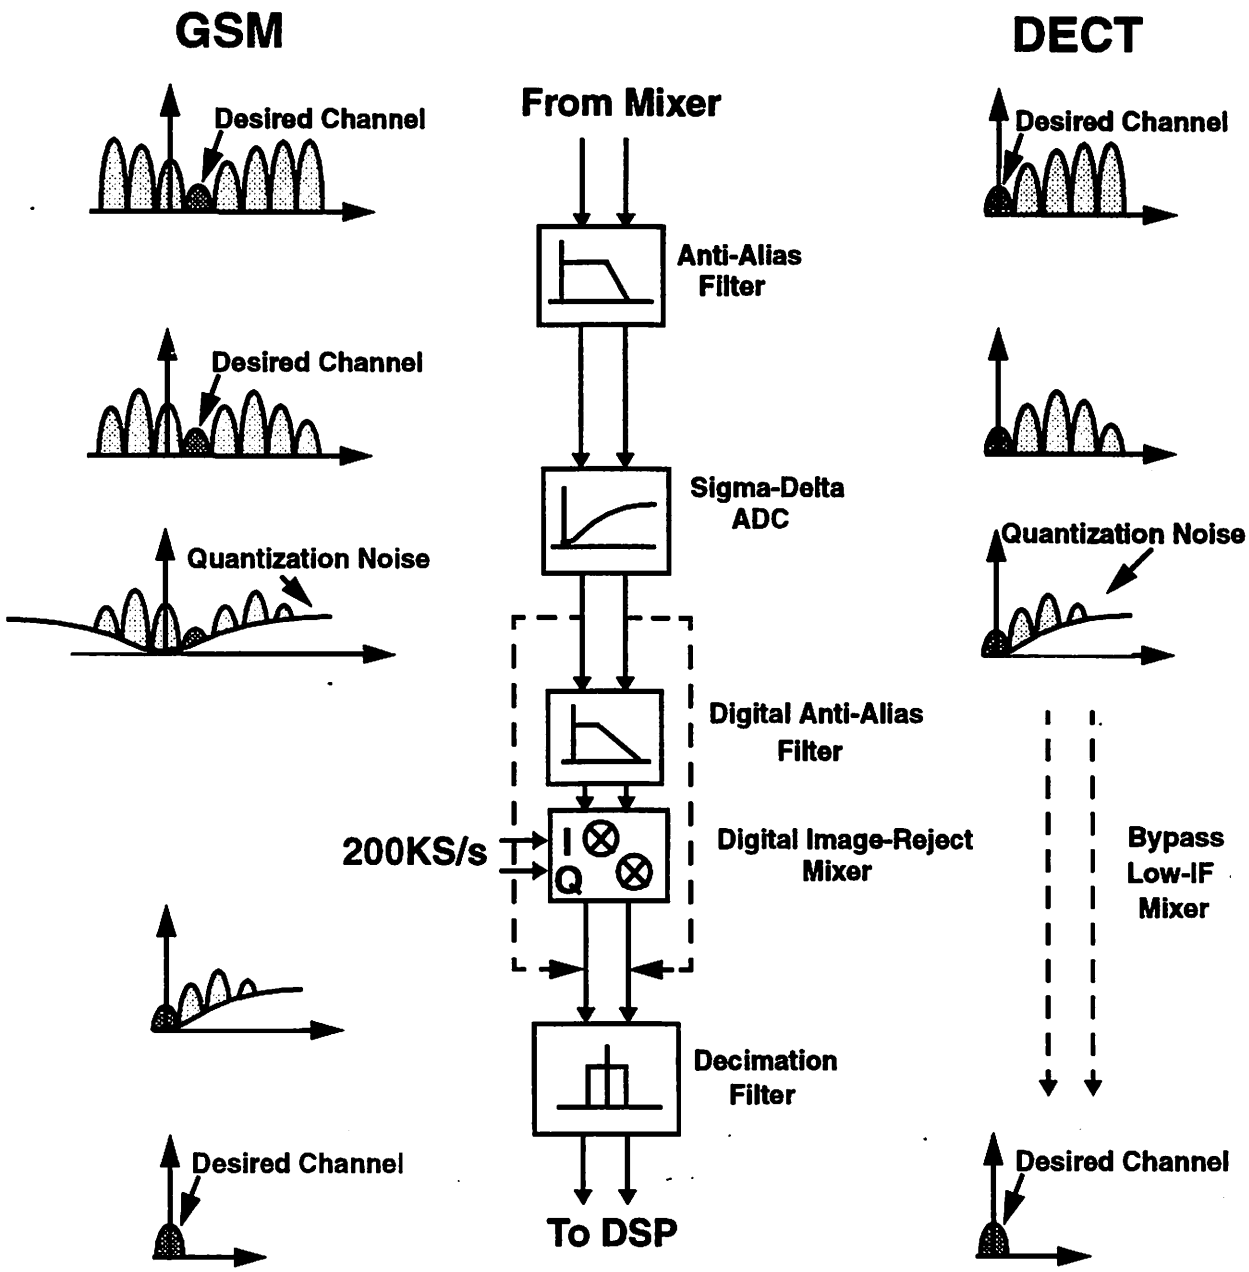

Fig. 2.5: Multi-Standard Baseband Architecture

2.3 Baseband Channel Select Filtering

The previous discussion on RF receivers, indicates that architectures that perform channel selection at baseband can be highly integrated and are amenable to multi-standard operation. For a channel select filter to be multi-standard, it must be able to adapt to different channel bandwidths and dynamic range requirements, implying some programmability. In general, channel select filtering can be divided into analog and digital. Analog filtering can be performed by continuous time R-C filters and by switched-capacitor filters. References [13] and [14] are examples of analog channel select filtering for DECT and IS-95 CDMA respectively.

Figure 2.6 and Figure 2.7 show how channel selection is performed in the analog and digital domains. In the analog case, a high dynamic range, highly linear channel select filter is needed to remove the large undesired signals. The input to the A/D is simply the desired channel, so the dynamic range can be much lower. Using automatic gain control (AGC) after the channel select filter can further reduce the ADC dynamic range. In the digital case, the scenario is reversed; the sigma-delta ADC must have a large dynamic range to accommodate the undesired channels. The decimation filter then removes the adjacent channels and noise. For this application, the digital approach using a sigma-delta modulator and decimation filter was chosen because it can be made programmable more easily than an analog solution.[11]

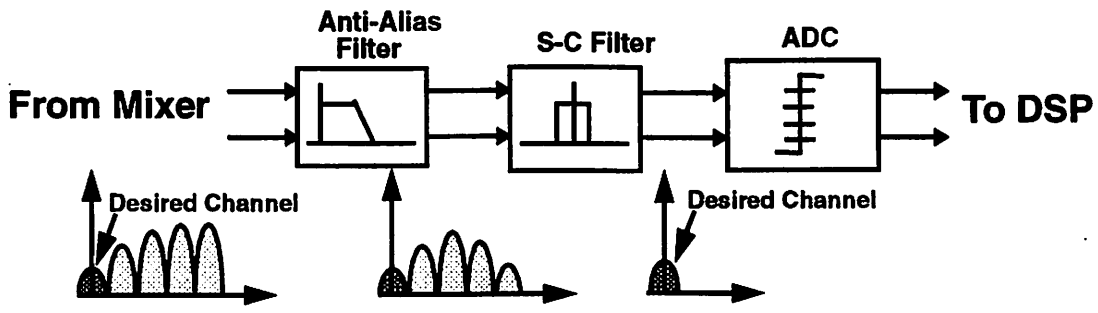

Fig. 2.6: Analog Channel Select Filtering

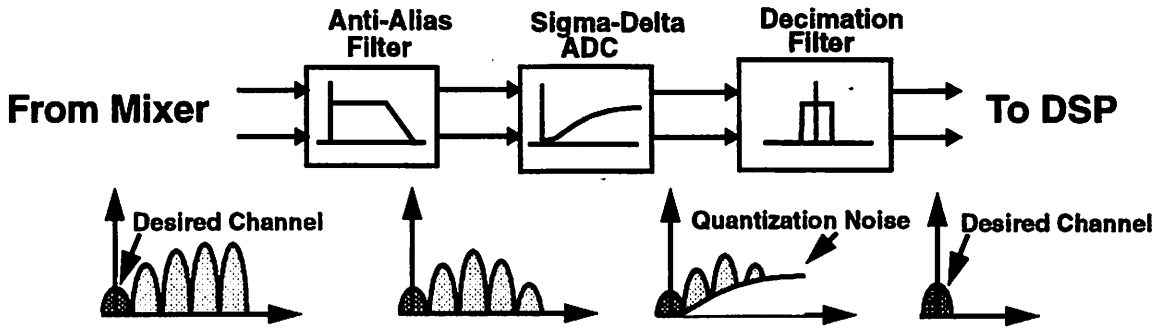

Fig. 2.7: Digital Channel Select Filtering

2.4 GSM and DECT Specifications

Before the specifications of the decimation filter can be determined, it is necessary to define some of the receiver specifications and blocking profiles. Table 2.1 gives the receiver specifications for GSM and DECT that are relevant to the decimation filter.[15] The sensitivity is the minimum detectable signal at the input of the receiver. The input noise is the thermal noise from the antenna (-173.8 dBm) times the channel bandwidth, which is 200kHz for GSM and 1.728MHz for DECT. The input SNR is the ratio of the input signal to the input noise and the required CNR is the ratio of the carrier to noise at the output of the receiver that is needed to meet the minimum BER requirement. The noise figure is the ratio of the input SNR to the

Table 2.1:

Standard	Sensitivity (dBm)	Input Noise (dBm)	Input SNR (dB)	Required CNR (dB)	Required NF (dB)
GSM	-102	-120.8	18.8	9	9.8
DECT	-83	-112.3	29.3	10.3	19.0

output SNR, and is a relative measure of how much noise is added by the receiver components to the desired signal. Appendix 1 discusses the calculation of these numbers in more detail.

Figure 2.8 and Figure 2.9 show the worst case blocking profile and adjacent channel interferers for GSM.[16] Blockers are large undesired signals within the same cell. A cell is composed of a cellular base station and the physical area that is within its transmit range. Adjacent channel interferers are large undesired signals from neighboring cells. All channels used within the same cell must be separated by two channels. For example, cell A may contain

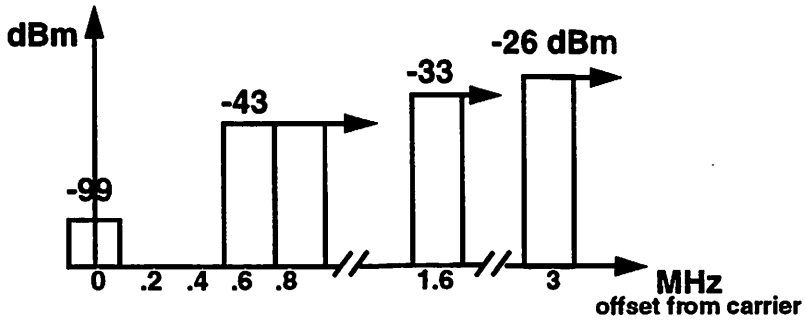

Fig. 2.8: GSM Blockers

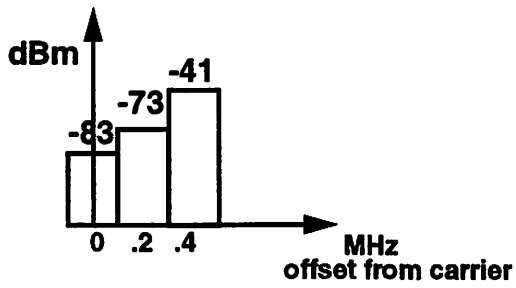

Fig. 2.9: GSM Adjacent Channel Interferers

channels 1, 4, and 7. Cell B, a cell next to cell A, can contain channels 2, 5, and 8. If channel 1 is the desired channel, channels 4 and 7 would be blockers and channel 2 would be an adjacent channel interferer. This example is further illustrated in Figure 2.10 which shows four cells and

the corresponding channels that may be used in each. The distinction between blockers and

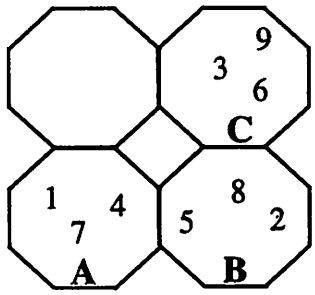

Fig. 2.10: Cell Reuse Scheme

adjacent channel interferers in the GSM specification occurs because the worst case undesired signals are referenced to different desired signal levels. The blocking profile for DECT [17] is shown in Figure 2.11.

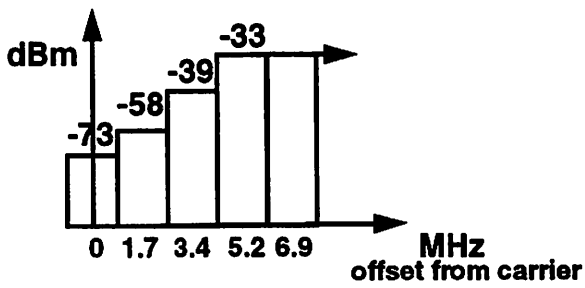

Fig. 2.11: DECT Blockers

2.5 Sigma-delta Modulator

If the proposed WIF radio architecture is to be multi-standard, an analog to digital converter (ADC) that can accommodate GSM and DECT must be used. Since the resolution for GSM is 16 bits, a high-resolution ADC is needed. A sigma-delta modulator is very attractive for this application for several reasons. The same modulator architecture, differing only in the oversampling ratio, 64x for GSM and 32x for DECT, can be used to meet both specifications. The OSR, and therefore the dynamic range, can be easily lowered for DECT so that the

minimum power solution for each standard is obtained. Because the data is oversampled, the requirements for the anti-alias filter are relaxed compared to an analog solution. [11]

The sigma-delta modulator coarsely quantizes the input at a high rate, trading accuracy in the analog circuitry for speed. The quantization noise is effectively pushed out of the desired band into higher frequencies, so that when the signal is decimated, a higher resolution can be obtained. The dynamic range of the modulator depends on the oversampling ratio (OSR), which is the sampling rate divided by twice the nyquist rate, and the order of the modulator. Increasing the order results in more complexity, while increasing the OSR requires faster op-amp settling.

For this work, a 4-th order sigma-delta using a 2-2 cascade architecture is assumed. Simulations are performed using the architecture in Figure 2.12. with non-idealities such as finite op-amp gain, and capacitor mismatch. Using (Eq 2-1), where L is the order and M is the oversampling ratio, it can be seen that if an OSR of 64 is used, the GSM dynamic range

$$DR = \frac{3}{2}\frac{2L+1}{\pi^{2L}}M^{2L+1} \qquad \text{(Eq 2-1)}$$

requirement of 98 dB can be achieved. An OSR of 32 will meet the DECT requirement of 85 dB. The design of the modulator is beyond the scope of this project; more in-depth discussions can be found in references [19], [20], and [21].

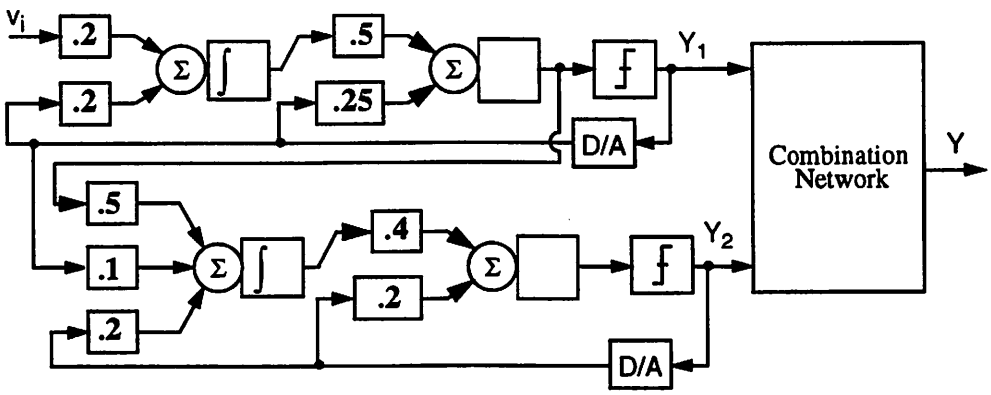

Fig. 2.12: Sigma-Delta Architecture

2.6 Decimation Filter

The function of a decimation filter is to remove all of the out-of-band signals and noise, and to reduce the sampling rate by M, where M is the over-sampling ratio (OSR). By averaging M values of the coarsely quantized sigma-delta output, the filter gives a high-resolution output at the low rate. The decimation filter in this project is designed to meet the GSM and DECT specifications for the receiver architecture described in section 2.2 and the sigma-delta modulator described in section 2.5 .

The input to the decimation filter is a six bit data stream with three components: the modulator quantization noise, the desired signal, and the blockers. The worst case blockers were described in section 2.4 and shown in Figure 2.8 and Figure 2.9 for GSM and Figure 2.11 for DECT. The worst case quantization noise can be obtained by exercising the sigma-delta with a large low-frequency tone. The FFT of the output is shown in Figure 2.13.

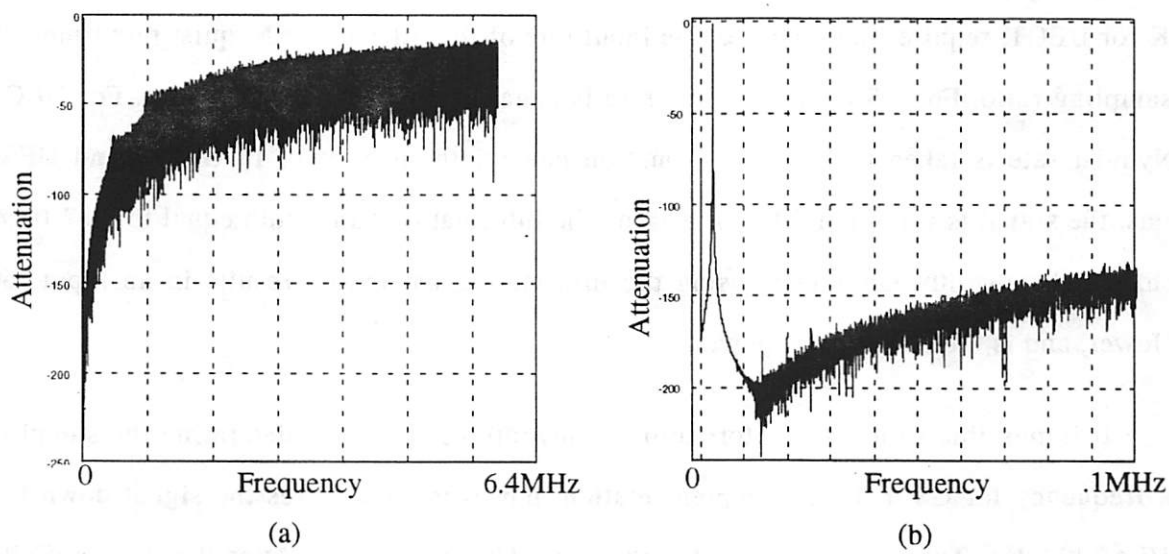

Fig. 2.13: Worst Case Quantization Noise at the Sigma-Delta Output (a) entire spectrum (b) zooming in on bandwidth of interest.

Now that the receiver specifications have been described and the worst case filter inputs have been determined, the specifications of the decimation filter can be calculated. Table 2.2 contains the important specifications. The filter contribution to the receiver noise figure should be negligible, so a maximum addition of 0.01 dB is the design goal. Assuming a receiver gain of 46 dB, the necessary carrier to noise ratio at the output of the filter is 21dB for GSM and 10.2dB for DECT. The noise figure and output CNR analysis is given in Appendix 1 .

Table 2.2:

Specification	GSM	DECT
CNR	21dB	10.2dB
Dynamic Range	16 bits	14 bits
OSR	64	32
Nyquist Rate	200kS/s	1.4MS/s
Power	minimize	minimize

Specification	GSM	DECT
Input Rate	12.8MS/s	44.8MS/s
Input Bits	6	6
Phase	linear	linear
Area	under 2mm x 2mm	under 2mm x 2mm

The dynamic range, determined by the minimum and maximum receiver input power, is 16 bits for GSM and 14 bits for DECT. The modulation schemes used, GMSK for GSM and GFSK for DECT, require linear phase. The input rate of the filter is the Nyquist rate times the oversampling ratio. For GSM, the Nyquist rate is equal to the channel bandwidth. For DECT, the Nyquist rate is taken to be the information bandwidth times two. In GMSK and GFSK systems, the signal is Gaussian filtered, making the information bandwidth equal to 0.82 times the channel bandwidth. Obviously, using the information bandwidth results in an input rate 18% lower, and is therefore lower power.

