THE UNIVERSITY OF CALGARY

Controllable, Analogue, Three-Dimensional Mixed Domain Linear Trajectory Filters for Video Signals

by

D. Cameron Taylor

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES

IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE

DEGREE OF MASTER OF SCIENCE

DEPARTMENT OF ELECTRICAL AND COMPUTER

ENGINEERING

CALGARY, ALBERTA

July, 1996

0-612-20886-9

Abstract

Beginning with the fundamentals of multidimensional systems theory and describing the design, implementation and testing of two prototypes, this thesis is an investigation into the use of analogue, mixed continuous-discrete domain circuits for real-time three-dimensional linear trajectory filtering of raster scanned video signals. A first order frequency planar pass filter and a second order frequency bowl pass filter have been constructed such that their parameters are controllable in real-time. A method to measure the three-dimensional frequency responses of these filters is applied for the first time and the results are presented. The sensitivity of these filters to errors in the delay elements is derived and their stability under variation of the delay lengths is investigated.

Acknowledgements

First, I would like to thank Dr. Bruton for the advice and encouragement he has given me in his role as supervisor. I have enjoyed and appreciated our interaction in both research and teaching situations.

I would like to thank Hai-Ling Margaret Cheng, Chia-Luh Chung, Andreas Dilger and Jophes Provine for their work in reviewing this thesis. Their suggestions and editing have made it both more readable and more informative. I would also like to thank Norm Bartley for his assistance in circuit design and programming. Warren Flaman, John Shelly and Garry Harrington have provided invaluable service in the construction of the filter. Finally, I would like to thank Christopher Kulach for the use of his controller module design and software and his assistance in adapting them to my use.

This research has been supported in part by a NSERC Postgraduate scholarship and by MICRONET, The Federal Centres of Excellence in Microelectronics, Devices and Systems.

This thesis is typeset in TeXand LaTeX, free document preparation systems by Donald Knuth, Leslie Lamport and many others.

To the glory of God,

Creator and Sustainer;

and to my wife, Cynthia, who's support

and encouragement have made all the difference

Contents

List of Tables

List of Figures

Chapter 1

Introduction

The recent explosion of interest in the creation, manipulation, measurement and broadcasting of video sequences along with such diverse industrial applications as seismic analysis, biomedical imaging and synthetic aperture radar has spurred the advance of multidimensional signal processing theory to the point where useful two and three dimensional filters can be designed using a number of straightforward methods. However, the problem of implementing these designs so that they meet such requirements as real-time operation, low cost and low power has yet to be solved. Most implementations of video filters currently available use digital multipliers and adders to make the required calculations. Because video sequences contain a very large amount of data and even the simplest 3-D filters require several operations on each datum or voxel,[1] these implementations either involve large amounts of hardware or operate at a lower than real-time rate.

For example, a standard 512 by 480 pixel image filtered at 30 frames per second requires one output voxel to be computed every 136 ns. In this thesis, this is referred to as real-time filtering. Discrete digital hardware or a high-speed workstation is capable of making 1 multiplication and 1 addition in that amount of time, but the first order recursive 3-D transfer function requires 15 multiply and 14 add operations

[1]In 2-D, a picture element is called a pixel. In 3-D, each discrete value in a signal is associated with a volume and so is known as a volume element, or voxel.

for each voxel [1]. Thus these filters must use massively parallel processing to achieve real-time operation.

Recently, systems using analogue operational amplifiers and reactance elements to make the computations have been implemented for real-time filtering of raster scanned video sequences [2–4]. These systems, known as mixed continuous-discrete domain systems, greatly reduce the amount of hardware required, especially for 3-D filtering; however, they have previously been tunable only by the adjustment of component values. Chapters 3 and 4 of this thesis describe the design and implementation of a mixed continuous-discrete domain filter that is controllable in real-time.