It is possible to use the information bandwidth of 82kHz to determine the sampling clock frequency for GSM. In this implementation, however, LO2 mixes the signal down to a low IF of 200kHz. This is done to combat the large 1/f noise at DC. After the A/D, a digital image-reject mixer is used to bring the signal down to baseband, as shown inFigure 2.5. Having the desired channel at low-IF prevents significant reduction in the A/D clock. The filter clock

can still be reduced to multiples of 82kHz, but it will not be a multiple of the A/D clock and the increased complexity in the synthesizers is not worth the power savings in the baseband.

2.7 Summary

The WIF receiver architecture is used for this application because it has no external components within the receive path, and is therefore smaller and lower-power, and because it can accommodate multiple telecommunications standards. The DECT (cordless) and GSM (cellular) telecommunication standards are chosen as a test vehicle for this project in order to demonstrate standards with significantly different channel bandwidths, carrier frequencies, and sensitivity requirements. Channel select filtering is performed at baseband in the digital domain because the sigma-delta ADC and the decimation filter can both adapt to a range of channel bandwidths and dynamic range requirements.

Chapter 3

Decimation Filter Architecture and Design

3.1 Introduction

The first step in designing a decimation filter is to decide which types of filters will be used and where decimation will occur. This chapter explores the issues involved in choosing a filter architecture for a low-power receiver application. The relative power of several architectures is compared, resulting in the four-stage architecture that is chosen to implement this filter. Due to the linear phase constraint of the DECT and GSM standards, this discussion will be limited to FIR filters. To further reduce the power consumption, automatic gain control is included in the filter. The final section discusses choosing and encoding the filter taps.

3.2 Decimation in Stages

It is possible to remove the undesired channels and noise with a single filter, and then decimate to the output rate. An example filter is shown in figure 3.1 for a given transition band, stop band rejection, and decimation ratio, M. As shown in [21], the number of taps in an FIR filter is proportional to the stopband rejection and to the sampling rate divided by the transition band. This means that if the transition band is a small percentage of the sampling rate, the filter

will have many taps, in this example, 100. The power is proportional to the number of taps and the rate at which they operate. By decimating in stages, the total number of taps in the filters is reduced, and subsequent filters operate at lower sampling rates, further reducing the power consumption.[22] Figure 3.2 shows decimation in two stages for the same overall transition band and oversampling ratio. Each filter only needs to reject the part of the undesired signal that will alias into the desired band or baseband. For a review of downsampling and aliasing,

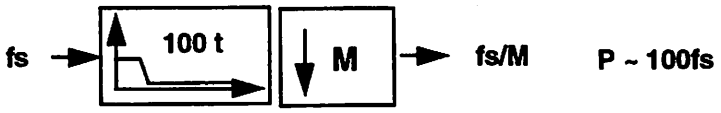

Fig. 3.1: Decimation in One Stage

refer to [21] and [34]. This allows the transition band of the first filter to be a much wider percentage of the sampling rate. As a result, the first filter only has 5 taps. The second filter decimates from the intermediate rate to the output rate, so it's transition band must be the same as the filter in 3.1. Because it is operating at a lower frequency, however, the transition band is a larger percentage of the sampling rate and the number of taps is smaller, in this example, 50.

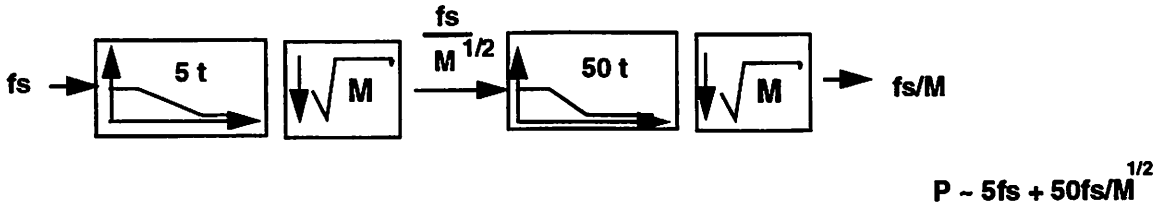

Fig. 3.2: Decimation in Two Stages

Assuming that the power consumed by a filter is proportional to the number of taps (capacitance) and the frequency, the power consumed by the multi-stage implementation would be between 8% and 15% of the power consumed by the single filter implementation for

oversampling ratios ranging from 32 to 128. If the filter is implemented such that the area is also proportional to the number of taps, there would be a savings of more than 60%.

In the example presented above, significant power savings was obtained by increasing the number of filter stages. The next step in the architecture design is to identify which types of filters would be amenable to a low-power solution, and then determine the possible combinations of filters and decimation ratios that could be used to meet the specifications.

3.3 Comb Filter

It has been shown in [26] that the comb filter is an efficient way to decimate the output of the sigma-delta modulator to four times the Nyquist rate. The comb filter is a running average of M_c input samples, where M_c is the decimation ratio of the comb filter. The magnitude response is given in (Eq 3-1) and the z-transform in (Eq 3-2), where k is the number of cascaded comb filters and T=1/fs is the input sampling period. [20]

$$|H(f)| = \left| \left(\frac{\sin \pi f M_c T}{M_c \sin \pi f T} \right) \right|^k \qquad \text{(Eq 3-1)}$$

$$H(z) = \left(\frac{1}{M_c} \cdot \frac{1 - z^{-M_c}}{1 - z^{-1}} \right)^k \qquad \text{(Eq 3-2)}$$

For a sigma-delta modulator of order L, a cascade of k=L+1 comb filters is needed to adequately attenuate the quantization noise that would alias into the desired band. [26] For a desired bandwidth fb, the out of band signal that will alias into that band is located at multiples of the decimated output frequency, fs/M_c, with a bandwidth equal to two times fb.

For this application, a 4th order modulator is used, so a 5th order comb filter is needed. Since the OSR is 64 for GSM, M_c is 16, and the output rate of the comb filter is

0.8MHz. The bandwidth of interest is 100KHz, so the out of band signals surrounding multiples of 0.8MHz with a bandwidth of 200KHz will alias into the desired band after decimation occurs. It can be seen from Figure 3.3, a plot of (Eq 3-1), that the comb filter provides a notch at each of the frequencies that will alias to baseband. The remaining frequencies will alias outside of 100KHz, and must be removed by subsequent filtering.

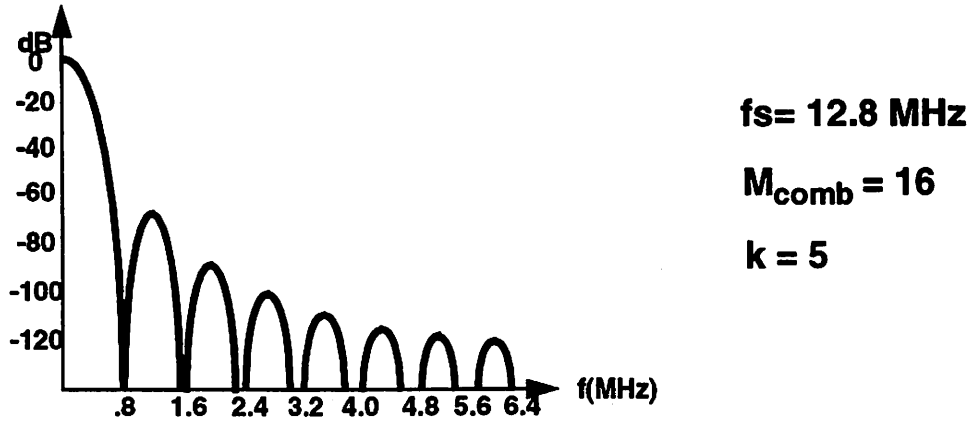

Fig. 3.3: Comb Filter Magnitude Response (GSM)

Comb filters can be efficiently implemented by separating (Eq 3-2) into numerator and denominator sections and by moving the numerator section after the resampling operation. (Eq 3-2) can be written as

$$H(z) = \left(\frac{1}{M_c}\right)^k \cdot (1 - z^{-M_c})^k \cdot \left(\frac{1}{1 - z^{-1}}\right)^k \qquad \text{(Eq 3-3)}$$

The numerator section followed by downsampling by M_c is equal to the same downsampling operation followed by a differentiator, $(1-z^{-1})$, as shown in Figure 3.4.(a) and (b). [22] The denominator is a delay free integrator. If a cascade of comb filters is used, it can be replaced by an integrator with delay, which pipelines the integrators so that the filter can run at a higher clock speed. Using integrators with delay as in (Eq 3-4) adds latency to the filter, but does not affect the magnitude response of the filter.

$$\left(\frac{z^{-1}}{1-z^{-1}}\right)^k \qquad \text{(Eq 3-4)}$$

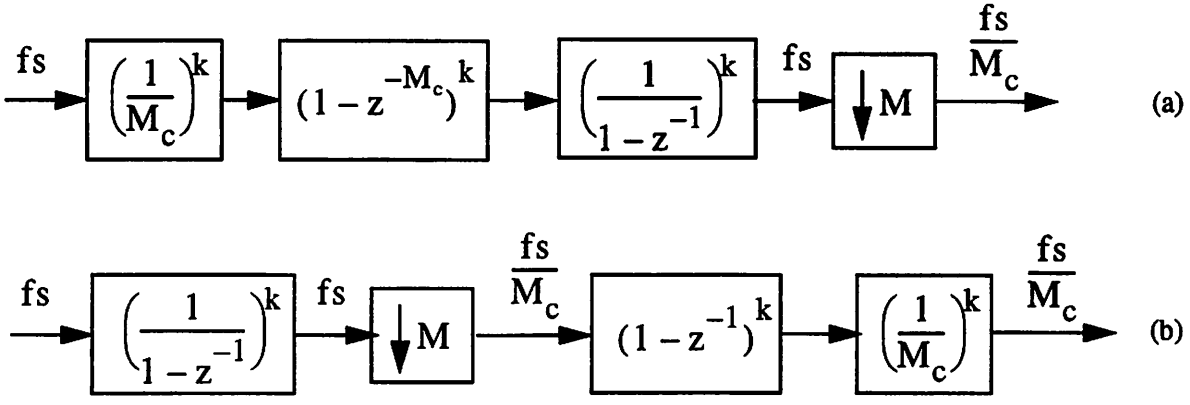

Fig. 3.4: (a) Direct implementation of (Eq 3-3). (b) Comb filter implementation with numerator section after the resampling operation.

Comb filters can be efficiently implemented without multipliers, using adders and registers. A hardware implementation of a 2nd order comb filter is shown in Figure 3.5. Moving the numerator section after the decimation reduces the amount of computation at the high rate, so that the power is minimized. The integrators have the tendency to grow without bound, which would eventually overflow the registers used in the implementation shown in Figure 3.5. In [22], it is shown that if 2's complement wrap-around arithmetic is used, the overflow problem can be avoided as long as the register width is greater than or equal to the value given by (Eq 3-5), where B_{in} is the number of input bits from the sigma-delta. For this application, Bin=6, M_c=16, and k=5, so the register width is 26 bits.

$$\text{RegisterWidth} = k \cdot \log_2(M_c) + B_{in} \qquad \text{(Eq 3-5)}$$

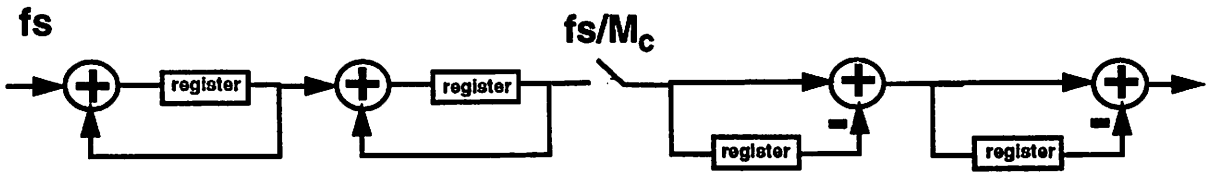

Fig. 3.5: Implementation of 2nd order Comb Filter

3.4 Halfband Filters

The remaining 4X of decimation can be performed in one stage by an FIR filter or in two stages by FIR and/or halfband filters. Halfband filters are a subset of symmetric FIR filters. All of their odd coefficients are equal to zero except for the center one, which is equal to 1/2. This results in fewer taps, less hardware, and lower power. Halfband filters are constrained to be equiripple filters with the property that their frequency response has a value of 0.5 at the frequency fs/4, where fs is the sampling rate. Because of this restriction, they are not suitable for decimation by more than two. An idealized halfband filter is shown in Figure 3.6, where t is the transition band. Note that noise in the second half of the transition band is only slightly attenuated by the filter before it aliases into the first half. The signal contained in the stop band aliases into the pass band.

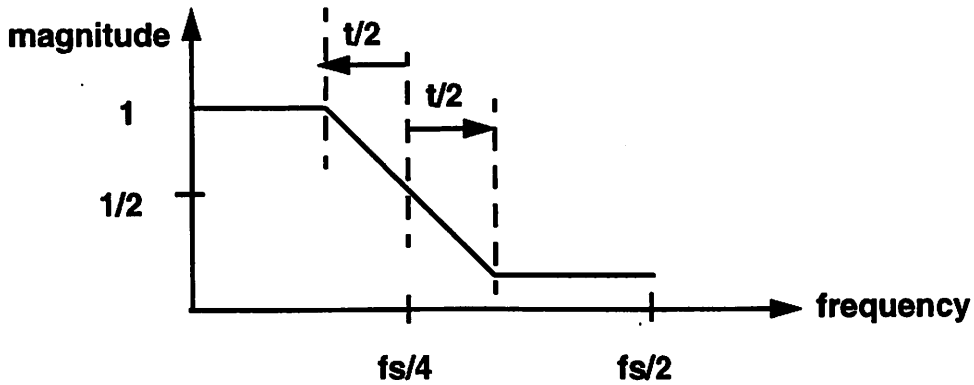

Fig. 3.6: Magnitude Response of a Halfband Filter

Halfband filters can be efficiently implemented with a polyphase direct-form filter, which allows the filter to run at the decimated rate instead of the input rate, reducing the power consumption by approximately one-half. [21] The polyphase structure of a 14th order halfband filter is shown in Figure 3.7. In this implementation, the effective number of taps, five in this example, is much less than the filter order. As with any FIR filter, the number of taps in a halfband filter depends on the stop band rejection and the relative transition bandwidth.

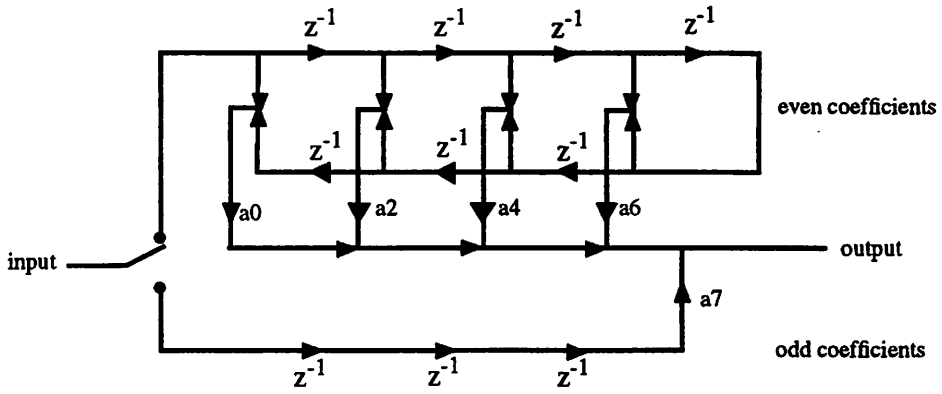

Fig. 3.7: Polyphase Direct-Form Implementation of a 14-tap Halfband Filter

3.5 Power Comparison of Architectures

Decimation filters used in audio applications are generally designed for a worst case out of band rejection that is not defined by large undesired blocking signals, as in [24] and [25]. This results in a uniform stopband attenuation for all of the filters following the comb filter. As explained in Section 2.4, the worst case out of band signals for a radio transceiver are more clearly specified. Since the goal is to meet the carrier to noise requirement with the minimum number of taps, the stop band should be tailored to the blocking specs, which may yield a non-uniform response.

Several possible filter architectures are compared to find the lowest power solution. In this calculation, it is assumed that the power is directly proportional to the number of filter taps, and that any overhead not proportional to filter length is negligible. It is also assumed that all architectures operate at the same voltage supply, so the supply is normalized to one. The operating frequency of each filter is also taken into account, so that the power equation $P=CV^2f$ [23] is modeled as in (Eq 3-6).

$$\text{Power} \propto \sum(\text{taps} \times \text{frequency} \times \text{wordlength}) \qquad \text{(Eq 3-6)}$$

The design specifications given in Table 2.2 must be met for each of the worst case blockers, applied separately to the receiver. The blocking specifications shown in Figure 2.8, Figure 2.9, and Figure 2.11 are given in dBm, which is power, for GSM and DECT. The blocking specifications in Figure 3.8 are given in dB*FS*, which is the magnitude relative to the full scale of the sigma-delta converter. At the decimation filter input, the blockers in Figure 3.8 will remain the same, with the addition of quantization noise. It is assumed that the 3MHz blocker is attenuated by 3dB by the anti-alias filter.

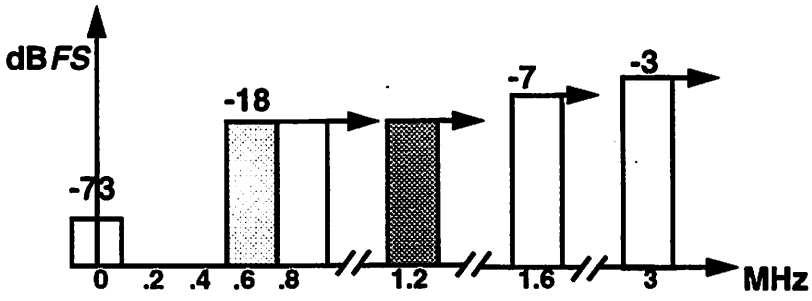

(a) GSM Blockers

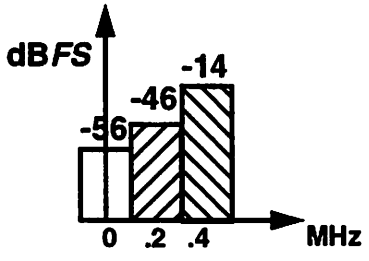

(b) GSM Adjacent Channel Interferers

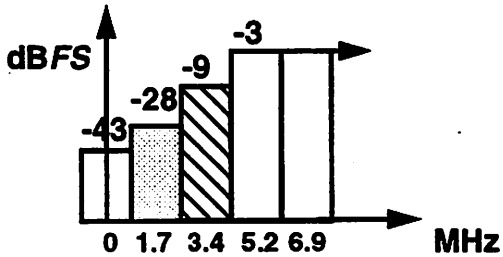

(c) DECT Blockers

Fig. 3.8: Blocking Specifications at Sigma-Delta Input

3.5.1 The First Stage: Comb Filter

A comb filter is used for the first stage for the reasons discussed in Section 3.3. Figure 3.9 shows the GSM blocking profile including quantization noise at the output of a fifth-order comb filter that decimates by 16. The magnitude of the quantization noise power within each channel is shown above an arrow. In Figure 3.9 (a), it can be seen that the 600kHz blocker has aliased to the channel adjacent to baseband. It has been attenuated by the comb filter from -18 dB to -49.6 dB. Many other blockers also alias into this band, but the magnitude of the aliased component is smaller than the 600kHz blocker. Since the worst case blockers are applied independently to the receiver along with the smaller desired channel, only the worst case must be considered. Figure 3.9 (a) also shows an undesired signal in the band between 300kHz and 400kHz. The blocker that causes the largest aliased component in this band is the 1.2MHz blocker. The adjacent channel interferers shown in Figure 3.9 (b) are slightly attenuated by the comb filter. The total noise budget is the maximum noise power, including quantization noise and finite word length effects, that can be added to the desired channel and still meet the CNR requirement.

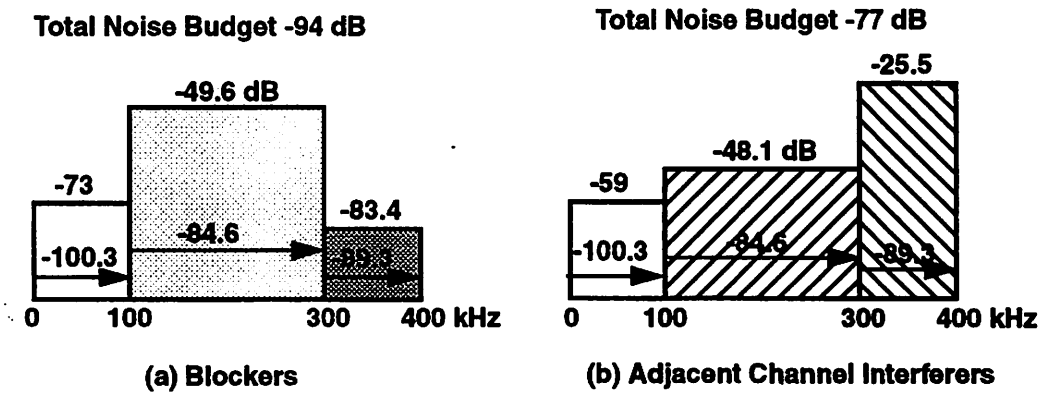

Fig. 3.9: Worst Case GSM Blockers and Noise at Comb Filter Output

In order to implement the DECT standard, the decimation ratio of the comb filter must be changed from 16 to 8, since the DECT OSR is 32. The DECT implementation uses the information bandwidth instead of the channel bandwidth to determine the sampling rate, as

explained in Section 2.6. This makes the calculation of the aliased components more difficult than to the GSM case. The bandwidth of interest is 700kHz and the output sampling rate is 5.6MHz, so the four bands of width 700K up to 2.8MHz must be considered. The two major contributors to undesired signal power in this region are the 1.7MHz blocker and the 3.4MHz blocker. These are shown in Figure 3.10(a) and (b), where the arrows represent the quantization noise power. Since the CNR requirement for DECT is 10.2dB, the maximum noise power added by the quantization noise and aliased blocker components is -53.2dB.

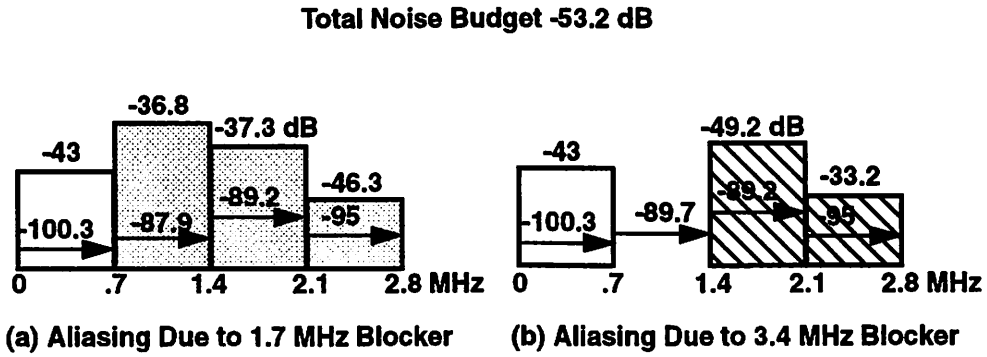

Fig. 3.10: Worst Case DECT Blockers and Noise at Comb Filter Output

The out of band components in Figure 3.9 and Figure 3.10 must be removed by subsequent filters with 4x decimation. If for GSM, a comb filter that decimated by 32 were used, later filtering would only be required to remove energy in the band from 100kHz to 200kHz. This approach will not work, however, because too much of the quantization noise and blocker power aliases into the desired band, degrading the overall receiver noise figure.

In (Eq 3-6) the power is dependent on the word length, the frequency, and the number of taps, which is five in the numerator and five in the denominator for a fifth order comb filter. The word length for the comb filter was given by (Eq 3-5) as $k\log_2(M_c)$ + Bin. Reducing any of these parameters will reduce the word length, and therefore, the power. Bin from the sigma-delta is fixed at 6 bits. k cannot be reduced because the filter would not provide enough quantization noise rejection. It is possible, however, to reduce M_c. If Mc is reduced to 8, the

word length can be reduced from 26 bits to 21. The differentiators of this implementation will run at twice the rate of those in the original implementation, but the overall comb filter power is lower. At the output of this filter, the remaining decimation is 8x instead of 4x. The comb filter with decimation by 8 provides less rejection of the lower frequency blockers, so there is more noise power to be rejected by the subsequent filters. This results in higher power consumption of the remaining stages and the net power of the entire filter is higher.

One way to reduce the word length without reducing the overall comb filter decimation ratio is to separate the comb filter into two filters, each with an M_c of 4. The first comb filter has a word length of 16 bits, with the integrators running at 64x and the differentiators running at 16x. The second comb filter has a word length of 26 bits, which is the same as the original filter, but the integrators are running at 16x instead of 64x. This is illustrated in Figure 3.11. A power comparison between this filter and the original comb filter is given in Table 3.1. The two-stage implementation is more power than the single cascaded comb filter due to the added integrators and differentiators running at 16x.

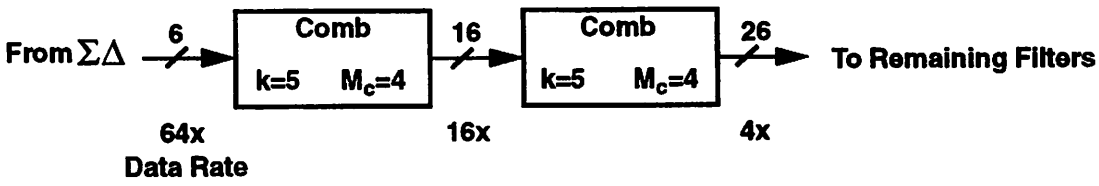

Fig. 3.11: Comb Filter With Two-Stage Implementation

Table 3.1:Power Comparison of Comb Filter Implementation

Single Comb				
Type	Number	Word Length	Rate	Power Metric
Integrator	5	26	64x	8320
Differentiator	5	26	4x	520
			Total:	8840

Table 3.1:Power Comparison of Comb Filter Implementation

Two-Stage Comb				
Type	Number	Word Length	Rate	Power Metric
Integrator	5	16	64x	5120
Differentiator	5	16	16x	1280
Integrator	5	26	16x	2080
Differentiator	5	26	4x	520
			Total:	9000

Another implementation of comb filters uses a bit-serial approach, which breaks the comb filter into several stages. [35] There is one stage for each 2x of comb filter decimation, and each stage operates at a progressively lower rate. The word length for the first stage is equal to Bi plus k. Each stage has a word length of k more bits than the last. For certain values of M_c , Bi, and k, this structure results in lower power. For a large Bin of 6, k=5, and M_c=16, however, this approach consumes more power than the conventional comb filter.

3.5.2 Remaining Filters

The worst case blockers in Figure 3.9 include four different interfering signals. Because the blockers are applied to the receiver separately, it is necessary to separate the profile into four cases. Figure 3.12 shows the desired and undesired signal powers for the 200k, 400k, 600k, and 1.2MHz blockers. Note that the quantization noise has been lumped with the blockers if they fall into the same channel. The remaining filters must reject all four of these scenarios so that the CNR is within the required level. The total allowable noise to meet this specification is given in the figure above each profile.

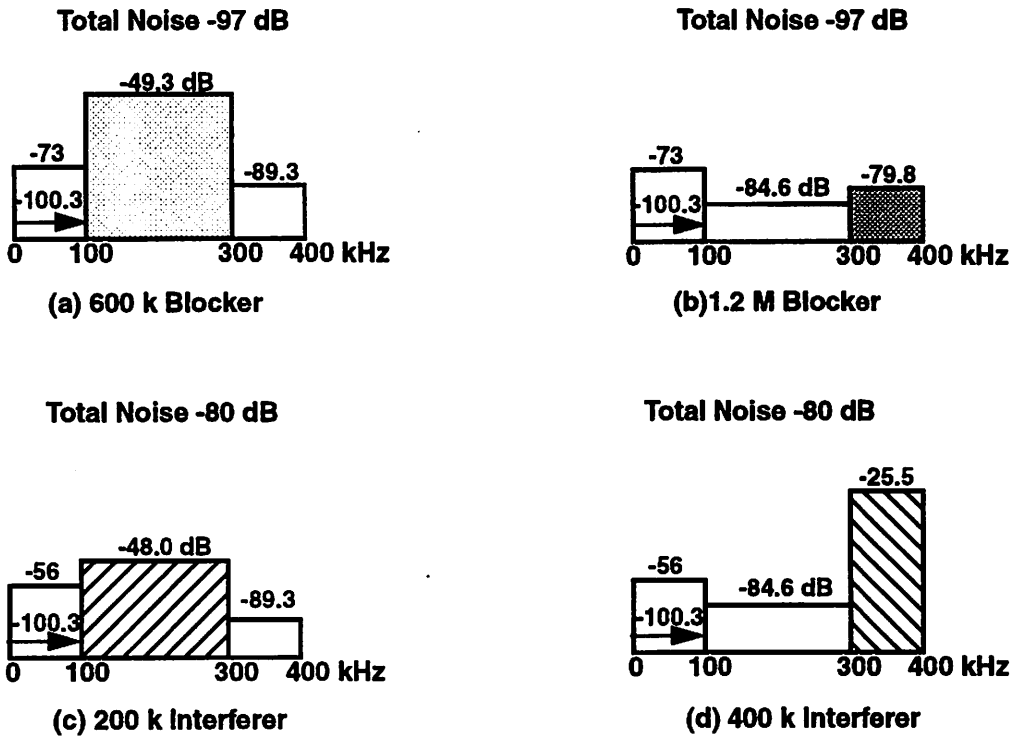

Fig. 3.12: Worst Case Second-Stage Inputs (GSM) From (a) 600kHz, (b) 1.2MHz, (c) 200kHz, and (d) 400kHz Channels

There are several possible filter combinations that can be used to achieve the last 4x of decimation. Recent low-power decimation filters [24] [25] have used an architecture consisting of a comb filter followed by two halfband filters. The solution with two halfbands consumes significantly less power than one using a single FIR followed by decimation by 4x. For this application, however, using two halfband filters is still very costly. Remember from Section 3.4 that the halfband filter has an attenuation of -6dB at fs/4. This means that some of the energy above 100kHz will remain after filtering and will be counted as part of the noise budget. In order to meet the specification with two halfbands, a very narrow transition band is needed in the second filter, resulting in a large number of taps.

Figure 3.13 illustrates the problem with using two halfband filters of a reasonable length. It can be seen in the figure that there is a large undesired signal at 200kHz at the output of the comb filter. The signal is only slightly attenuated by the first halfband, and must

therefore be removed by the second halfband. At 100kHz, the filter attenuates the blocker by 6dB. A transition band from 100k to 107kHz, contains a noise power that is greater than the total allowable noise for the entire filter. The transition band of the second halfband would have to be 3kHz in order to reject enough of the neighboring blocker, which results in a filter with 371 taps!

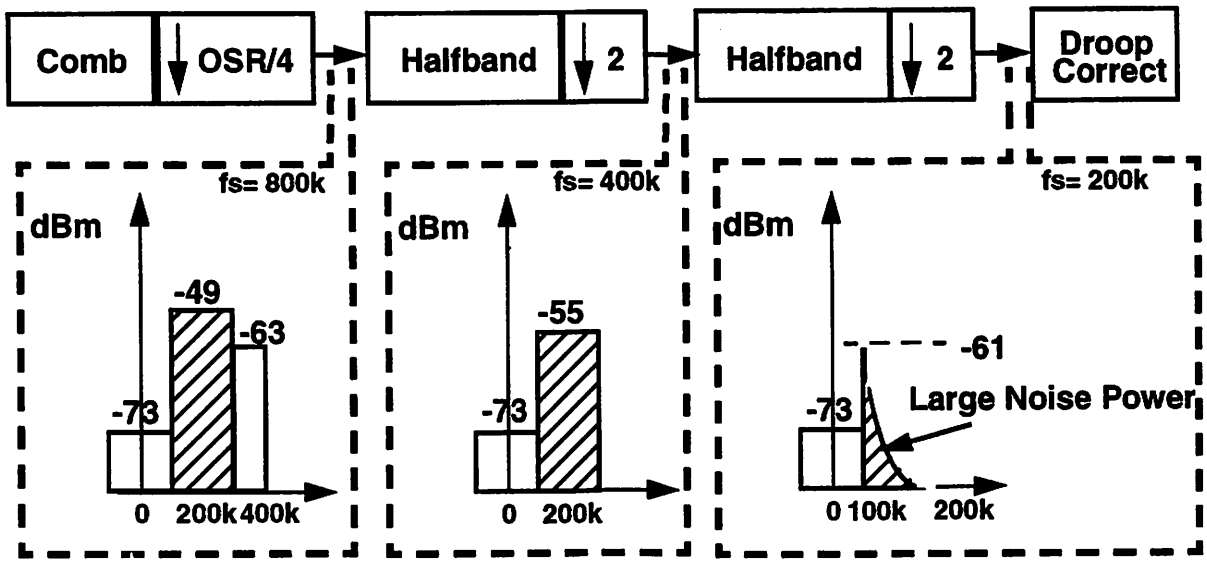

Fig. 3.13: Problem With Two-Halfband Architecture

This problem can be overcome by several methods. If an FIR is used in place of the second halfband, the transition band need no longer be symmetric about 100k. The transition band can be shifted lower in frequency to reject undesired signal power above 100kHz. Another possible solution is to add an FIR filter after the two halfbands to reject the transition band power. Recall from Section 2.6 that the information bandwidth of the desired signal is 82kHz. The un-rejected noise power in the second halfband filter's transition band (see Figure 3.13) will alias between 93k and 100kHz. Because this is above the information bandwidth, it can be removed with a FIR filter.

Another possibility is to reduce the power of the 600kHz blocker *before* it aliases to the 200k channel by using a sixth order comb filter. Increasing the order of the comb filter adds

an extra integrator and differentiator, increasing the number of adders and registers from 10 to 12. It also increases the required word length according to (Eq 3-5) from 26 to 30. The rest of the filter can be implemented with two halfbands, with about one-fourth the taps in the second halfband as compared to the fifth order comb, two halfband filter implementation.

It can be seen from Figure 3.3 that the comb filter is not flat in the passband, but has some droop. This droop must be corrected by an FIR filter. If an FIR filter is used in the design, it can include the frequency response necessary to correct the droop. If there is no FIR, an extra droop-correct filter must be added to the end of the filter. As can be seen from the following comparison, the droop-correct filter contributes negligible power consumption to the overall filter.

The FIR and droop-correct filters are designed using the Remez exchange algorithm. The halfband filters are designed using the Remez algorithm and a technique from [27]. The stopband of each filter is chosen so that the overall filter can meet the CNR requirement. Each filter within the decimation filter will contribute some noise to the total. In the architecture comparisons, the total noise is distributed to the individual filters such that the power of the overall filter is minimized. In order to do this, some iteration is required.

3.5.3 Power Metric Calculation

The power of all of the filter architectures mentioned in Section 3.5.2 can be compared using (Eq 3-6), as in Section 3.5.1, where the power of two comb architectures was compared. For the filters after the comb filter, a word length of 13 bits is assumed. It is also assumed that each tap can be implemented with three adders instead of a multiply. A halfband filter of length 15, for example, has nine non-zero taps. Because the filter is symmetric, four of those nine taps are identical to another tap, resulting in 5 unique taps that must be implemented. In this comparison, only the effective number of taps is used since that more closely emulates the actual hardware. For a comb filter, each tap is one adder and one register.

To make the comparisons consistent, the number of taps should be multiplied by the number of adders per tap. This yields a slightly different power metric, given in (Eq 3-7).

$$\text{Power} \propto \sum(\text{taps} \times \text{adds/tap} \times \text{frequency} \times \text{wordlength}) \qquad \text{(Eq 3-7)}$$

The power metric of the sixth order comb filter is calculated as follows. There are 6 integrators operating at 64x and 6 differentiators operating at 4x. Both have a word length of 30 bits, resulting in a power proportional to 11,832. The fifth order comb filter was calculated to be 8840. The sixth order comb is approximately one-third more power.

Table 3.2 shows the five architectures that are compared, the power metric for each, and the normalized value relative to the lowest power architecture. The first architecture consists of a comb filter and an FIR followed by decimation by 4x. The FIR operates at it's input rate, which is four times the output rate, and therefore consumes a lot of power. The second architecture uses a sixth-order comb filter to attenuate the large 600kHz blocker before it aliases into the adjacent channel. The extra comb stages and increased word length consume much more power than is saved in the later halfband stages. The third architecture is the commonly used comb filter followed by two halfband stages. It's power is 17% higher than the lowest power solution. Architecture 4 replaces the second halfband with an FIR filter. While the total number of taps is reduced, the FIR must run at the input rate of 2x, unlike the second halfband, which operates at the decimated rate of 1x. The power of these two architectures is approximately the same, although the one with fewer taps will have a smaller area. The final architecture uses an FIR to filter out the undesired signal that has aliased into the region between the information bandwidth of 82k and the channel bandwidth of 100kHz. This solution has the lowest power of the five and is therefore used in this work.

Because this comparison makes several simplifying assumptions and does not take clock distribution and other overhead circuitry into account, the normalized power numbers given in Table 3.2 are very approximate. A more accurate analysis was not performed because,

in addition to potential power savings, Architecture 5 can be easily adapted from the GSM specification to DECT. This adaptation is further explained later in the chapter.

Table 3.2:Power Comparison of Filter Architectures

	Architecture	Power Metric	Normalized Power
1	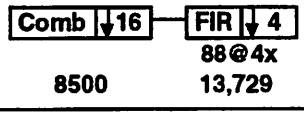	22,228	2.03
2	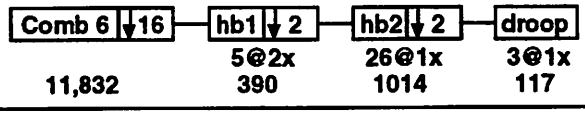	13,353	1.22
3	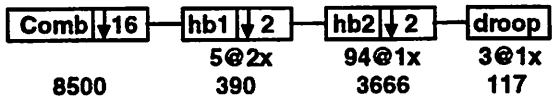	12,637	1.17
4	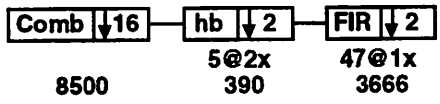	12,556	1.15
5	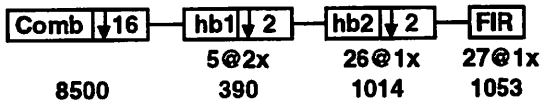	10,957	1

3.6 Power Reduction with Automatic Gain Control

Recall that the smallest desired channel that the decimation filter sees is -73dB*FS*. For this case the total allowable noise is -94 db*FS*. An equation for the truncation error is given in (Eq 3-8) where b is the number of bits retained. [27]

$$error < \pm 2^{-b} \qquad \text{(Eq 3-8)}$$

In the above power comparison, it was assumed that all of the filters following the comb filter had 13 bits. This yields an error of -78.3dB according to (Eq 3-8), which is obviously not below the total allowable noise requirement. To meet this requirement, 21 bits would be needed, assuming a 5 bit accuracy, 2 bits for error, and 13 bits to meet the -73dB dynamic range for the worst case input. If gain is applied to the desired signal, the total allowable noise will scale with the signal. This allows a reduced dynamic range, requiring fewer bits to be used, resulting in a lower power implementation.

Under worst case conditions, the desired channel for GSM is -73dB*FS* and the largest blocker is -3dB*FS* at the input of the decimation filter. Gain cannot be applied to the signal because it is already close to full scale. At the output of the comb filter (see Figure 3.9), however, the largest blocker is -25.5dB for a desired signal of -56dB. If the desired signal is greater than -56dB, the worst case blocker is -30.6dB higher. For a desired signal smaller than -56dB, it can be assumed that the adjacent channel interferer specification no longer applies. In the blocking profile, the worst case blocker is 23.4dB larger than the desired signal. Since, for small desired signals, the largest blockers are well below full scale, gain can be applied until the largest signal approaches full scale.