1.1 Mixed Continuous-Discrete Domain Signals

A mixed continuous-discrete domain (mixD) signal is simply a multidimensional signal that is continuous in one or more of the dimensions it is defined over and discrete in the others. Thus, where a digital domain signal has a region of support that is a set of points, a mixD system's region of support is a set of lines, planes, solids, etc. Systems that operate on mixD signals are known as mixD systems.

MixD signals are common in video applications. The three most common formats for television broadcast, NTSC (North America), PAL (Europe) and SECAM (Europe) are all raster scan video formats, and therefore mixD signals. The three dimensional images represented by these signals are continuous in the horizontal direction, but discrete both vertically and temporally. Digital systems discretize these signals by sampling in the horizontal dimension, and then often end up constructing a mixD output signal. Both of these operations are complex and require costly hardware. Dedicated analogue early vision systems [6–8], on the other hand, involve expensive VLSI techniques and can not operate on a stored or transmitted signal. The analogue filters described in this thesis operate directly on a NTSC raster scanned video signal in real-time and use mixD signals internally.

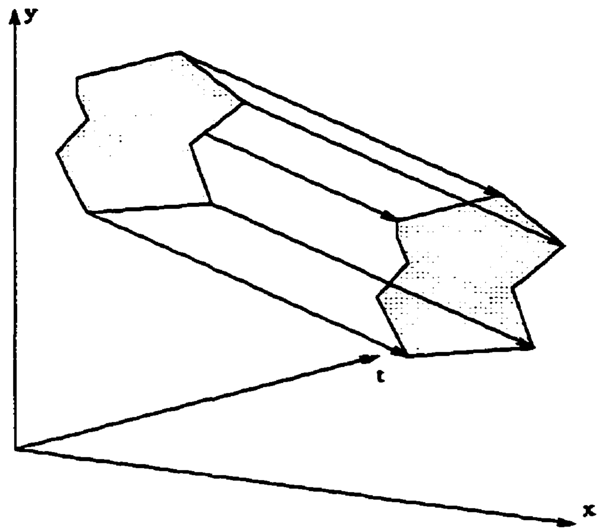

Figure 1.1: An Example of a 3-D Continuous Domain Linear Trajectory Signal

1.2 Linear Trajectory Filters

Many video processing applications make use of velocity information about objects in the image; therefore, a filter that can distinguish between objects on the basis of their motion is useful. Solid objects moving in a straight line at a constant speed, with no rotation or change of size, as shown in Figure 1.1, correspond to linear trajectory (LT) signals [9]. While motion in video sequences is seldom purely translational, it is nearly always smooth; objects usually have inertia. A smoothly curved trajectory can be approximated by a decomposition into a set of piecewise linear trajectory signals [10]. This makes a filter that can be tuned in real-time to selectively pass LT signals—a 3-D LT filter—useful for extracting velocity information.

LT filters, and velocity selective filters in general, are costly in terms of computation, so highly efficient methods for their implementation are of interest. The analogue filters designed and tested in this thesis are first and second order linear trajectory filters. It is shown that they can be built using much less hardware than is required for a digital implementation. The first order filter—known as the integrator-differentiator double loop (IDD) filter, after the form of the final signal flow graph—is the simplest known LT filter [3, 4, 11]. The second order filter has better directional selectivity [12], that is, an improved ability to enhance or reject a LT signal on the basis of its 3-D trajectory [13]. These two filters have been implemented on a common platform known as the Endeavour Analogue Video Filter. The system includes interfaces to a video camera, a monitor and a workstation along with a graphical user interface that allows it to be tuned easily and precisely while in operation.

The filters developed here are intended to be used as building blocks in a variety of video processing applications including video compression, object tracking and classification and computer vision. One application would be to use a number of LT filters to measure the energy in the signal associated with objects moving in different directions and so determine the dominant motional components for video compression [5]. Another would be to use an adaptive LT filter to extract a moving object from noise and clutter for a tracking algorithm [10, 14]. A final possibility is to combine these two functions to track and classify multiple moving objects for machine vision.

1.3 Notation

Multidimensional equations often contain many terms, compared with one dimensional equations of similar order, which can make them difficult to read and may hide their general structure. This is especially true of equations involving summations and integrations, so a more compact form of these has been adopted from [3] for this thesis. A vector is typeset in boldface

$$\mathbf{t}^{(m)} \equiv [t_1, t_2, \ldots, t_m]^T \tag{1.1}$$

where m is usually clear from the context and omitted. Multidimensional signals are represented as functions of more than one variable and are thus written as functions of vectors or M-tuples

$$x(\mathbf{t}^{(m)}) \equiv x(t_1, t_2, \ldots, t_m) \tag{1.2}$$

and mixD signals are written as functions of two M-tuples, one continuous and one discrete:

$$x(\mathbf{t}^{(p)}, \mathbf{n}^{(m-p)}) = x(t_1, \ldots, t_p, n_{p+1}, \ldots, n_m). \tag{1.3}$$

A summation over a vector is taken to be a multiple summation over each member of the vector

$$\sum_{\mathbf{i}=\mathbf{A}}^{\mathbf{B}} x(\mathbf{i}^{(m)}) \equiv \sum_{i_1=A_1}^{B_1} \sum_{i_2=A_2}^{B_2} \cdots \sum_{i_m=A_m}^{B_m} x(i_1, i_2, \ldots, i_m) \tag{1.4}$$

and similarly for integration or differentiation over a vector

$$\begin{gathered} \int_{\mathbf{t}=\mathbf{A}}^{\mathbf{B}} x(\mathbf{t}^{(m)}) d\mathbf{t} \equiv \int_{t_1=A_1}^{B_1} \int_{t_2=A_2}^{B_2} \cdots \int_{t_m=A_m}^{B_m} x(t_1, t_2, \ldots, t_m)\, dt_m \ldots dt_2\, dt_1 \\ \frac{\partial^{\mathbf{i}}}{\partial \mathbf{t}^{\mathbf{i}}} y(\mathbf{t}^{(m)}) \equiv \frac{\partial^{i_1}}{\partial t_1^{i_1}} \cdots \frac{\partial^{i_m}}{\partial t_m^{i_m}} y(t_1, \ldots, t_m). \end{gathered} \tag{1.5}$$

The vector exponent of a vector is

$$\mathbf{z}^{\mathbf{i}} = z_1^{i_1} z_2^{i_2} \cdots z_n^{i_n}. \tag{1.6}$$

1.4 Thesis Organisation

In this thesis the design and implementation of a controllable, analogue, 3-D, mixed continuous-discrete domain linear trajectory filter for NTSC video signals is proposed, described and tested. The spatio-temporal and frequency domain characteristics, sensitivity to parameter changes and stability of the filter are investigated and recommendations are made for improvements.

Chapter two is a review of multidimensional signals and systems. Useful tools for the description and design of continuous, discrete and mixed domain systems are covered briefly and some important issues such as stability are touched upon. Also, linear trajectory signals and the raster scanning process are discussed.

In chapter three the designs of the first order planar pass or IDD filter [3,4] and the second order bowl filter [5,12] are reviewed and modified slightly. The theoretical effectiveness of each in enhancing or rejecting signals on the basis of their 3-D velocity is also discussed. Manipulating the filter trajectory by changing the delay lengths overcomes one of the basic limitations of the technique [3] and is investigated here.

Chapter four describes the implementation of the signal flow graphs designed in chapter three. The hardware and software modules are described individually and in the context of the overall design. Some elements are adapted from previous designs [3,5,15] while others are original. The major improvement in this implementation and a principal contribution of this thesis is that the filter parameters are controllable in real-time via input currents to analogue multipliers. These currents are presently supplied by digital to analogue converters controlled via a workstation, but the option of control by analogue circuitry exists. Suggestions for improvements in future implementations are also given.