To minimize overhead, it is assumed that the gain is provided by a programmable bit-shifter, which is controlled by the DSP based on the received signal strength (RSSI) of the desired channel. Therefore, the gain is applied in steps of approximately 6dB. The worst case dynamic range needed will be 30.6dB, the difference between the desired signal and the largest blocker, plus 6dB, the gain step, plus 3dB to keep the largest signal 3dB below full scale. Suppose a blocker is 26.9dB. A bit-shift of 4 provides 24dB of gain, which would make the signal above full scale minus 3dB. So, a bit-shift of 3 must be used, resulting in a magnitude of -8.9dB. The worst case desired signal is 30.6 dB below this, or -39.5dB. Now, just over 6 bits

are needed to achieve the necessary dynamic range. Adding the 21dB CNR and 2 bits for error, gives a 13 bit word length.

Using the same method as in Table 3.2, it can be calculated that the power of filter architecture 5 would be 14% higher with 21 bit word lengths in the final filters. This comparison neglects the power consumed by the AGC, which is only a few percent of the entire filter if standard cell muxes are used, and less if a logarithmic shifter is used.

Because the worst case difference between the desired signal and the largest blocker is small for DECT, resulting in a nine bit word length, the implemented word length is set by the GSM specifications. Appendix 2 gives the DSP controls needed to operate the AGC in GSM and DECT modes.

3.7 FIR and Halfband Filter Design

The FIR and halfband filter taps are designed in MATLAB using the Remez-exchange algorithm. The halfband filter program also utilizes the method explained in [28]. In order to meet the CNR specified in Table 2.2, the total noise allowed is -94dB for the out of band blockers and -77 dB for the adjacent channel interferers. This includes the ADC quantization noise, remaining transition band noise, and the signal power that aliases into baseband after decimation. To meet this specification and minimize the number of taps needed, the stopband rejection of the first halfband is -55dB, the second halfband is -50dB, and the fir is -26dB. The frequency response for the filters is given in Figure 3.14.Given these rejections, the filter lengths are 11, 75, and 26 respectively. The effective number of taps, as explained in Section 3.5.3, are 4, 20, and 13. The tap values are listed in appendix3.

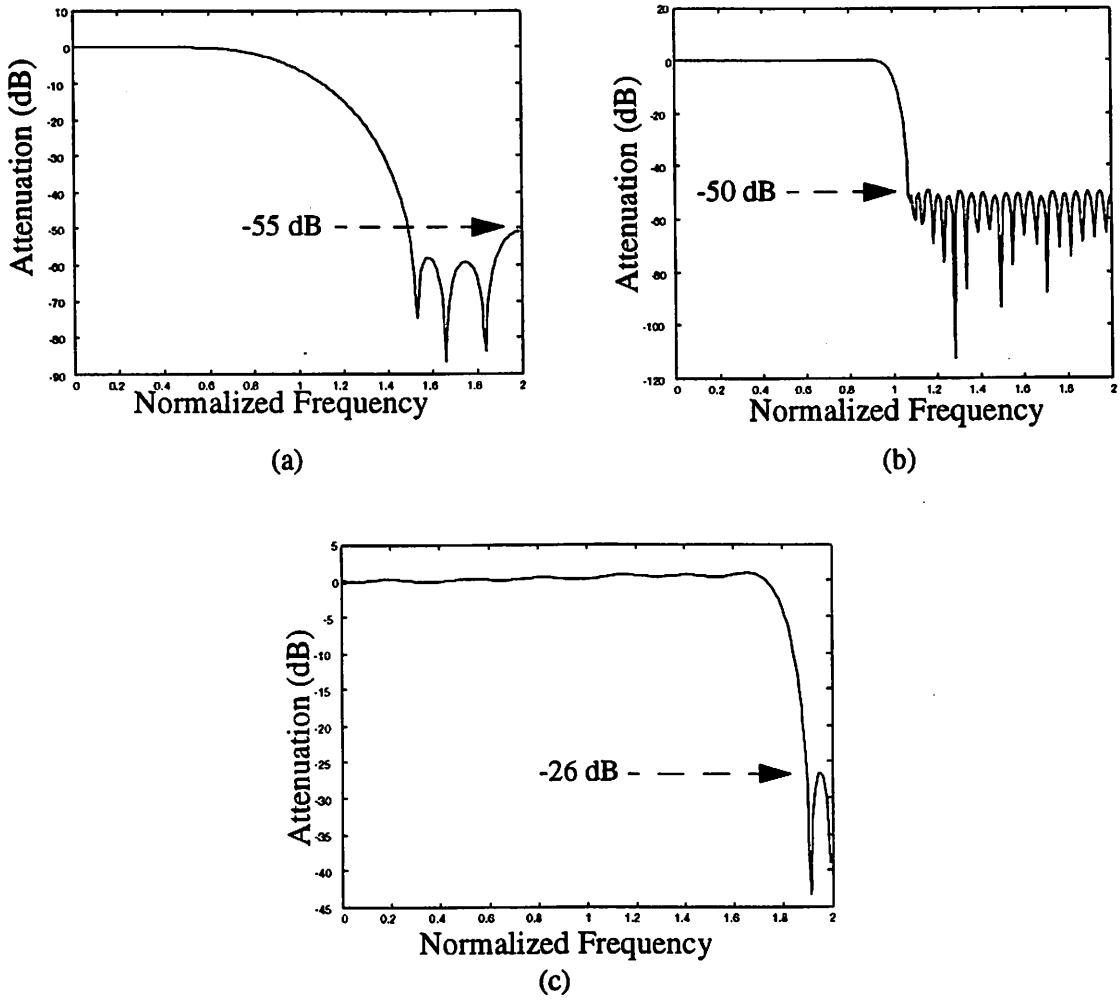

Fig. 3.14: GSM Filters (a) First Halfband, (b) Second Halfband, (c) FIR

3.7.1 Tap Quantization and Encoding

Multipliers are expensive in terms of both power and area. One way to implement taps without using multipliers is to use cannonic signed digit encoding (CSD). Cannonic signed digit encoding quantizes each tap to a finite number of powers of 2, and takes advantage of the fact that multiplication by a power of two is simply a bit shift. For example, $0.47045898 = 2^{-1} - 2^{-4} + 2^{-5} + 2^{-9} - 2^{-12}$. Instead of using a multiplier, this tap can be implemented with five shift

and add operations. To save even more power, the list of powers of 2 can be truncated. The amount of truncation depends on how much quantization error can be tolerated in the taps and how much degradation of the filter frequency response can be tolerated. Complex algorithms have been proposed to optimize the number of CSD's used in a filter for a given frequency response. [30] For this work, a simple MATLAB program was written to expand the filter taps into powers of 2, and quantize the number of CSD's to within a given error for each tap. The amount of error was then made as large as possible, reducing the number of CSD's, without degrading the frequency response of the filter too much. The net result was an average of 2.5 CSD's per tap for the halfband and FIR filters. According to the rule-of-thumb given in [30], each tap should have one CSD for each 20dB of rejection, plus, one extra for large tap values. Using this rule, the first halfband filter quantized with the MATLAB program has 2 extra CSD's. The second halfband has 8 extra, and the fir has 6 extra CSD's. So, it may be possible to eliminate 16 of the 93 CSD's if a complex local search algorithm were designed, resulting in 12.7% fewer adders (16/126 adders). Designing such an algorithm is complex, and though it would be useful, it is beyond the scope of this project.

The Goodman filters [31], a collection of hardware efficient decimation and interpolation filters, were considered as a possibility for the first halfband. The F8 filter has a frequency response that is closest to the specification, but its stopband rejection is over-designed by 5dB. As a result, it has more registers and more CSD's than the filter generated by MATLAB.

3.7.2 Nested Multiplication

Nested Multiplication reduces roundoff noise by combining smaller partial products and then shifting the result instead of shifting completely before adding. [24] For example, multiplying a number x by 0.3125, where $0.3125 = 2^{-2} + 2^{-4}$, would normally be written as

$$0.3125x = 2^{-2}x + 2^{-4}x \,, \qquad \text{(Eq 3-9)}$$

which requires shifts of 2 and 4. The multiplication can also be performed as

$$0.3125x = 2^{-2}(x + 2^{-2}x) \,, \qquad \text{(Eq 3-10)}$$

which requires a maximum shift of 2. For a given word length, the nested multiplication in (Eq 3-10) will have less roundoff noise than the multiplication in (Eq 3-9). Using nested multiplication also reduces hardware complexity by reducing the maximum adder word length and largest shifter needed.

3.8 Modifying the Filters for DECT Operation

The filters discussed in the last few sections were for the GSM specification. If the filter is operating in DECT mode, the requirements are much less stringent. Rejections of only -26dB and -28dB are needed in the first and second halfband filters to reject the blockers and quantization noise. Recall that halfband filters are equiripple filters. For this application, the signal is phase modulated, so a large degree of ripple can be tolerated. The filter is constrained to 0.6dB of ripple, so each halfband is limited to 0.2dB. This causes the design to be limited by the ripple, not by the blockers as is the GSM case. Therefore, the stopband rejections for the DECT halfbands is set to -33dB. The FIR filter is not needed at all, so a droop-correct filter must be added. There are several options for designing the DECT filters, depending on the implementation used.

The first option is to simply use the GSM filters for DECT operation. This requires no additional area, but consumes excess power. Another option is to design new halfband filters

for the above rejections. This yields the lowest power consumption, but can have a large impact on the area, depending on the implementation chosen. A third option is to use the GSM filters, eliminating the excess CSD's so that the stopband rejections just meet the DECT specifications. This can be done by using a separate power supply, and shutting off the un necessary hardware, resulting in the filters shown in Figure 3.15. This will not reduce the number of registers as a custom solution would, but it does reduce the number of adders in the halfband filters by approximately 50%. Similarly, the effective word length for the halfband filters can be reduced to the nine bits needed to meet the DECT dynamic range requirement. This solution has minimal impact on area, and its power consumption is within 2% of a custom solution. The best choice for the DECT filters is implementation dependent and will be discussed in Chapter 4.

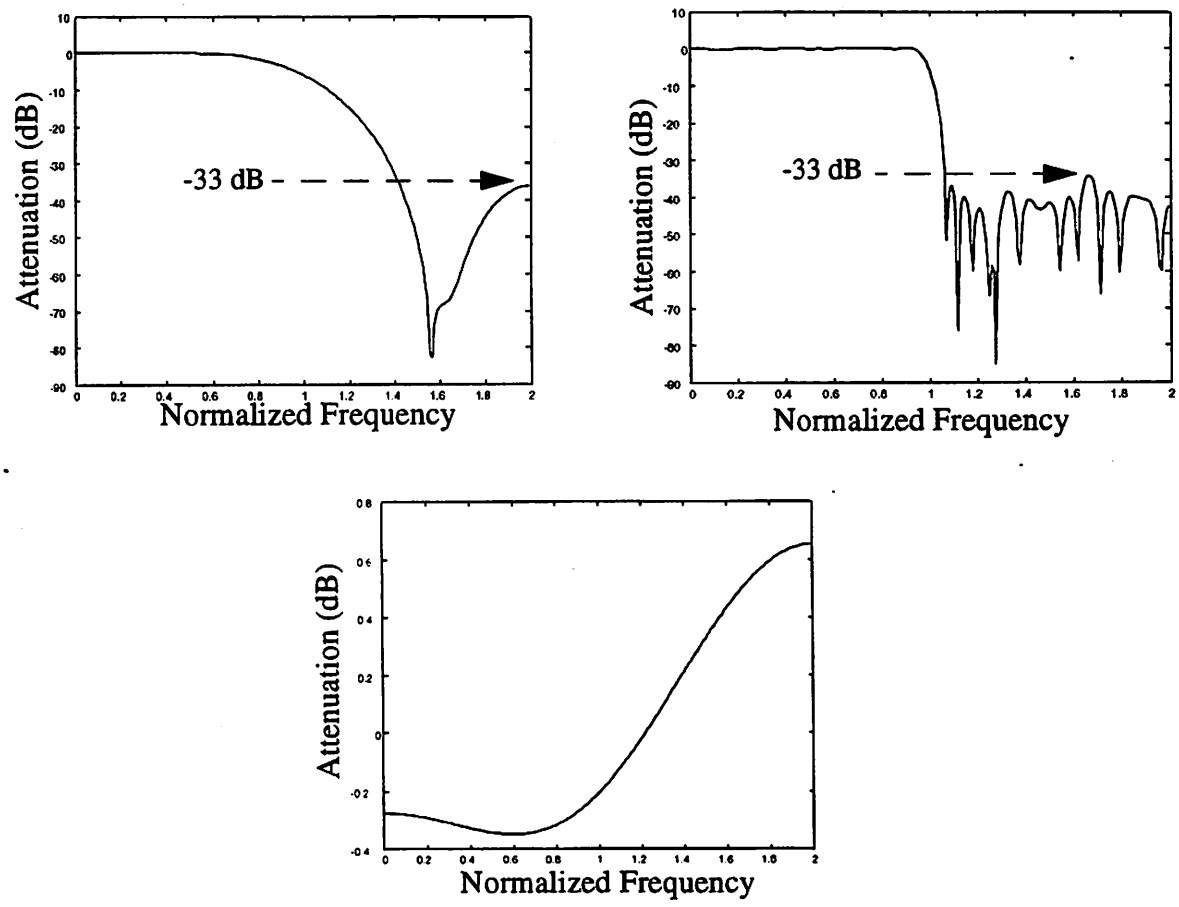

Fig. 3.15: DECT Filters (a) First Halfband, (b) Second Halfband, (c) Droop-Correct

3.9 Summary

This chapter discusses decimation filter architecture design for a low-power receiver application. The important conclusions are summarized below.

- Decimation in stages saves both power and area.
- The filter architecture should be tailored to the worst case blocking profile, which can result in a non-uniform stopband rejection, and non-traditional filter architectures.
- Automatic gain control should be used as early in the filter as possible to reduce the word length, and therefore the power, of later stages.
- Cannonic signed digit tap encoding should be used to avoid power hungry multipliers.
- Nested multiplication should be used to reduce the amount of truncation noise and the maximum adder and shifter widths.
- If a multi-standard solution is needed, the power used in each mode can be minimized by powering down all unused hardware blocks.

Figure 3.16 shows the architecture chosen to implement the low-power decimation filter. The filter operates in two modes, GSM and DECT. Several controls from the DSP determine what mode the filter is in and how much AGC is needed. In GSM mode, 16x of decimation follows the comb filter. The amount of AGC is based on the RSSI and the worst case blocking profile, and is given in Appendix 2. Following the AGC are two halfband filters and an FIR filter. In DECT mode, the amount of decimation after the comb filter is changed to 8x. The amount of AGC is set by the RSSI and the DECT blocking profile. In the halfband filters, unnecessary hardware is shut down, reducing the word length from 13 bits to 9, and degrading the stopband rejection of both filters to the values shown in brackets in Figure 3.16. The FIR filter is replaced by a droop-correction filter.

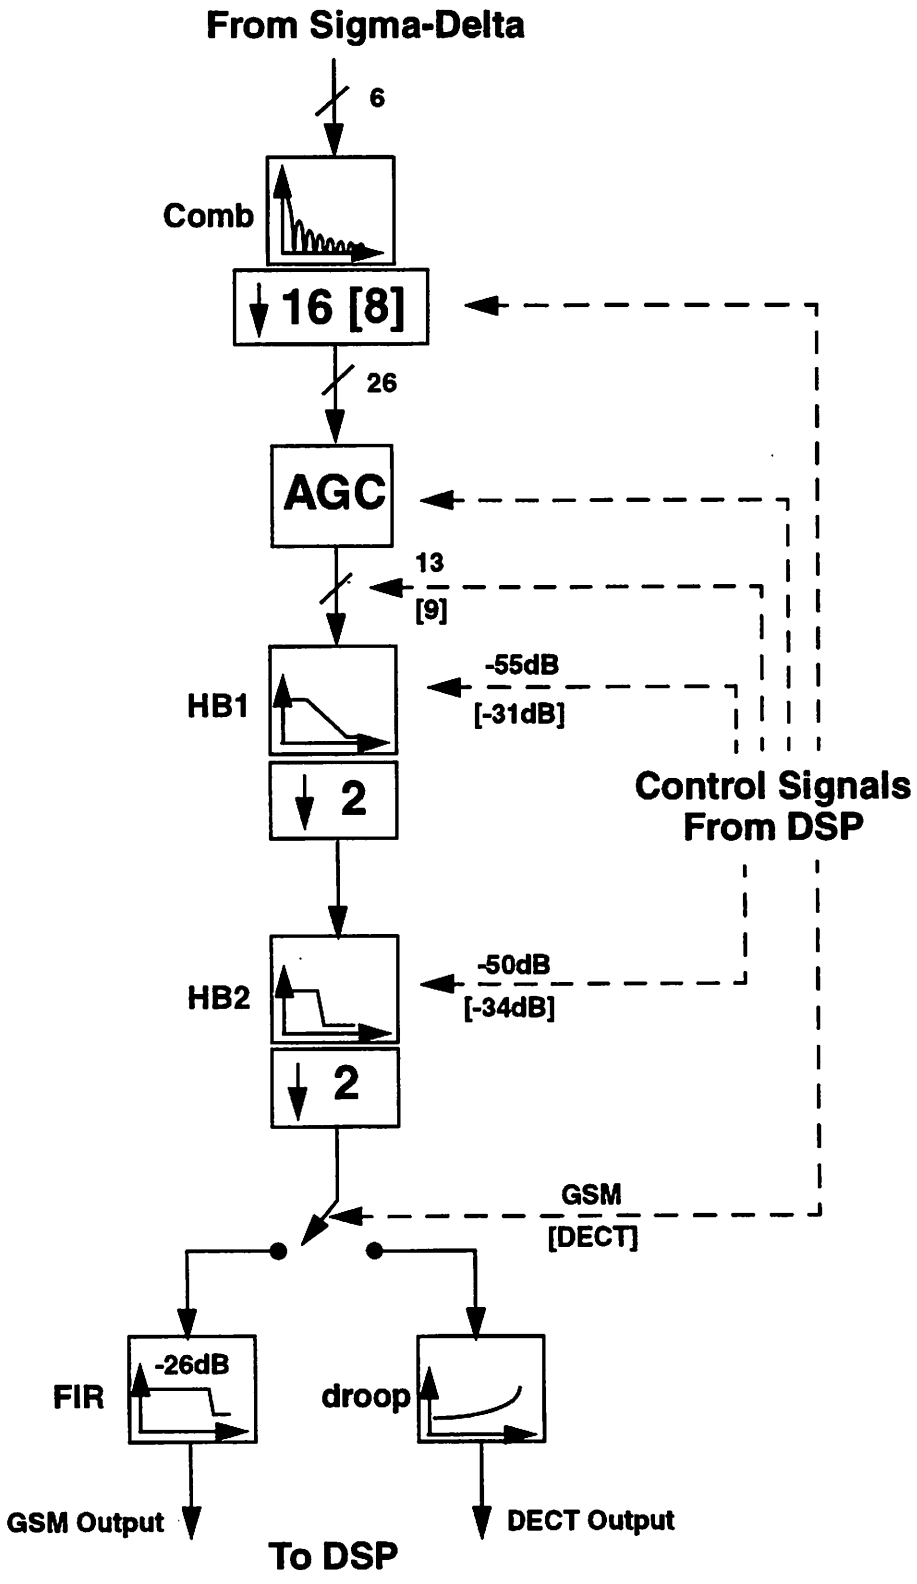

Fig. 3.16: Decimation Filter Architecture for GSM and DECT Standards

Chapter 4

Decimation Filter Implementation

4.1 Introduction

The implementation of a decimation filter should be chosen based on requirements such as power, area, performance, and, depending on the application, multi-standard operation and substrate noise. For example, a filter requiring extremely high performance will be implemented differently than one in which minimizing area is important. This chapter presents two implementations suitable for a multi-standard decimation filter and explains why a fixed multi-standard implementation is used to implement the GSM/DECT decimation filter.

The three general categories for decimation filter implementation are hard wired, reconfigurable, and programmable. [32] A hard wired implementation uses dedicated hardware, and will therefore have lower power and higher performance. In a reconfigurable implementation, the filter coefficients can be changed with the ROM mask. In general, this implementation has a smaller area and higher power than the hard wired implementation. A programmable implementation can be advantageous if a fast enough DSP is already available. The filter in this project is to be implemented with the receiver, so the programmable implementation is not considered. Both the ROM programmable (reconfigurable) and the fixed (hard wired) implementations are considered in this chapter. Several blocks in the receiver

architecture, such as the AGC, comb filter, and clock dividers, are the same for both implementations. They will, therefore, be discussed only once in Section 4.2 where the hard wired implementation is described.

4.2 Fixed Implementation

The fixed implementation uses dedicated hardware and is therefore specific to the GSM and DECT standards. The only programmability included in the filter is that which is necessary to switch between the GSM and DECT modes, and to operate the AGC. In this implementation, each tap has its own adder(s) to implement the cannonic signed digits, and all of the shifts within the individual filters are hard wired. Hardware is re-used between the GSM and DECT modes, but not within or between the filters. So, although the implementation is fixed, *one* filter can effectively implement both standards, keeping the area reasonably small.

4.2.1 Comb Filter

The comb filter is implemented as a cascade of five integrators followed by decimation (performed by the shaded register) and five differentiators as shown in Figure 4.1. It is possible to replace an integrator, a differentiator, and the resampling switch with an accumulate-and-dump circuit as explained in [32]. This replaces the resampling switch and a low-rate adder and register with some simple logic, slightly reducing the power and area of the comb filter. In order to use the accumulate-and-dump circuit, a non 50% duty cycle clock is needed at fs/M_c, the output rate of the comb filter. For simplicity, the accumulate-and-dump circuit will not be used in this implementation.

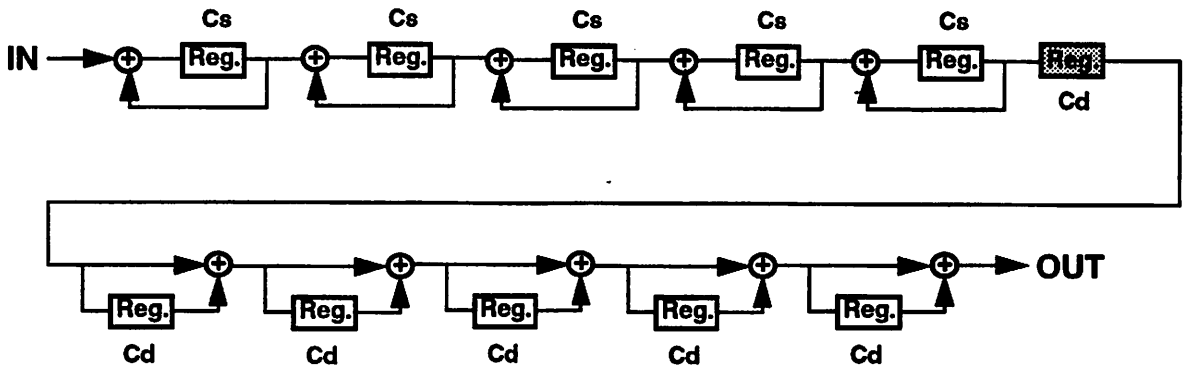

Fig. 4.1: Comb Filter Implementation

The input sampling rates for GSM and DECT are 12.8MHz and 44.8MHz respectively. Since the bandwidth and the sampling rate for DECT are higher, it will determine the worst case timing requirements for the filter. Because the integrators are pipelined, the worst case delay of a 26 bit register and a 26 bit adder must be completed in 22.3ns. Ripple carry adders are extremely slow and the standard cell adders in a 0.5um and a 0.35um process could not meet the DECT speed requirement. For the 0.35um process, a cascade of 2-bit standard cell adders and a D-flip flop were fast enough to meet the specification. If standard cell adders with sufficient speed are not available, a logarithmic adder as described in [23] could be used.

4.2.2 Halfband and FIR Filters

The halfband filters are implemented using the polyphase structure described in Section 3.4. The FIR filters are implemented using the simple direct form structure. For a fixed implementation, the shifters and adders needed to implement the taps are hard wired as shown in Figure 4.2, a diagram of the first halfband filter. CSD tap encoding and nested multiplication can both be seen in the figure. The input rate for the first halfband filter is 800kHz for GSM and 5.6MHz for DECT. The effective filter rates are 400kHz and 2.8MHz respectively, so the filter can easily meet the timing requirements using the same standard cells as the comb filter. The commutator shown in the figure can be implemented using two standard cell D-flip flops, each clocked on every other input sample. For rising edge triggered flip flops, the clock need

only be inverted for one of the latches. As described in Section 3.8, some of the adders in the halfband filters are connected to different power supplies that can be shut down when the filter is operating in DECT mode. This degrades the stopband rejection of the filter to meet the less stringent DECT specification. The adders that are powered down in DECT mode are shaded in Figure 4.2. The dashed arrows show where an adder needs to be bypassed. This can be done by adding a pass gate or a mux, that is set to adder or bypass mode, depending on which standard is being used. Seven adders can be shut down, leaving eight in operation in DECT mode. The over head for this programmability is four muxes and the extra power supply line. The muxes consume very little power because they follow the input and do not switch. Using this technique results in approximately 30% less power in the first halfband and 36% in the second.

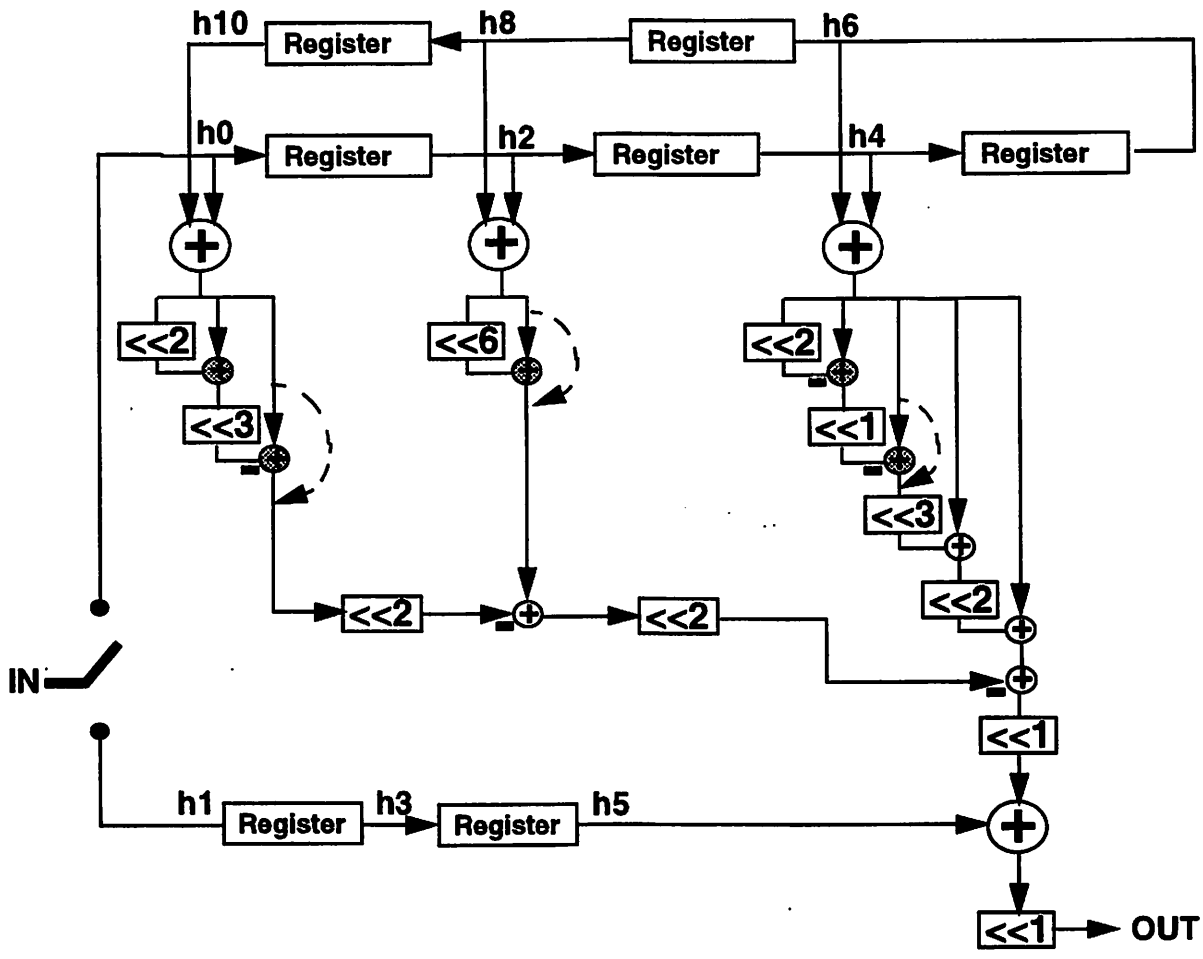

Fig. 4.2: First Halfband Filter Implementation

4.2.3 AGC

The AGC is implemented as a logarithmic shifter as described in [23]. Since this filter is implemented with standard cells, the pass transistors are replaced by an array of muxes. The output word length is 13 and the largest shift value is 7, so three levels of muxes are needed. The shift value, and therefore the four select values, are controlled by the DSP, based on the RSSI as explained in Section 3.6. The AGC structure is shown in Figure 4.3, where several of the interconnections in the third row are left out for clarity. Note that input bits 5 through 0 are not seen on the diagram because they are not found in the output for shifts of 7 or less. Each of the select values controls the entire row of muxes. It can be seen from the diagram that if sel_1 selects the right path, no shift is performed by the first row. If sel_1 selects the left path, a shift of one bit occurs. The second and third rows perform shifts of two, and four bits respectively, or no shift at all. In this way, shifts can be performed using combinations of the three rows to achieve overall shifts with values from zero to 7.

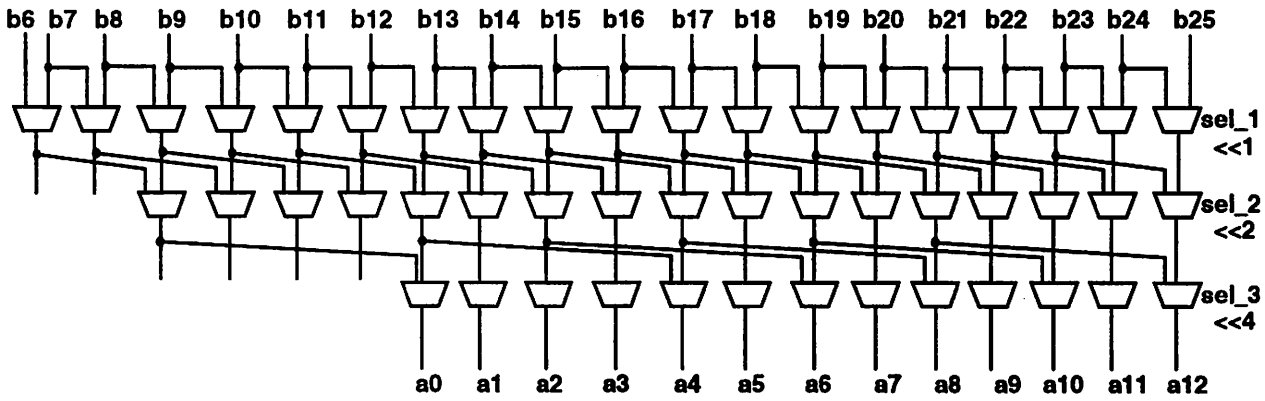

Fig. 4.3: Partial Diagram of Mux-Based AGC

4.2.4 Clock Dividers and Peripheral Circuits

It is assumed that the input clocks will be generated by the frequency synthesizer for the sigma-delta and the decimation filter. So, the remaining 50% duty cycle clocks needed by the decimation filter can be generated using a simple divide by 2 circuits described in [23].

After the comb filter, the clock must be divided by 16 for GSM and 8 for DECT. If four divide by 2 circuits are cascaded, the correct output can be chosen using a mux as shown in Figure 4.4. The select is set by the DSP, and will only switch if the mode of operation changes. The shaded divide by 2 circuit can be powered down in DECT mode to conserve power.

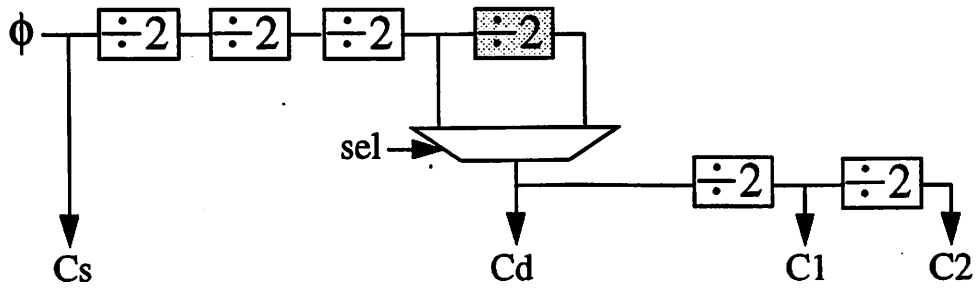

Fig. 4.4: Clock Dividers

4.2.5 Summary of Fixed Multi-Standard Implementation

The fixed implementation uses dedicated hardware for low-power consumption, and good performance. The filter that implements the GSM standard is used as a template for the DECT filter. The filter requirements for GSM are more stringent than those for DECT, so several blocks can be powered down, degrading the filter response and further reducing the power consumption in DECT mode.

4.3 ROM Programmable Implementation

In this implementation, the filter coefficients for the halfband and FIR filters are stored in a read only memory (ROM), and implemented in an arithmetic logic unit (ALU). The output data of each filter is stored in random access memory (RAM), and is read into the ALU when it is needed by the subsequent filter. In addition to specifying the filter coefficients and order of operations, the ROM codes also control the read and write functions of the RAM. With a ROM-based implementation, the ROM can be broken into blocks for each standard. A signal

from the DSP is needed to choose the block corresponding to the mode that the receiver is operating in. With this architecture, new filters can be implemented by simply changing the ROM mask, and ensuring that the counters, clock multipliers, and RAM can accommodate the new filter. A block diagram of the ROM programmable implementation is shown in Figure 4.5.

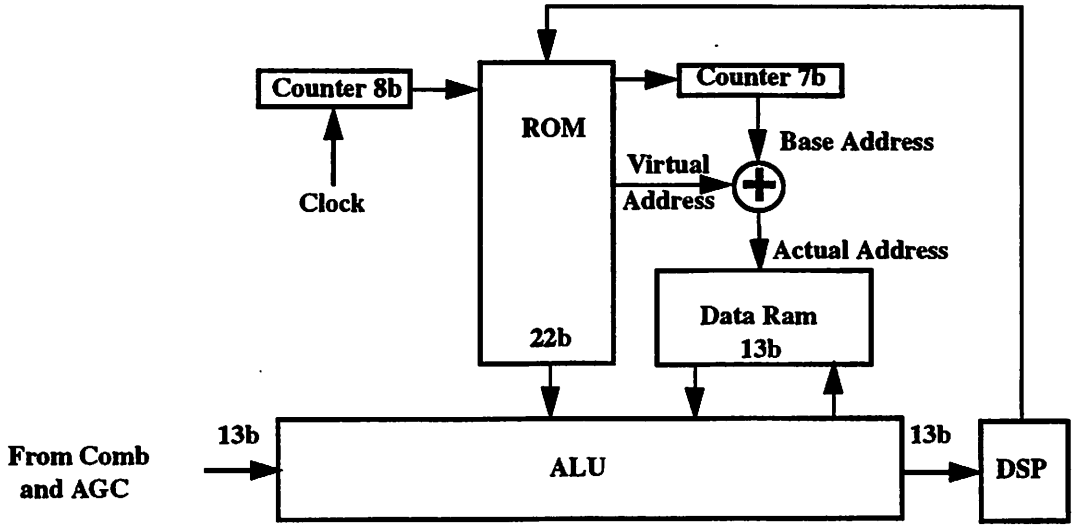

Fig. 4.5: Architecture of ROM Mask-Programmable Implementation [24]

4.3.1 Arithmetic Unit

The ALU performs all of the arithmetic needed to implement the halfband and FIR filters. It stores data in and reads data from the RAM, shifts, negates, adds numbers, and delivers output data to the DSP. A block diagram of the ALU is given in Figure 4.6. [24] In the figure, data_in is the input data from the AGC. When the data is available, it is read into the RAM through mux1. Input data to the filter can be read into reg2 and reg3 through mux2 and mux3 respectively. The contents of reg1 and reg2 can be inverted, shifted, and added or subtracted from the contents of reg3. The result can then be stored in reg1, reg2, or reg3 for further computation, or in reg4, where it is read by the DSP as the filter output. If the result is an intermediate output, for example of the first halfband, it is latched into reg1 and then stored into the RAM through mux1. Bits b0 through b12 in the diagram are control bits from the ROM

that set the ALU to the correct configuration for each operation. The shift requires three control bits in order to perform logarithmic shifts of seven or less.

As an example, consider implementing $h = 2^{-3}(x_1 + 2^{-1} x_2)$, where x_1 and x_2 are data stored in the ROM. First x1 and x2 must be loaded into reg3 and reg2 respectively. x2 is passed from reg2 through mux4 and mux5 where it is shifted by 1 bit. From reg3, x1 passes through mux6 where it is added to the output of the shifter. Now, $(x_1 + 2^{-1} x_2)$ is at the output of the adder. This value can be stored in reg1 and then shifted by 3 on the next cycle. The result, h, is then located at the input to the adder. It can be added to another value, or added to zero and either stored in the RAM or latched into reg4 as an input to the DSP.

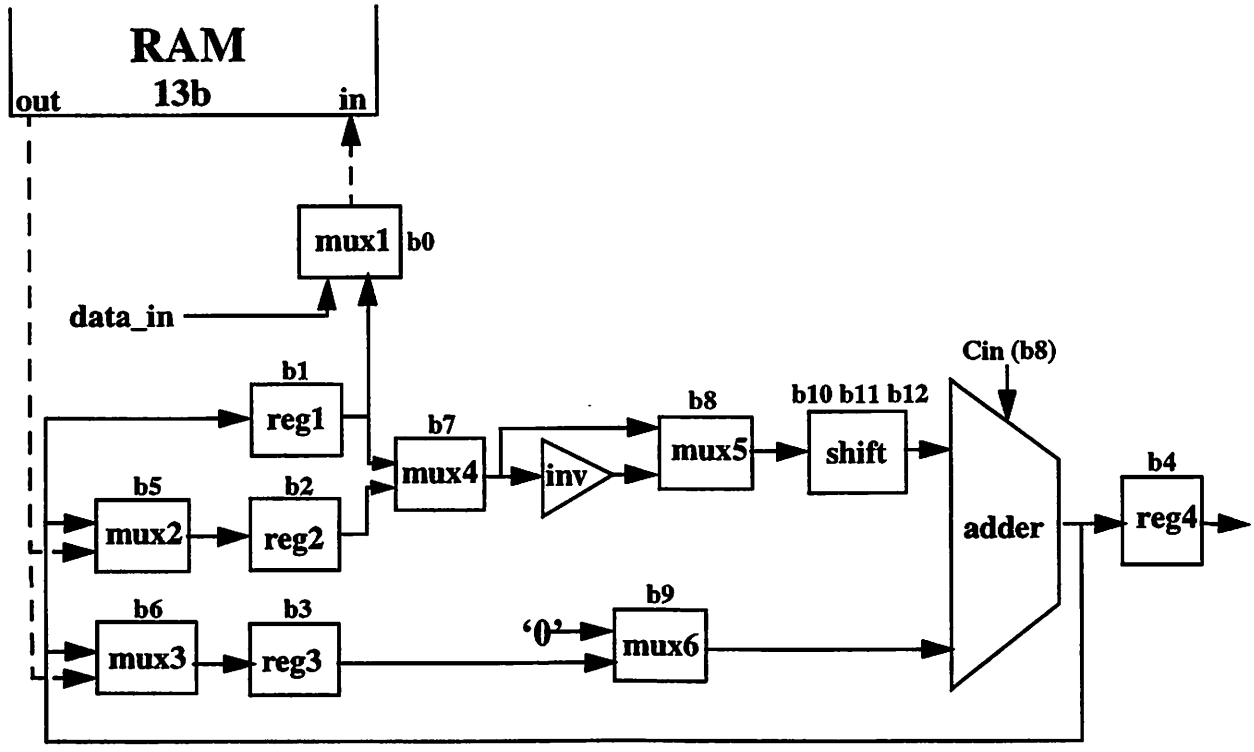

Fig. 4.6: Block Diagram of Arithmetic Unit [24]

4.3.2 ROM Operation

The ROM contains a series of 22b codes that, when accessed cyclically, implement the final stages of the decimation filter. Each cycle results in a single output at the Nyquist rate. In Figure 4.6, it can be seen that bits 0 through 12 are used for controlling the ALU. Bit 13 is used to update the base address by incrementing the 7-bit counter in Figure 4.5. Bit 14 controls the RAM write operation, and bits 15 to 21 are the virtual RAM address. It can be seen from Figure 4.5 that the base address, at the output of the 7-bit counter, is added to the virtual address to form the actual RAM address. In this way, the RAM is written to and accessed in a circular fashion. The RAM is described in more detail in the following section.