The filters are thoroughly characterized in chapter five, both in the spatio-temporal domain and in the 3-D frequency domain. Non-ideal effects such as overflow instability are discussed along with methods to reduce them. The 1-D to 3-D frequency response transformation for raster scanned systems proposed in [5] is applied to both filters to provide a complete frequency domain characterization, which has not been done before. A 1-D high pass post filter is also suggested to improve directional selectivity at low frequencies and to remove a non-ideal dominant pole characteristic in the measured responses.

In chapter six the sensitivity of mixD systems to delay element errors is investigated. Lower bounds on the first order and Schoeffler sensitivities of structures implementing a given transfer function are found. The direct form and ladder form implementations of the first order planar pass filter are compared to the lower bound and it is shown that the ladder form implementation is superior.

Chapter seven investigates the stability of mixD systems. Conditions for practical BIBO stability are found for various regions of support and a technique for designing practically BIBO stable mixD systems from continuous positive M-D networks is developed. The filters designed in chapter three conform to these conditions. The question of the stability of these filters under delay length variations is also addressed.

Finally, chapter eight gives a summary of the thesis and makes suggestions for future research.

Chapter 2

Review of Multidimensional Signals and Systems

A Multidimensional (M-D) signal can be respresented by a function of more than one independent variable. Examples of such functions include the intensity or luminosity of a photograph, seismic data, biomedical imaging data and synthetic aperture radar sensor data; but the signal of interest in this thesis is a time varying image, such as a video sequence. A video sequence is a three dimensional (3-D) signal, with the three dimensions being horizontal position, vertical position and time, and is said to be in the spatio-temporal domain.

M-D signals can be classified [3] as continuous domain signals, which are functions of continuous variables; discrete domain signals, which take on values only at discrete values of the independent variables; or mixed continuous-discrete domain (mixD) signals, for which some of the independent variables are continuous and others are discrete. Continuous and discrete domain signals can be thought of as special cases of mixD signals, but that is unnecessarily complex in most situations.

Systems which process M-D signals are called M-D systems and are classified in the same way. This chapter will review some basic concepts and tools in M-D systems theory. For more information the reader is referred to [16–21]. The concepts will be covered for continuous, discrete and mixed domain signals in that order.

2.1 M-D Continuous Domain Signals and Circuits

It is common, in the design of M-D systems, to begin with a continuous domain model and move to the appropriate domain through a number of linear transformations [3, 4, 12, 22–26] (see section 3.4). The main advantage of this approach is the ability to draw on some well known properties of continuous domain circuits, which may be described by M-D differential equations, through Kirchoff's laws and Tellegen's theorem [24].

2.1.1 M-D Circuit Elements and Differential Equations

M-D circuits imply the conservation of current and voltage surfaces, $i(\mathbf{t}^{(m)})$ and $v(\mathbf{t}^{(m)})$, in a manner completely analogous to 1-D circuits, and may be constructed and solved in a similar manner. M-D circuit elements are also defined in a manner very similar to 1-D elements [24]. Resistors support a voltage surface relative to the current flowing through them as

$$v(\mathbf{t}) = R\,i(\mathbf{t}). \tag{2.1}$$

Inductors perform partial directional differentiation of the current along one dimension, t_i, as

$$v(\mathbf{t}) = L_i \frac{\partial}{\partial t_i} i(\mathbf{t}) \tag{2.2}$$

and capacitors do the same with voltage as

$$i(\mathbf{t}) = C_i \frac{\partial}{\partial t_i} v(\mathbf{t}). \tag{2.3}$$

The schematic symbols for M-D inductors and capacitors are shown in Figure 2.1, while an M-D resistor can be drawn the same as a 1-D resistor. M-D transformers, gyrators and independent sources are the same as their 1-D counterparts.

Systems made of interconnections of these elements can be described by the M-D differential equation

$$\sum_{\mathbf{i}=0}^{\mathbf{N}} b_{\mathbf{i}} \frac{\partial^{\mathbf{i}}}{\partial \mathbf{t}^{\mathbf{i}}} y(\mathbf{t}) = \sum_{\mathbf{i}=0}^{\mathbf{M}} a_{\mathbf{i}} \frac{\partial^{\mathbf{i}}}{\partial \mathbf{t}^{\mathbf{i}}} x(\mathbf{t}) \tag{2.4}$$

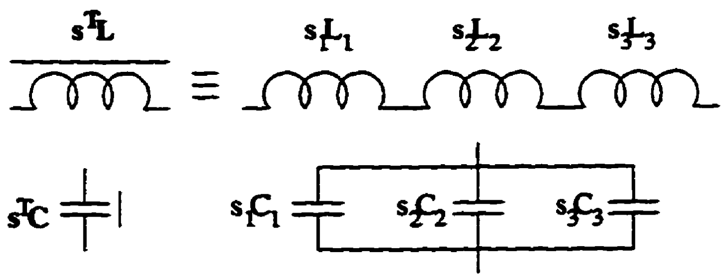

Figure 2.1: The Schematic Symbol for an M-D Inductor and an M-D Capacitor

where $b_0 = 1$, $x(\mathbf{t}) \in \mathcal{R}$ is the input and $y(\mathbf{t}) \in \mathcal{R}$ is the output of the system, which may be a voltage at a node or a current through an element. This class of systems is linear and shift-invariant and they can therefore be fully characterized by their impulse responses [12].

2.1.2 M-D Impulse Response and Convolution

The manipulation of differential equations is difficult, so a number of techniques for simplifying or avoiding their use in the design of linear shift-invariant (LSI) systems have been developed. Central to these techniques is the system impulse response. The M-D continuous domain impulse function $\delta(\mathbf{t})$ is defined by the two properties

$$\begin{gathered}\int_{-\infty}^{\infty} \delta(\mathbf{t})\, d\mathbf{t} = 1 \\ \text{and} \quad \delta(\mathbf{t}) = 0,\ \forall \mathbf{t} \neq \mathbf{0}.\end{gathered} \tag{2.5}$$

If an LSI system is modeled by the operator $\Phi[\bullet]$ as $y(\mathbf{t}) = \Phi[x(\mathbf{t})]$ then the impulse response of the system is defined to be [12]

$$h(\mathbf{t}) \equiv \Phi[\delta(\mathbf{t})]. \tag{2.6}$$

It can be shown [27,28] that the zero initial condition output response of the system is fully determined by the impulse response, $h(\mathbf{t})$, and the input, $x(\mathbf{t})$. The zero initial condition response to any absolutely integrable input is given by the M-D convolution integral:

$$y(\mathbf{t}) = \int_{-\infty}^{\infty} x(\mathbf{r}) h(\mathbf{t} - \mathbf{r})\, d\mathbf{r}. \tag{2.7}$$

Conceptually, the convolution integral implies decomposing $x(\mathbf{t})$ into an infinite set of weighted and shifted impulses, each of which is then applied to the filter separately, with the results added together to form the output. It is also possible to decompose the input into an infinite set of weighted and shifted complex exponential functions in the form $e^{\mathbf{s}^T\mathbf{t}}$, where $\mathbf{s}$ is the M-D Laplace variable and $\mathbf{s}^T\mathbf{t}$ is the dot product of $\mathbf{s}$ and $\mathbf{t}$, which results in the M-D Laplace transform.

2.1.3 M-D Laplace Transform and Transfer Functions

The M-D Laplace transform $X(\mathbf{s})$ of the signal $x(\mathbf{t})$ is defined by [27]

$$X(\mathbf{s}) = \int_{-\infty}^{\infty} x(\mathbf{t})e^{-\mathbf{s}^T\mathbf{t}}\,d\mathbf{t} \tag{2.8}$$

where the region $\mathbf{s} \in \mathcal{C}^m$ such that the integrals converge to the same function for all values of $\mathbf{s}$ is known as the region of convergence of $X(\mathbf{s})$. The inverse transform is given by

$$x(\mathbf{t}) = \frac{1}{(2\pi j)^m} \int_{\sigma-j\infty}^{\sigma+j\infty} X(\mathbf{s})e^{\mathbf{s}^T\mathbf{t}}\,d\mathbf{s} \tag{2.9}$$

and $\boldsymbol{\sigma}$ is chosen so that the integrals will converge [27]. The M-D Fourier transform of an absolutely integrable signal $x(\mathbf{t})$ can be found by replacing $\mathbf{s}$ by $j\omega$ in $X(\mathbf{s})$.