The clock that increments the 8-bit wrap around counter in Figure 4.5 is determined by the number of ROM instructions and by the filter output rate. For an output rate of 200kHz, and N ROM codes, the clock must be N times 200kHz, since one complete cycle of the ROM is needed for each filter output. To implement the GSM filter, the ROM must include instructions for the two halfband filters and the FIR filter. These require 19, 86, and 47 instructions respectively. The FIR filter has no decimation, so for each input, one output is produced. Each of the halfband filters decimate by two, so two inputs are needed for one output. Thus, the first halfband requires four inputs for one filter output, and must operate twice during each cycle. The total number of instructions becomes 19(2) + 86 + 47 = 171. This results in a clock rate of 171 times 200kHz for GSM, or 34.2MHz for DECT. To make the clock design easier, the clock should be 34.4MHz, a multiple of 800kHz, the comb filter output rate. It seems easiest to operate the first halfband twice, then the second halfband, and then the FIR as shown in Figure 4.7. But, since an input arrives every 43 instructions, three inputs, numbers 2, 3, and 4, are received in between halfband 1 operations. The halfband only has room for two inputs, so the third input will overwrite relevant data and give an erroneous response.

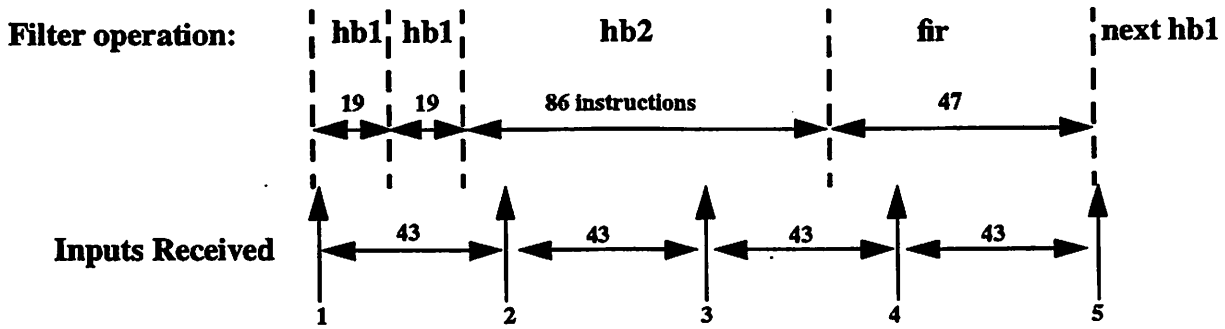

Fig. 4.7: Incorrect ROM Ordering of Filter Operations

Figure 4.8 shows the correct ordering of the filters. In order for each operation of halfband1 to receive the correct input data, the filter operation must be spread out in time and staggered with other filter cycles. Suppose filter cycles x, y, z, and a occur consecutively. At the beginning of Figure 4.8, the first halfband 1 operates. Before the second halfband 1 can operate, two more inputs are needed. The time before the inputs arrive can be used for the operation of the FIR from cycle x, the previous cycle. After halfband 1 operates both times, the inputs to halfband 2 are available. When halfband 2 is finished, two new inputs have arrived from the AGC. To make sure that these inputs are not overwritten, the first operation of halfband 1 for the next cycle, z, occurs. The operation of halfband 1 does not affect the FIR filter at all, so the FIR for cycle y can be implemented after the first operation of halfband 1 for cycle z. A complete filter cycle (cycle z) is shown in italics in Figure 4.8. Note that this is different from a ROM cycle, which is simply the instructions which must be cyclically repeated in order to implement the filter. Appendix 3 gives a list of the ROM instructions needed to implement the filter.

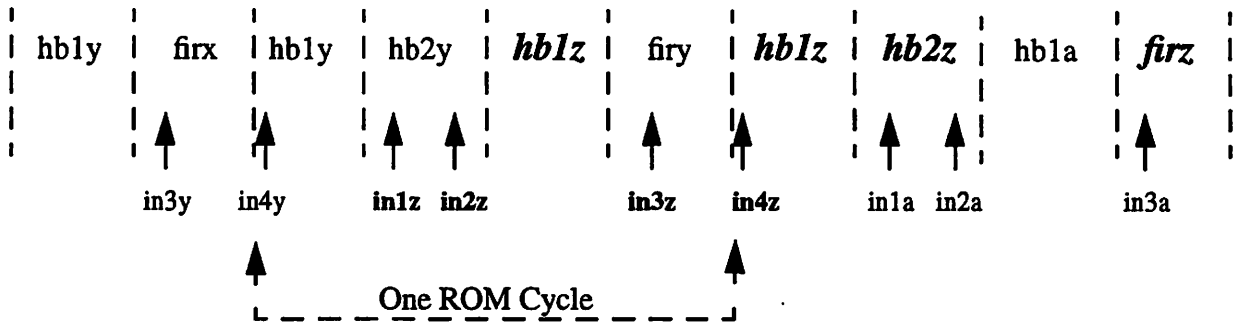

Fig. 4.8: Correct Ordering of ROM Coefficients

4.3.3 RAM Partitioning and Multi-Rate Addressing

The RAM stores all of the relevant data for each of the filters. This data corresponds to what would be stored in registers in a fixed implementation. But, unlike a fixed or direct-form implementation, each data word is stored in the RAM once, and remains in the same location until it is no longer needed. After the data is no longer needed, it is replaced by new data. It is possible to implement a ROM programmable filter using registers instead of a RAM, however, it has been shown in [25] that the power dissipated in the registers is significantly greater than the power dissipated by the RAM. This is because the register based filter must shift data between registers on every clock cycle, instead of simple storing and reading them once.

The RAM in this application is circularly addressed as shown in Figure 4.9. The RAM is divided into blocks for each filter. The block boundaries and the base address are incremented once for each filter output, moving circularly throughout the RAM. As data at the top of each block becomes obsolete, it is replaced by new data from the block above. In Figure 4.5, it can be seen that the actual RAM address is the base address plus the fixed virtual address that is stored in the ROM.

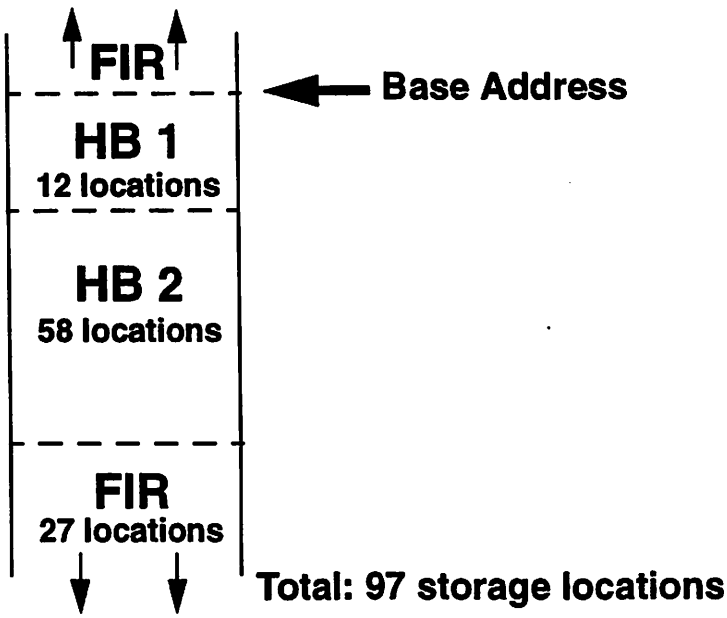

Fig. 4.9: Circularly Addressed RAM

Each block must increment once for each overall filter output, regardless of the filter rate. So, the first halfband should only move one location for every four inputs, and the second halfband moves once for every two inputs. In order to prevent the blocks that run at higher sampling rates from overwriting the blocks that run at slower rates, they must be separated into sub-blocks, with the number of sub-blocks corresponding to the amount of oversampling. The first halfband, for example, runs at four times the output rate. Therefore, it is partitioned into four sub-blocks. The second halfband runs at twice the output rate and the FIR runs at the output rate, so they will have two sub-blocks and one block respectively.

One sub-block increments for each filter input. To prevent data from one sub-block from overwriting current data in another sub-block, the filter block must be partitioned into even and odd sub-blocks. These sub-blocks can be further divided into even and odd blocks until the correct number of sub-blocks are obtained. As an example of this, the first halfband is shown in Figure 4.10 in its direct-form, polyphase implementation. An input data stream of a1, a2, a3... is assumed, and ki represents the filter coefficients.

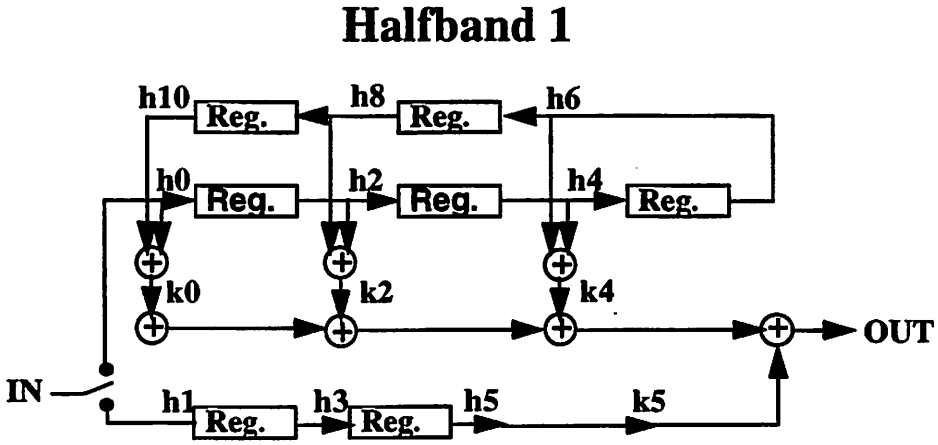

Fig. 4.10: Direct-Form Polyphase Implementation of Halfband 1

Figure 4.11 shows the RAM contents of the first halfband filter for a ROM programmable implementation, corresponding to the direct-form implementation shown in Figure 4.10. The filter has been divided into even and odd blocks 1 and 2. Each of these has been further divided into even and odd blocks, as shown on the left side of the figure. The corresponding register locations are listed to the right. The third column shows the contents of the registers for an input stream of a1, a2, a3, etc. Note that the register location furthest from the input (h14) contains the oldest input datum a1. Suppose that a15 has just been stored in the location corresponding to h0. The filter computes an output, which is stored as an input to the second halfband. The data stored in h14 and h7 are no longer be needed for future computation, so they can be overwritten. It is assumed that the address just below the bottom of the block is also replaceable. So the next inputs a16, and a17, can be placed in these locations. The right most column shows how the boundary has shifted by one address. Now, the datum that corresponded to h4, for example, has "moved" to h8. Another output can be computed for the second halfband, and the data that was previously stored at locations h5 and h12 are now obsolete. The even_1 sub-block has moved down by one address, replacing a10 with a18. The

datum at the top of the block, a3, can be replaced by the an input to the block above. Appendix 3 gives the RAM partitioning in more detail.

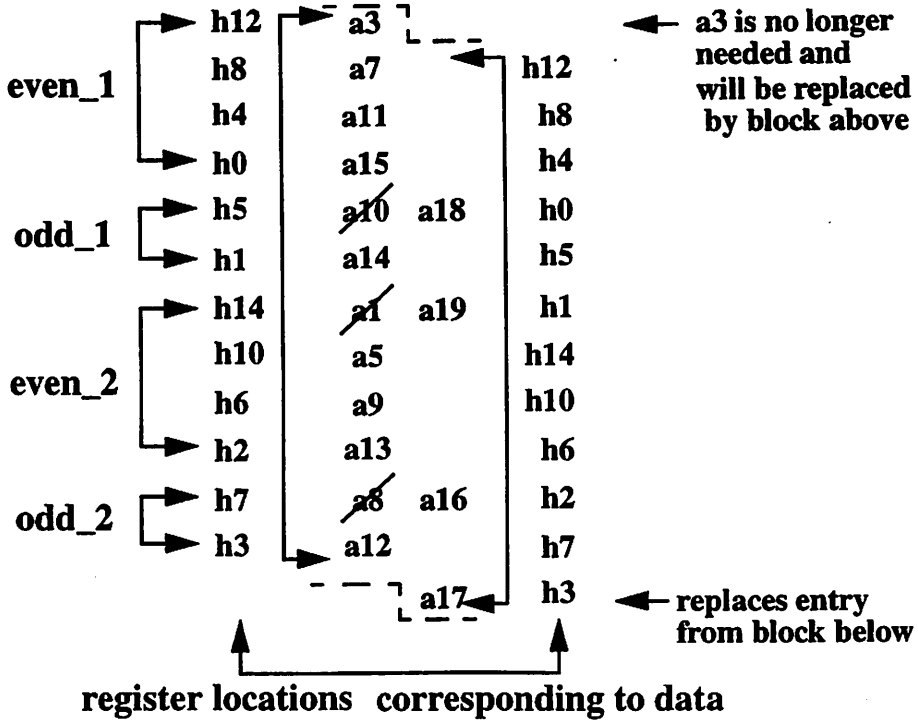

Fig. 4.11: RAM Contents for Halfband 1

4.3.4 Summary of ROM Programmable Implementation

The ROM programmable implementation is mask programmable, and is therefore easy to change or add to. The hardware in the ALU is reused for each tap and each filter. As a result the area is very small, but the clock rate is substantially higher than a fixed implementation. The comb filter and AGC circuits are the same as for the fixed implementation. In the clock generation circuitry, the last two dividers are replaced by a multiplier. The structure of the ALU and the operation of this implementation are based on [20] and [24]. See these references for further information.

4.4 Comparison of Implementations

For this project, the most important factor in determining which implementation to use is the power consumption. Issues of secondary importance are area and adaptability to other standards. The power and area numbers in this section were calculated using the power and area numbers from a 0.35um standard cell library. The power estimates do not take the activity factor into account, and are therefore pessimistic. The library is not designed for low power, so significant power savings could be obtained if the cells were re-designed for low-power and low-voltage. The area estimates ignore the routing overhead since there are five layers of metal and all of the routing can be done on top of the cells. [33]

Table 4.1: Power Comparison of Filter Architectures

Implementation	Power - GSM Mode	Power - DECT Mode	Comparison to This Work
Fixed Multi-Standard (This Work)	4.37 mW	16.4 mW	--
Fixed Multi-Standard Without AGC or Block Power Down	4.77 mW	19.4 mW	Saves 9% - GSM and 19% - DECT
Fixed With All Custom Filters	4.37 mW	16 mW	Negligible savings for DECT
ROM Mask-Programmable	26.3 mW	64.8 mW	6x higher - GSM 4x higher - DECT

Table 4.2: Area Comparison of Filter Architectures

Implementation	Area Estimate	Comparison to this Work
Fixed Multi-Standard (This Work)	1.4 mm^2	--
Fixed With All Custom Filters	1.8 mm^2	30% larger
ROM Mask-Programmable	.58 mm^2	60% smaller

Tables 4.1 and 4.2 give the power and area for the ROM programmable and the fixed architectures. The fixed implementation used in this work is labeled as multi-standard, to differentiate it from an implementation that uses separate filters for each standard (Fixed With All Custom Filters). Using custom filters saves a negligible amount of power, but has 30% more area. If more than two standards need to be accommodated, the area would be even larger. In the second row of Table 4.1, the power for the fixed implementation is compared to the same implementation without any of the power saving techniques, such as AGC and powering down unused blocks. The power savings is 9% and 19% for GSM and DECT respectively.

The most important comparison is between the fixed and ROM programmable implementation. The power of the fixed implementation is six times smaller for GSM and four times smaller for DECT. The area, however, is more than twice that of the ROM programmable implementation. Since power is more important than area in this project, the fixed implementation is chosen. Another factor that may be of importance, depending on the application, is substrate noise. A lower-power solution will have inject less noise into the substrate. Furthermore, in the fixed implementation, all of the clocks are less than or equal to the sigma-delta clock. In the ROM programmable architecture, the clock is greater than the sampling rate of the ADC. As a result, the substrate may not settle before the sigma-delta takes an input sample, which causes further degradation in the dynamic range. [36] [37]

The next chapter will discuss these issues in more detail, and try to address decimation filter design for multi-standard applications in general.

4.5 Simulation Results

Simulation results indicate that the decimation filter will achieve the required CNR for the worst case quantization noise and the worst case blocking profile for both standards. Figure 4.12 and Figure 4.13 show the worst case quantization noise at the output of the comb filter. The noise power within 100kHz is -100.3dB, which is 27.3dB below the minimum input signal.

The AGC applies gain to the desired channel, the noise at the output of the comb filter, and the blockers. Figure 4.14 shows the output spectrum after the remaining filters have been applied. For the minimum desired signal, the in-band noise power is -69.5 dB, which includes the quantization noise and truncation effects. The worst case blocker aliasing can then be added to this number to obtain the total noise power.

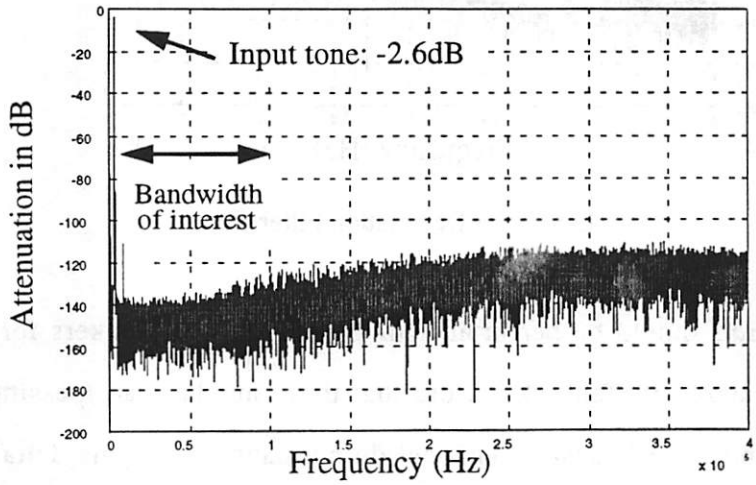

Fig. 4.12: Quantization Noise at Comb Filter Output

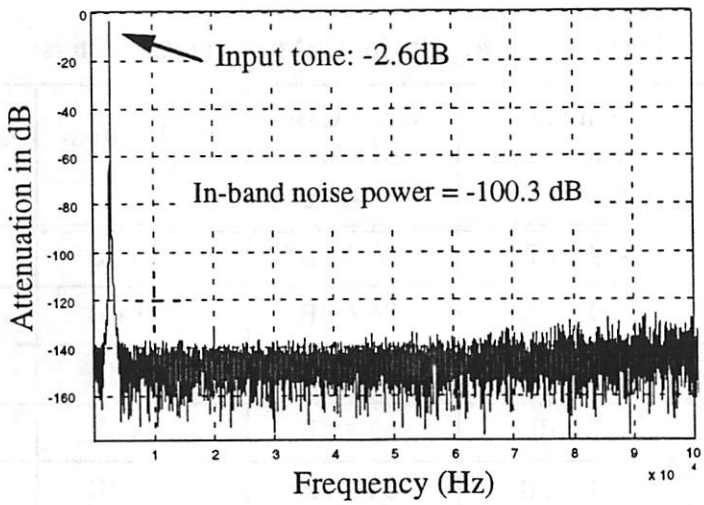

Fig. 4.13: Noise at Comb Filter Output Within Baseband (Zoom of Figure 4.12)

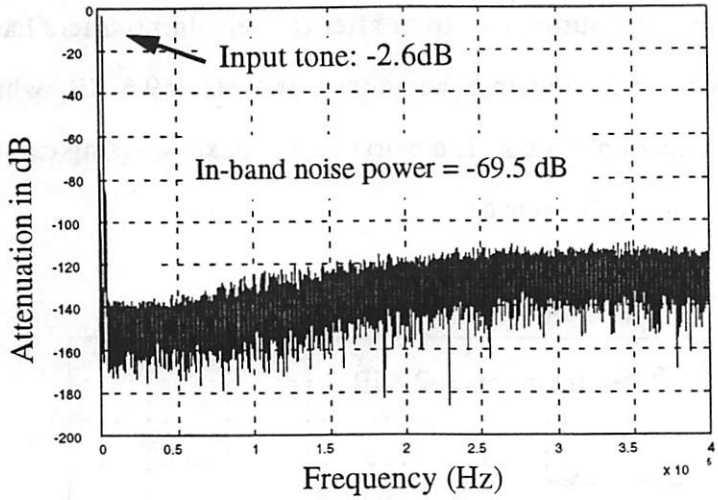

Fig. 4.14: Noise at Decimation Filter Output

This calculation should be performed for the worst case blockers for each standard. The results are summarized in Table 4.3. Note that these numbers are pessimistic, since the worst case blockers and a small desired channel do not cause the sigma-delta to produce the worst case quantization noise. The blocker attenuation for the DECT filters is slightly overdesigned to meet the ripple requirement.