Given these definitions, a LSI system can be characterized by the Laplace transform of the system impulse function $h(\mathbf{t})$, if it exists, which is known as the transfer function of the system. Given that the initial conditions of the system are zero, that is

$$\begin{aligned} \frac{\partial^i}{\partial t_k^i}x(\mathbf{t}) = 0, i = 0, \ldots, M_k, \forall k \\ \frac{\partial^i}{\partial t_k^i}y(\mathbf{t}) = 0, i = 0, \ldots, N_k, \forall k \end{aligned} \tag{2.10}$$

the system described by equation 2.4 will have the transfer function

$$H(\mathbf{s}) = \frac{Y(\mathbf{s})}{X(\mathbf{s})} = \frac{\sum_{\mathbf{i}=0}^{\mathbf{M}} a_{\mathbf{i}}\mathbf{s}^{\mathbf{i}}}{\sum_{\mathbf{i}=0}^{\mathbf{N}} b_{\mathbf{i}}\mathbf{s}^{\mathbf{i}}}. \tag{2.11}$$

which relates the Laplace transform of the output of the system to that of the input. These transforms exist and are invertible if the system is stable and the input is square integrable [27].

2.1.4 Stability, Passivity and Losslessness

A system is said to be bounded-input bounded-output (BIBO) stable if, for every magnitude bounded input applied to the system, the output is also magnitude bounded. A continuous domain system is BIBO stable iff

$$\int_{-\infty}^{\infty} |h(\mathbf{t})|\, d\mathbf{t} \leq S_1 < \infty \tag{2.12}$$

where S_1 is a real finite number. Unfortunately this is, in general, difficult to test given the coefficients of equation 2.4. A sufficient condition on $H(\mathbf{s})$ for BIBO stability is that for the real parts of all s_i greater than or equal to zero, the denominator polynomial is non-zero [16]. However, because M-D polynomials cannot, in general, be completely factored, testing this condition is also usually difficult. Several classes of functions have been shown to be BIBO stable, notably the driving-point functions of M-D passive networks [22].

M-D resistors, inductors and capacitors with positive values as well as transformers and gyrators are passive [24], that is, the net energy delivered to any of these components is non-negative over all $\mathbf{t}$. Thus, systems containing only these elements are BIBO stable, and have been used to design useful discrete and mixD filters [3–5,9,11,12,20–24,29–32]. The filters that will be described in this thesis have all been designed from passive continuous domain circuit prototypes, and derive their stability from that fact (see section 7.4).

2.2 M-D Discrete Domain Signals and Systems

Discrete domain signals are generally obtained by sampling continuous domain signals on a rectangular grid such that

$$x(\mathbf{n}^{(m)}) = x_c(n_1T_1, n_2T_2, \ldots, n_mT_m) \tag{2.13}$$

where x_c is the continuous signal and T_i are the sampling intervals in each dimension. Other sampling methods are possible [17], but do not lend themselves as easily to processing video signals. The original signal can be reconstructed from the sampled signal if the original signal is bandlimited, that is, the magnitude of its Fourier transform is zero for $\omega_i > \pi/T_i$ [17]. The most common reason for sampling a signal is to digitize it and process it with a computer, or digital signal processor (DSP). DSP based filters generally apply a discrete difference equation.

2.2.1 M-D Discrete Difference Equation

The class of discrete domain systems of interest in this thesis can be described by the discrete difference equation [12, 17]

$$\sum_{\mathbf{i}=\mathbf{0}}^{\mathbf{N}} b_{\mathbf{i}} y(\mathbf{n}-\mathbf{i}) = \sum_{\mathbf{i}=\mathbf{0}}^{\mathbf{M}} a_{\mathbf{i}} x(\mathbf{n}-\mathbf{i}) \tag{2.14}$$

where, without loss of generality, $b_0 = 1$, $x(\mathbf{n})$ is the input and $y(\mathbf{n})$ is the output. These systems are first quadrant (or hyper-quadrant) causal, linear and shift-invariant if a and b are constant and so can be fully characterized by their impulse response [16].

2.2.2 M-D Impulse Response and Convolution

The M-D discrete domain impulse function $\delta(\mathbf{n})$ is defined as

$$\delta(\mathbf{n}) = \begin{cases} 1, & \mathbf{n} = \mathbf{0} \\ 0, & \text{otherwise} \end{cases} \tag{2.15}$$

and, given that the system can be modeled by the operator $\Phi[\bullet]$ such that $y(\mathbf{n}) = \Phi[x(\mathbf{n})]$, the impulse response is defined to be [16]

$$h(\mathbf{n}) = \Phi[\delta(\mathbf{n})]. \tag{2.16}$$

Discrete domain systems are classified as finite impulse response (FIR) if the impulse response is of bounded extent in all dimensions, and as infinite impulse response (IIR) otherwise [12].

For a LSI system, the input can be decomposed into a set of weighted, shifted impulses, which can be applied separately to the system and then summed, to arrive at the output. This process results in the M-D convolution sum:

$$y(\mathbf{n}) = \sum_{\mathbf{k}=-\infty}^{\infty} x(\mathbf{k})h(\mathbf{n}-\mathbf{k}) \tag{2.17}$$

It is also possible to decompose the input into an infinite set of weighted and shifted complex exponential functions in the form $\mathbf{z}^{\mathbf{n}}$, leading to the M-D Z Transform.

2.2.3 M-D Z Transform and Transfer Functions

The M-D Z transform $X(\mathbf{z})$ of the discrete domain signal $x(\mathbf{n})$ is defined to be [17]

$$X(\mathbf{z}) = \sum_{\mathbf{k}=-\infty}^{\infty} x(\mathbf{n})\mathbf{z}^{-\mathbf{k}} \tag{2.18}$$

where $\mathbf{z} \in \mathcal{C}^m$. The region in $\mathcal{C}^m$ where this sum converges to the same finite function for all choices of $\mathbf{z}$ is known as the region of uniform convergence (ROC) of $X(\mathbf{z})$ [16]. The Fourier transform of $x(\mathbf{n}^{(m)})$ can be found from the Z transform, $X(\mathbf{z}^{(m)})$ by substituting $e^{j\Omega_i}$ for z_i, $i = 1, 2, \ldots, m$ [17], if the ROC includes the distinguished boundary of the unit disk, which it will if the signal is absolutely summable [16].

The inverse Z transform is given by

$$x(\mathbf{n}) = \frac{1}{(2\pi j)^m} \oint_C X(\mathbf{z})\mathbf{z}^{\mathbf{n}-1}\, d\mathbf{z} \tag{2.19}$$

where each contour of integration must be closed, lie completely within the ROC of $X(\mathbf{z})$ and encircle the origin counterclockwise [17].

Given these definitions, an LSI system can be characterized by the Z transform of the system impulse response $h(\mathbf{n})$, if it exists, which is known as the Z transform transfer function of the system. Given that the initial conditions of the system are zero, that is

$$\begin{aligned} x(\mathbf{n}) &= 0, \forall \mathbf{n} | n_i < 0 \\ y(\mathbf{n}) &= 0, \forall \mathbf{n} | n_i < 0 \end{aligned} \tag{2.20}$$

the system described by equation 2.14 has the Z transform transfer function [12, 16]

$$H(\mathbf{z}) = \frac{Y(\mathbf{z})}{X(\mathbf{z})} = \frac{\sum_{\mathbf{k}=0}^{\mathbf{M}} a_{\mathbf{k}} \mathbf{z}^{-\mathbf{k}}}{\sum_{\mathbf{k}=0}^{\mathbf{N}} b_{\mathbf{k}} \mathbf{z}^{-\mathbf{k}}} \tag{2.21}$$

which relates the Z transform of the output of the system to that of the input. These transforms exist and are invertible if the system is stable and the input is absolutely summable [17].