Table 4.3: Performance Results for GSM and DECT Filters

Standard	Blocker	Output Quantization Noise Power	Worst Case Blocker Aliasing	Total Output Noise	Output Noise Budget
GSM	600kHz	-69.5 dB	-60.1 dB	-52.0 dB	-51.9 dB
GSM	1.2MHz	-69.5 dB	-86.6 dB	-56.1 dB	-51.9 dB
GSM	200kHz	-67.7 dB	-80.0 dB	-65.8 dB	-48.2 dB
GSM	400kHz	-67.7 dB	-62.4 dB	-58.6 dB	-48.2 dB
DECT	1.7MHz	-44.4 dB	-33.7 dB	-31.5 dB	-23.1 dB
DECT	3.4MHz	-44.4 dB	-35.2 dB	-32.6 dB	-23.1 dB

4.6 Summary

This chapter discusses two decimation filter implementations for the GSM and DECT standards. The lower-power solution is used to implement the design, and the performance results are given in Table 4.3. The important conclusions are summarized below.

- In the fixed implementation, the filter can be modified to meet each standard by powering down blocks that are not needed. These blocks should be connected to a separate power supply that can be shut down with a control signal from the DSP. This implementation is very power efficient.
- In the ROM programmable implementation, the ROM codes must be carefully ordered to ensure correct operation. The RAM must be partitioned in a way that accommodates multi-rate addressing. This implementation is very area efficient.
- The fixed multi-standard implementation was chosen for this design because power is the most important design constraint. The estimated power consumption for the filter is 4.37mW in GSM mode and 16.4mW for DECT.
- Using AGC and powering down unnecessary circuitry yields a power savings of 9% for GSM and 19% for DECT. The savings for DECT is greater since the filter requirements are relaxed compared to GSM.
- Simulation results indicate that the filter will meet the specifications in GSM and DECT modes for worst case inputs.

The power and area estimates given in this chapter are from standard cells in a 0.35um CMOS process with five layers of metal and a 3.3V power supply. Custom circuit design would reduce the power consumption and area, and enable the filter to run at a higher speeds. In this case, the power supply could be reduced, resulting in further power savings. A 1V power supply, for example, would reduce the power from 4.37mW to 0.4mW in GSM mode, and from 16.4mW to 1.51mW for DECT.

Chapter 5

Multi-Standard Issues

5.1 Introduction

Implementing two or more wireless standards in a personal communications device increases the functionality and flexibility of the device. This work is part of a larger project in which the GSM cellular and DECT cordless standards are implemented together on a single chip. Other multi-standard implementations of interest are to combine cordless and cellular standards for other geographical locations, to combine several cellular standards (for example United States, Europe, and Japan,) or to add GPS to any cellular standard to receive and transmit information involving position. Obviously, separate solutions can be designed for each standard, but that would not be optimum for area, power, or small form factor.

Like the transceiver, the decimation filter can meet the different standards with separate filters or with a single multi-standard filter as illustrated in Figure 5.1. Figure 5.2 shows this for a more general case. If several standards must be implemented, the area in (a) will become very large. This project implements the decimation filter in Figure 5.1 (b). This chapter discusses some of the issues involved in designing and implementing a more general multi-standard filter as in Figure 5.2 (b), and proposes future work in this area.

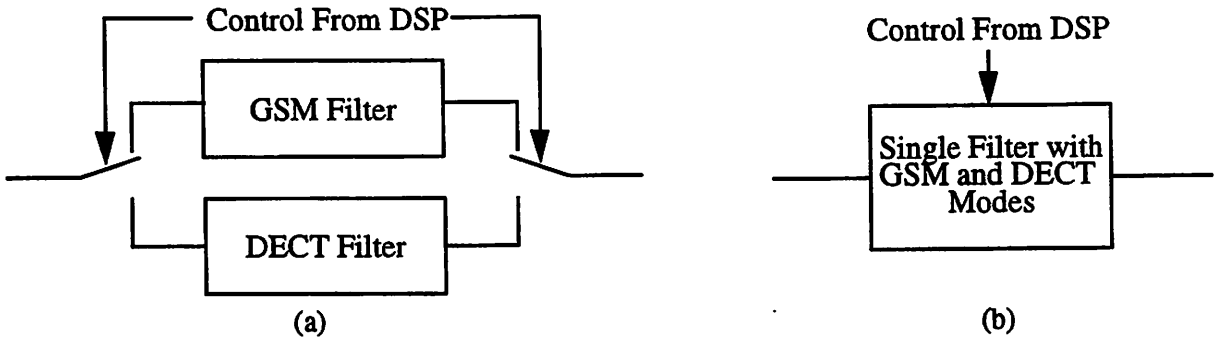

Fig. 5.1: Implementing GSM and DECT Standards With (a) Two Separate Filters or With (b) a Single Dual-Mode Filter

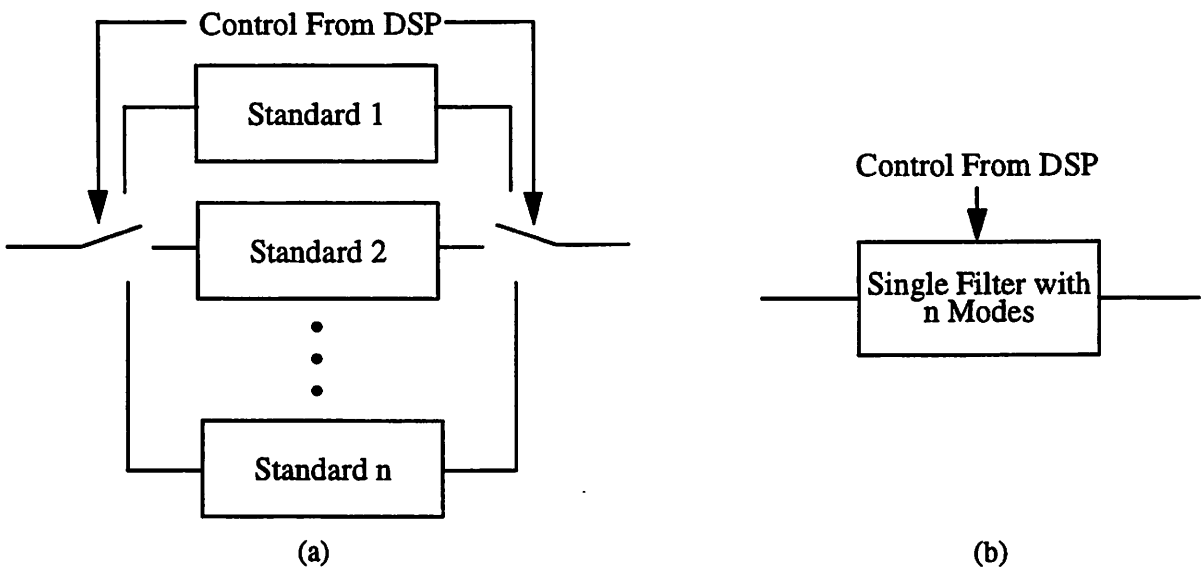

Fig. 5.2: Multi-Standard Decimation Filter Concept (a) Separate Custom Filter for Each Standard (b) Single Filter Capable of Implementing Several Standards

5.2 Choosing an Implementation

Some of the factors that affect the implementation choice of a decimation filter are the modulation schemes and the blocking profiles of the different standards, the number of standards, power, area, bandwidth, sampling rate, decimation ratio, sigma-delta architecture, and sensitivity to substrate noise. Each standard specifies a modulation scheme, a BER,

blocking profiles, and channel bandwidth. The output CNR can be calculated from the BER and modulation scheme. The sigma-delta modulator architecture and oversampling ratio are designed to meet the requirements of the different standards in such a way that the power consumption (or area, or both) is minimized. For more information on designing a sigma-delta modulator see references [11] and [21]. The importance of power, area, and substrate noise depend on the application. For a battery powered application, such as a cellular phone, the power consumption is extremely important and the area must be reasonably small. In other applications, however, area might be much more important. If the decimation filter is integrated on the same chip as the modulator or other receiver blocks, the substrate noise becomes important.

Given the specification for each standard and corresponding sigma-delta modulator, a decimation filter with the lowest power (or lowest area) can be designed for each standard as in Chapter 3. The similarities and differences between the resulting filters, along with the power and area requirements, will help determine which implementation to use. For this work, the GSM filter could be simplified to meet the DECT requirements. Since the FIR was not needed for DECT, a small droop-correct filter was added. Both standards used a fourth order sigma-delta, so the same comb filter could be used for both. As long as the area was reasonable (under $2mm^2$) the goal of this work was to minimize the power consumption. Thus, the fixed implementation is a better choice than the ROM programmable. If the area had been more important than power, however, the ROM programmable implementation would have been more appropriate.

To design a decimation filter using a fixed implementation, the filter should be designed for the worst case stopband rejection and ripple requirements. Then, unnecessary blocks can be powered down in each mode to save power. This solution is best if the power consumption is more important than area, and if the requirements of the different standards are similar, or are relaxed versions of each other. This implementation is also attractive for

applications where substrate noise is important, since the power consumption is low and all of the clocks operate at or below the sampling frequency.

It is possible that the filters following the comb filter are so different that very little hardware sharing is possible in a fixed implementation. If the resulting area is too large, or if the area is the most important, a ROM programmable implementation should be used. With this implementation, it is easy to change or add standards by altering the ROM mask, assuming that the RAM and the programmable shifter are big enough. The wrap around counters and the AGC may also need to be modified. It is possible to further reduce the area of the filter by adding the comb filter instructions to the ROM or by replacing it with other halfband filters. This may, however, result in a clock speed that is impossible to meet using current technologies. The clock for this implementation is proportional to the Nyquist rate of the channel and the number of instructions needed in the ROM. So, for a wide-bandwidth standard, or for one needing a large stop band rejection, the clock may be very high. It is important to estimate the clock rate early in the design to ensure that the speed requirements can be met. This is especially critical if the ability to design using standard cells is important. If the ROM programmable implementation is applied to this application, the DECT timing requirements cannot be met by the standard cells in a 0.35um technology even if the counter, ROM, and ALU are pipelined. A preliminary design of a logarithmic adder showed that the ALU would be able to meet the timing requirements with custom circuit design.

The basic trade-off between the two solutions is power vs. area. The difference in power consumption between a fixed and ROM programmable solution will probably decrease as the number of standards increases. This is because the fixed implementation has added complexity with each new standard. The separate supply lines and power down circuitry increase the total capacitance and therefore the power consumption. There also may be a need to switch between two filters, like this implementation where the FIR or the droop correct filter was selected depending on the mode.(GSM/DECT)

5.3 Future Work

The filter presented in this work only addresses two telecommunications standards and two possible implementations. Other issues to consider are implementing standards with different modulation schemes, implementing a larger number of standards, and implementing standards where it is optimal to use different modulator architectures. Some ideas to consider are variable order comb filters, an architecture that uses a cascade of filters that can be switched in or out to program the amount of rejection, and applications that do not require linear phase. Design of multi-standard tranceivers is a relatively new area, so the possible research issues to be explored are numerous.

Chapter 6

Conclusions

This thesis presents the design of a low-power decimation filter that performs channel select filtering for a highly integrated multi-standard receiver. The more important research results are summarized below.

- Designed a decimation filter that meets the GSM and DECT specifications without significantly degrading the receiver noise figure. Estimated power consumption is 4.4mW and 16.4mW for GSM and DECT modes respectively.
- Showed that automatic gain control can be used within a decimation filter that performs channel select filtering to reduce power consumption and area.
- Developed a power saving technique for a fixed multi-standard implementation, where filter blocks and sub-blocks can be powered down, enabling the filter to adapt to each standard with reduced power consumption.

The design of a multi-standard decimation filter depends heavily on the standards that are implemented. The two implementations presented in Chapter 4, and the multi-standard issues discussed in Chapter 5, are only a first step in understanding this complex area.

Appendix 1

Noise Figure and Decimation Filter CNR Calculations

This appendix shows how the receiver noise figure, the decimation filter CNR, and the allowable decimation filter noise are calculated.

Receiver Noise Figure

The thermal noise power available from the source under a matched condition is kTBw, where Bw is the bandwidth of interest.[15] This results in the following equation for the input noise floor.

$$kTBw = -173.8 + 10\log Bw \quad \text{(dBm)} \qquad \text{(Eq 1)}$$

Which yeilds a noise floor of -120.8 dBm for GSM and -112.3 dBm for DECT. The sensitivity, or minimum input signal, is given as -102 dBm for GSM and -83 dBm for DECT. From this, the input signal to noise ratio (SNR) can be calculated as

$$SNR_{in} = \frac{S_{in}}{N_{in}} \qquad \text{(Eq 2)}$$

where S_{in} is the sensitivity and N_{in} is the input noise floor. This equation gives 18.79 dB of input SNR for GSM and 29.3 dB for DECT.

From the required bit error rate (BER) given in the specification and given the modulation scheme used, the output carrier to noise ratio (CNR) can be calculated to be 9dB for GSM and 10.3 dB for DECT. [15] The noise figure (NF) can then be calculated as

$$NF = 20\log\left(\frac{SNR_{in}}{CNR_{out}}\right) \qquad \text{(Eq 3)}$$

where the noise figure is in dB. The noise figure is 9.79dB for GSM and 19dB for DECT.

Decimation Filter Noise Figure Relative to Receiver NF

The decimation filter should degrade the overall noise figure by no more than 0.01dB. For two cascaded blocks as shown in Fig. 1, the overall noise figure can be written an in (Eq 4) as a function of the individual noise figures of the two blocks, F1 and F2, and the gain of the receiver G1. [15]

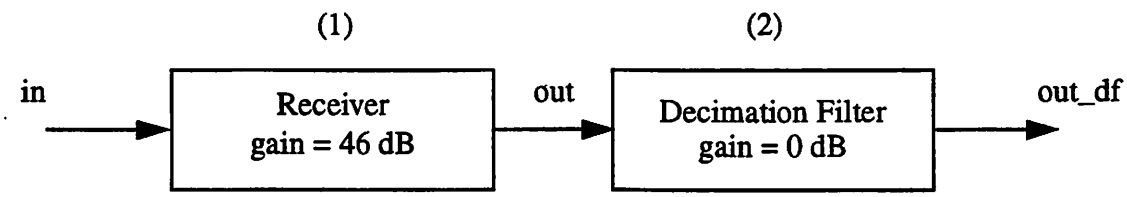

Fig. 1: Receiver and Decimation Filter

$$F = F1 + \frac{(F2 - 1)}{G1} \qquad \text{(Eq 4)}$$

The noise figure specification for GSM is 9.79 dB. If the decimation filter can add .01dB to this value, then F=9.8dB and F1=9.79dB in (Eq 4). For G1=46dB, F2 is 4.26. Similarly, for DECT, F1=19dB and F=19.01dB. This yields an F2 of 9.68dB.

Relative Decimation Filter Noise Contribution

From the decimation filter noise figure, the allowable noise contribution can be calculated. Using (Eq 3) with NF = F2 and the SNR_{in} equal to the output SNR of the receiver, the noise figure can be rewritten as

$$F2 = \frac{Si \times No}{So \times Ni} \qquad \text{(Eq 5)}$$

where Si is the input signal, So is the output signal, Ni is the input noise, and No is the output noise. Si is equal to the sensitivity multiplied by the receiver gain. Ni can be found using (Eq 2) since the CNR at the output of the receiver is known. The decimation filter does not have any gain, so So is equal to Si. These cancell, resulting in

$$F2 = (Ndf + Ni)/(Ni) \qquad \text{(Eq 6)}$$

where the output noise is equal to the input noise added to the decimation filter noise contribution, Ndf. Since F2 and Ni are known, Ndf can be calculated using (Eq 6). The difference between the signal and Ndf is the relative noise contrbution, and is used to determine the filter noise contribution relative to full scale. A 6dB margin is added to this value, to ensure that the specification is met. The important parameters related to the noise figure and decimation filter noise contribuiton are summarized in Table 1.

Table 1:

	GSM	DECT
Input Noise Floor	-120.8 dBm	-112.3 dBm
Reciever Sensitivity	-102 dBm	-83 dBm
Reveiver Output CNR	9 dB	10.3 dB
Receiver Noise Figure (F1)	9.79 dB	19 dB

Table 1:

	GSM	DECT
Decimation Filter Noise Figure (F2)	4.26 dB	9.68 dB
Receiver Minimum Output Signal (Si)	-56 dB	-37 dB
Receiver Output Noise (Ni)	-65 dB	-47 dB
Decimation Filter Noise (Ndf)	-67.98 dB	-41.17 dB
Relative Noise Contribuiton	12 dB	4.1 dB
Filter Noise Contribution with 6dB Margin	18 dB	10.1 dB
Sensitivity Relative to Full Scale	-76 dB	-43 dB
Minimum Signal for Blocking Test	-73 dB	-43 dB
Total Filter Noise Contribution Relative to Full Scale	-94 dB	-53.1
Output CNR Relative to Minimum Signal (Blocking Test)	21 dB	10.2 dB

Appendix 2

DSP Controls

This appendix gives all of the control signals from the DSP for the decimation filter.

Automatic Gain Control

By providing gain when the desired signal is small, the automatic gain control reduces the dynamic range, and therefore the wordlength, of the filters following the comb filter. Because the blocking specifications are different for GSM and DECT, the amount of gain provided for a given desired signal strength (RSSI) will be different for the two standards.

The AGC is provided after the comb filter, while the signal contains the desired channel as well as blockers and noise. The amount of gain that can be applied is based on the RSSI, which is calculated in the DSP. For a given RSSI, the worst case blocking profile at the AGC can be determined. The gain applied must not raise the largest possible blocker above 3dB below full scale. The GSM specifications contain separate profiles for the blockers and the adjacent channel interferers as explained in Section

2.4. The largest blocker is 23.7 dB above the desired signal in the blocking profile, and 30.5 dB above for the adjacent channel interferers. Since the minimum deisred signals are -73dB and -56dB respectively for the two profiles, it can be assumed that for a RSSI below -56dB, the worst case signal is 23.7dB larger, and for an RSSI above -56dB, it is 30.5dB larger.

Table 2 gives the gain, the number of bit shifts, and the DSP control values for the logarithmic shifter given in Section 4.2.3, for all possible RSSI values of the GSM signal. Table 3 contains the same information for the DECT specification, where the largest undesired signal is 6.2dB above the desired.

Table 2: AGC for GSM

RSSI (dB)	Gain (dB)	Bit Shift	DSP Control
below -68.8	42.1	7	111
-69.8 to -62.8	36.1	6	011
-62.8 to -56.8	30.1	5	101
-56.8 to -51.5	18	3	110
-51.5 to -45.5	12	2	010
-45.5 to -39.5	6	1	100
above -39.5	0	0	000

Table 3: AGC for DECT

RSSI (dB)	Gain (dB)	Bit Shift	DSP Control
below -39.3	30.1	5	101
-39.3 to -33.2	24	4	001
-33.2 to -27.2	18	3	110

Table 3: AGC for DECT

RSSI (dB)	Gain (dB)	Bit Shift	DSP Control
-27.2 to -21.2	12	2	010
-21.2 to -15.2	6	1	100
above -15.2	0	0	000

Decimation After Comb Filter

Because the oversampling ratios for GSM and DECT are different, the amount of decimation after the comb filter is also different for the two modes. The decimation ratio of the comb filter, 16 for GSM and 8 for DECT, is determined by the input and resampling clocks, or by the clock divider. Fig. 4.4 shows a clock divider circuit that provides division by 16 or 8 for the comb filter. The DSP must provide a one bit signal for the mux that chooses between these two clocks.

Powering Down Blocks in DECT Mode

When the filter is switched from GSM to DECT mode, much of the hardware becomes unnecessary. The blocks that are unused can be powered down to reduce the power consumed in the DECT mode. This is especially important because the Nyquist rate for DECT is much higher than for GSM. Following is a list of blocks that can be shut down in DECT mode:

- The extra divide by two that is used to create the comb filter resampling clock for GSM.
- The FIR filter. (Note that the droop correct filter can be shut down in GSM mode.)
- Several adders in the halfband filters. (Some of the adders that are shut down must be bypassed using a transmision gate or a buffer. These can be shut down in GSM mode.)

- Four of the least significant bits (LSBs) in the halfband filters are not needed since the dynamic range after the AGC is lower for DECT. These adders and regesters can be powered down, leaving a wordlength of nine.

Each block should be connected to one of three power supply lines. One supply is connected to the actual power supply and these blocks are always on. The other two lines are for those blocks that must be powered down in GSM or DECT modes. They are derrived from the main supply, but contain control curcuitry to shut them off and turn them back on. The signal for this is a single bit, since there are two modes, GSM and DECT.

The extra power supply lines create extra capacitence to be charged. But, this capacitence is rarely switched, so the impact on the power is negligible. The process used for this chip has five metal layers, so the impact on area is small.

Summary of DSP Signals

Four signals are needed from the DSP, three to control the AGC and one to set the filter to GSM or DECT mode. The bit that determines the mode is used to power down unused blocks and set the correct comb filter decimation ratio.

Appendix 3

Filter Design and Implementation

This appendix gives a list of the filter taps with the corresponding CSD's, ROM codes for the first halfband filter, and the RAM allocation.

Fixed Implementation

Table 1: Halfband 1 - GSM

Coefficients	Ideal Tap Value	Actual Tap Value	Cannonic Signed Digits
h(0), h(10)	0.0130492	0.01318359375	2-6 + -2-9 + -2-11
h(2), h(8)	-0.0638715	-0.0634765625	-2-4 + -2-10
h(4), h(6)	0.3016129	0.3017578125	2-2 + 2-4 + -2-7 + -2-8 + 2-10
h(5)	0.50	0.5	2-1

Table 2: Halfband 2 - GSM

Coefficients	Ideal Tap Value	Actual Tap Value	Cannonic Signed Digits
h(0), h(74)	0.0022057	0.002197265625	2-9 + 2-12
h(2), h(72)	-0.0015871	-0.0015869140625	-2-9 + 2-12 + 2-13
h(4), h(70)	0.0021457	0.002197265625	2-9 + 2-12
h(6), h(68)	-0.0028234	-0.0028076171875	-2-9 + -2-10 + 2-13
h(8), h(66)	0.0036384	0.003662109375	2-8 + -2-12
h(10), h(64)	-0.0046142	-0.004638671875	-2-8 + -2-11 + -2-12
h(12), h(62)	0.0057778	0.00579833984375	2-8 + 2-9 + -2-14
h(14), h(60)	-0.0071624	-0.0072021484375	-2-7 + 2-11 + 2-13
h(16), h(58)	0.0088168	0.0087890625	2-7 + 2-10
h(18), h(56)	-0.0108092	-0.0107421875	-2-7 + -2-8 + 2-10
h(20), h(54)	0.0132286	0.01318359375	2-6 + -2-9 + -2-11
h(22), h(52)	-0.0162307	-0.0162353515625	-2-6 + -2-11 + -2-13
h(24), h(50)	0.0200490	0.02001953125	2-6 + 2-8 + 2-11
h(26), h(48)	-0.0251007	-0.025146484375	-2-5 + 2-7 + -2-9 + 2-12
h(28), h(46)	0.0321715	0.0322265625	2-5 + 2-10
h(30), h(44)	-0.0429509	-0.04296875	-2-5 + -2-7 + -2-8
h(32), h(42)	0.0618411	0.0618896484375	2-4 + -2-11 + -2-13
h(34), h(40)	-0.1050028	-0.10498046875	-2-3 + 2-6 + 2-8 + 2-11
h(36), h(38)	0.3179423	0.31787109375	2-2 + 2-4 + 2-8 + 2-9 + -2-11
h(37)	0.5	0.5	2-1

Table 3: FIR - GSM Only

h(0), h(27)	0.0070553	0.0078125	2-7

Table 3: FIR - GSM Only

h(1), h(26)	-0.0211789	-0.01953125	-2-6 + -2-8
h(2), h(25)	0.0127450	0.015625	2-6
h(3), h(24)	-0.0222934	-0.0234375	-2-6 + -2-7
h(4), h(23)	0.0256602	0.02734375	2-5 + -2-8
h(5), h(22)	-0.0274444	-0.03125	-2-5
h(6), h(21)	0.0251610	0.0234375	2-5 + -2-7
h(7), h(20)	-0.0184301	-0.015625	-2-6
h(8), h(19)	0.0052606	0.00390625	2-8
h(9), h(18)	0.0156721	0.015625	2-6
h(10), h(17)	-0.0487916	-0.046875	-2-4 + 2-6
h(11), h(16)	0.1021211	0.1015625	2-3 + -2-6 + -2-7
h(12), h(15)	-0.2152712	-0.21875	-2-2 + 2-5
h(13), h(14)	0.6480249	0.6484375	2-1 + 2-3 + 2-6 + 2-7

The next filters are based on the filters above, with several of the blocks turned off to save power in DECT mode. Note that the ideal coefficients are the same, but the number of CSD's is smaller. The FIR filter is not needed, and is replaced by a simpler droop-correct filter.

Table 4: Halfband 1 - DECT Modification

Coefficients	Ideal Tap Value	Actual Tap Value	Cannonic Signed Digits
h(0), h(10)	0.0130492	0.015625	2-6
h(2), h(8)	-0.0638715	-0.0625	-2-4
h(4), h(6)	0.3016129	0.3046875	2-2 + 2-4 + -2-7

Table 4: Halfband 1 - DECT Modification

Coefficients	Ideal Tap Value	Actual Tap Value	Cannonic Signed Digits
h(5)	0.5	0.5	2-1

Table 5: Halfband 2 - DECT Modification

Coefficients	Ideal Tap Value	Actual Tap Value	Cannonic Signed Digits
h(0), h(74)	0.0022057	0.00	0.00
h(2), h(72)	-0.0015871	0.00	0.00
h(4), h(70)	0.0021457	0.00	0.00
h(6), h(68)	-0.0028234	0.00	0.00
h(8), h(66)	0.0036384	0.00390625	2-8
h(10), h(64)	-0.0046142	-0.00390625	-2-8
h(12), h(62)	0.0057778	0.00390625	2-8
h(14), h(60)	-0.0071624	-0.0078125	-2-7
h(16), h(58)	0.0088168	0.0078125	2-7
h(18), h(56)	-0.0108092	-0.0078125	-2-7
h(20), h(54)	0.0132286	0.015625	2-6
h(22), h(52)	-0.0162307	-0.015625	-2-6
h(24), h(50)	0.0200490	0.01953125	2-6 + 2-8
h(26), h(48)	-0.0251007	-0.0234375	-2-5 + 2-7
h(28), h(46)	0.0321715	0.03125	2-5
h(30), h(44)	-0.0429509	-0.04296875	-2-5 + -2-7 + -2-8
h(32), h(42)	0.0618411	0.0625	2-4
h(34), h(40)	-0.1050028	-0.10546875	-2-3 + 2-6 + 2-8
h(36), h(38)	0.3179423	0.31640625	2-2 + 2-4 + 2-8
h(37)	0.5	0.5	2-1

Table 6: Droop-Correct - DECT Only

Coefficients	Ideal Tap Value	Actual Tap Value	Cannonic Signed Digits
h(0), h(4)	0.00110113404	0.011718750	2-7 + 2-8
h(1), h(3)	-0.0277349943	-0.027343750	-2-5 + 2-8
h(2)	1	1	2-0

ROM Programmable Implementation

The next filters are optimized to the DECT specification and are not based on the GSM filters. These should be used in a ROM programmable implementaton, followed by the droop-correct filter given above.

Table 7: Halfband 1 - DECT Custom Filter

Coefficients	Ideal Tap Value	Actual Tap Value	Cannonic Signed Digits
h(0), h(6)	-0.0506242	-0.046875	-2-4 + 2-6
h(2), h(4)	0.2950593	0.2968750	2-2 + 2-5 + 2-6
h(3)	0.5	0.5	2-1

Table 8: Halfband 2 - DECT Custom Filter

Coefficients	Ideal Tap Values	Actual Tap Values	Cannonic Signed Digits
h(0), h(50)	0.0092077	0.0087890625	2-7 + 2-10
h(2), h(48)	-0.0055188	-0.005859375	-2-8 + -2-9
h(4), h(46)	0.0071845	0.0078125	2-7
h(6), h(44)	-0.0092141	-0.0087890625	-2-7 + -2-10

Table 8: Halfband 2 - DECT Custom Filter

Coefficients	Ideal Tap Values	Actual Tap Values	Cannonic Signed Digits
h(8), h(42)	0.0117121	0.011718750	2-7 + 2-8
h(10), h(40)	-0.0148136	-0.015625	2-6
h(12), h(38)	0.0187580	0.01953125	2-6 + 2-8
h(14), h(36)	-0.0239604	-0.0234375	-2-5 + 2-7
h(16), h(34)	0.0312072	0.031250	2-5
h(18), h(32)	-0.0421789	-0.042968750	-2-5 + -2-7 + -2-8
h(20), h(30)	0.0612676	0.0615234375	2-4 + -2-10
h(22), h(28)	-0.1046595	-0.10546875	-2-3 + 2-6 + 2-8
h(24), h(26)	0.3178218	0.3173828125	2-2 + 2-4 + 2-8 + 2-10
h(25)	0.5	0.5	2-1

RAM Allocation

For the GSM specification, the RAM should be allocated as in Fig. 6.1 for the multi-rate addressing to function properly. Each filter is partitioned into N sub-blocks, where N is the oversampling ratio at the filter input. The sub-blocks are partitioned into even and odd blocks, and ordered so that unnecessary data is quickly replaced by new data. This is explained in greater detail in Section 4.3.3. and is also described in [20].

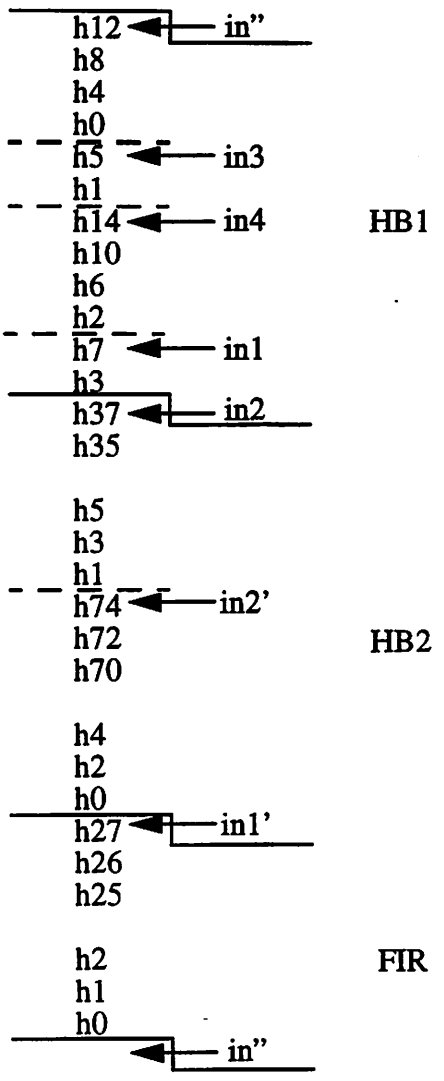

Fig. 6.1: Allocation of RAM for GSM Specification

ROM Instructions

The 22 bit instrucitons coded into the ROM control the filter operation. An example is given in Fig. 6.2 for the first halfband filter. These instructions can be translated to binary code using the RAM partitioning to find the addresses, and the ALU diagram for the first 13 bits.

```
(previous operation)  load h0 to R2
                      load h10 to R3
                  1)  store new data        //hb1 input
                  2)  R3=R2=R2+R3
                  3)  R2=R3+R2>>2
                      load h2 to R3
                  4)  R1=R3-R2>>3           //x0
                      load h2 to R3
                  5)  R3=R2=R2+R3
                  6)  R3=R3+R2>>6           //x2
                      load h4 to R2
                  7)  R1=R3-R1>>2
                      load h6 to R3
                  8)  R3=R2=R2+R3
                  9)  R2=R3-R2>>2
                 10)  R2=R3-R2>>1
                 11)  R2=R3+R2>>3
                 12)  R3=R3+R2>>2           //x4
                 13)  R1=R3-R1>>2
                      load h5 to R3
                 14)  R1=R3+R1>>1
                 15)  R1=R2>>1
                 16)  store R1 in RAM       //hb2 input
```

Fig. 6.2: ROM Instrucitons for Halfband 1 (Second Operation)

The ROM instructions for the individual filters must be carefully ordered as explained in Section 4.3.2 to ensure correct operation.

References

[1] Paul Gray, and Robert Meyer. "Future Directions in Silicon ICs for RF Personal Communications," *Proceedings, 1995 Custom Integrated Circuits Conference*, pp. 83-90, May 1995.

[2] Theodore Rappaport,*Wireless Communications Principles and Practices*, IEEE Press, 1996.

[3] Trudy Stetzler, et al. "A 2.7V to 4.5V Single-Chip GSM Transceiver RF Integrated Circuit," *Digest of Technical Papers, International Solid-State Circuits Conference*, pp. 150-151, February 1995.

[4] Chris Marshall, et al. "A2.7V GSM Transceiver IC with On-Chip Filtering," *Digest of Technical Papers, International Solid-State Circuits Conference*, pp. 148-149, February 1995.

[5] Hisayasu Sato, et al. "A 1.9GHz Single-Chip IF Transceiver for Digital Cordless Phones," *Digest of Technical Papers, International Solid-State Circuits Conference*, pp. 342-343, February 1996.

[6] K. Irie et al. "A 2.7V GSM RF Transceiver IC," *1997 IEEE International Solids-State Circuits Conference. Digest of Technical Papers*, pp. 302-303, February 1997.

[7] T. Weigandt, *"Low-Phase-Noise Low-Timing-Jitter Design Techniques for Delay Cell Based VCOs and Frequency Synthesizers,"* Ph.D. Thesis, Electronics Research Laboratory, University of California at Berkeley, December 1997.

[8] Christopher Hull, et al. "A Direct-Conversion Receiver for 900 MHz (ISM Band) Spread-Sectrum Digital Cordless Telephone," *Digest of Technical Papers, International Solid-State Circuits Conference*, pp344-345, February 1996.

[9] Jan Sevenhans, et al. "An Integrated Si Bipolar RF Transceiver for a Zero IF 900 MHz GSM Digital Movile Radio Front End of a Hand Portable Phone," *Proceedings of the Fifth Annual IEEE International ASIC Conference and Exhibit*, pp. 561-564, September 1992.

[10] J. Crols and M. Steyaert, "A Single-Chip 900 MHz CMOS Receiver Fromt-End with a High Performance Low-IF Topology," *IEEE Journal of Solid-State Circuits*, Vol. 30, No. 12. pp. 1483-1492, December 1995.

[11] A. Feldman, *High-Speed, Low-Power Sigma-Delta Modulators for RF Baseband Channel Applications*, Ph.D. Thesis, Electronics Research Laboratory, University of California at Berkeley, 1997.

[12] Jacques Rudell, et al. "A 1.9 GHz Wide-Band IF Double Conversion CMOS Integrated Receiver for Cordless Telephone Applications," *Digest of Technical Papers, International Solid-State Circuits Conference*, pp.304-305, February 1997.

[13] T. Cho. et al. "A Power-Optimized CMOS Baseband Channel Filter and ADC for Cordless Applications," *1996 Symposium on VLSI Circuits Digest of Technical Paters*, pp. 64-65, June 1996.

[14] H. Khorramabadi, et al. "Beseband Filters for IS-95 CDMA Receiver Aplications Featuring Digital Automatic Frequency Tuning," *Digest of Technical Papers, Intermational Dolid-State Circuits Conference*, pp. 172-173 February 1996.

[15] "Jacques Rudell, et. al. *Specifications for an Integrated CMOS GSM / DECT Receiver*, Electronics Research Laboratory M97/82, University of California at Berkeley, 1997.

[16] *European Digital Cellular Telecommunications System (Phase 2)*, European Telecommunications Standards Institute, 1994.

[17] *Digital European Cordless Telecommunications (DECT) Approval Test Specification*, European Telecommunications Standards Institute, 1992.

[18] Louis Williams, *Modeling and Design of High-Resolution Sigma-Delta Modulators*, Ph.D. Dissertation, Stanford University, 1993.

[19] James Candy and Gabor Temes, "Oversampling Methods for A/D and D/A Conversion," in *Oversampling Delta-Sigma Data Converters*. IEEE Press, New York, 1992.

[20] Brian Brandt, *Oversampled Analog-to-Digital Conversion*, Ph.D. Dissertation, Stanford University, 1991.

[21] R. E. Crochiere and L. R. Rabiner, *Multirate Digital Signal Processing*, Prentice Hall, 1983.

[22] Eugene Hogenauer, 'An Economical Class of Digital Filters for Decimation and Interpolation," *IEEE Trans. on Acoustics, Speech, and Signal Processing*, Vol ASSP-29, No. 2, April 1981

[23] Jan Rabaey, *Digital Integrated Circuits*, Prentice Hall, 1996.

[24] Brian Brandt and Bruce Wooley, "A Low-Power, Area-Efficient Digital Filter for Decimation and Interpolation," *IEEE Journal of Solid-State Circuits*, Vol. 29, No. 6. pp. 679-687, June 1994

[25] Edwin Chan, *Low Power Audio Codec*, Masters Thesis, University of California at Berkeley, 1996.

[26] James Candy, "Decimation for Sigma-Delta Modulation," *IEEE Transaction on Communications*, Vol. COM-34, pp. 72-26, January 1986.

[27] L. R. Rabiner and G. Gold, *Theory and Application of Digital Signal Processing*, Prentice Hall, 1975.

[28] P. Vaidyanthan and T. Nguyen, "A "trick" for the design of FIR and Half-band Filters," *IEEE Trans. on Circuits and Systems*, Vol. CAS-34, No. 3, March 1987.

[29] James Peroulas, *Design and Implementation of a High Speed CDMA Modulator for the INFOPAD Basestations*, Masters Thesis, University of California at Berkeley, 1995

[30] Henry Samueli, "An Improves Search Algorithm for the Design of Multiplierless FIR Filters with Powers-of-Two Coefficients,"*IEEE Trans. on Circuits and Systems*, Vol, 36, No. 7, July 1989.

[31] David Goodman, "Nine Digital Filters for Decimation and Interpolation," *IEEE Trans. on Acoustics, Speech and Signal Processing*, Vol, ASSP-25, No. 2, April 1977

[32] S. Norsworthy et al., *Delta-Sigma Data Converters*. IEEE Press, 1997.

[33] Arthur Abnous, private communication.

[34] A. Oppenheim and R. Schafer, *Discreet-Time Signal Processing*, Prentice Hall, 1989.

[35] N. Tan, et. al. "A Novel Bit-Serial Design of Comb Filters for Oversampling A/D Converters," *Proceedings of IEEE International Symposium on Circuits and Systems*, pp. 259-262, May 1994.

[36] I. Fujumori et al. "A 5V Single-Chip Delta-Sigma Audio A/D Converter with 111dB Dynamic-Range," *IEEE Journal of Solid-State Circuits*, pp. 329-336, March 1997.

[37] K. Sienski et al. "Hardware Efficient FIR Filter Structures From Linear Program Design Constraints," *1996 IEEE International Conference Proceedings on Acoustics, Speech, and Signal Processing*, pp. 1280-2, May 1996.