It is common, in M-D filter design as well as in 1-D filter design, to use the well known properties of continuous domain circuits discussed in section 2.1.4 in the initial design by beginning with a continuous domain prototype and moving into the discrete domain via a frequency domain transform such as the bilinear transform. The design is then completed in the discrete domain. The bilinear transform in dimension i is given by the substitution

$$s_i \Rightarrow \frac{2}{T_i} \frac{z_i - 1}{z_i + 1}. \tag{2.22}$$

2.2.4 Stability

A discrete domain system is BIBO stable (see section 2.1.4) iff [16, 31]

$$\sum_{\mathbf{n}=-\infty}^{\infty} |h(\mathbf{n})| \leq S_1 < \infty \tag{2.23}$$

where S_1 is a finite real number. However, this implies that the output of the filter is calculated for points extending to infinity in all directions. It has been shown [33] that this is not possible with a finite state machine, so the BIBO stability condition is too restrictive. It is sufficient, in practice, to ensure stability in a region of finite extent in all but one direction, which is known as practical bounded input bounded output (PBIBO) stability [31]. If that direction is along one of the axes, the condition

$$\sum_{\mathbf{n}=\mathbf{0}}^{\mathbf{N}_{(k)}} |h(\mathbf{n})| \leq S_1 < \infty, \forall k \tag{2.24}$$

where all elements of $\mathbf{N}_{(k)}$ are bounded integers except for element k, which is ∞, is necessary and sufficient for PBIBO stability [31]. Chapter 7 extends this condition to the case where the direction of infinite extent is not along one of the axes.

All continuous voltage transfer functions of an M-D circuit containing only positive resistors, capacitors, inductors, transformers and gyrators with the output voltage measured across a terminating resistor lead to PBIBO stable discrete filters after bilinear transformation in a rectangular region of support [31]. This fact is used extensively through out the designs, along with the extensions to mixD systems and non-rectangular regions of support discussed in chapter 7.

2.3 M-D Mixed Continuous-Discrete Domain Signals and Systems

MixD system models[1] have been useful in describing continuous 1-D systems containing both reactance elements and delay elements [3,4,16,37,38], such as those designed in this thesis. While the input to the filters is, in fact, a 1-D signal, it may be interpreted, through a raster scan transformation (see section 2.5), to be a 3-D mixD signal. It is then appropriate to associate a complex frequency with the reactance elements and different complex frequencies with delays of two different lengths [21,37].

Because the filters implement the differentiation and integration associated with the continuous dimensions of the mixD differential/difference equation with lumped elements that differentiate with respect to time, they can only process signals that have one continuous dimension [3]. However, for the sake of completeness this review

[1]This is not mixed DFT/LDE filtering described in [34–36].

will treat mixD signals and systems with p continuous dimensions and $m-p$ discrete dimensions.

MixD signals will be indicated as such by being functions of two l-tuples, one of continuous variables and one of discrete.

2.3.1 Differential/Difference Equations

All of the mixD systems in this thesis are approximately linear and shift-invariant within the region of calculation (the signals are all bounded spatially) and may be described by the differential/difference equation with constant coefficients [2, 3]

$$\sum_{\mathbf{i}=0}^{\mathbf{N}_c}\sum_{\mathbf{k}=0}^{\mathbf{N}_d} b_{\mathbf{i},\mathbf{k}} \frac{\partial^{\mathbf{i}}}{\partial \mathbf{t}^{\mathbf{i}}} y(\mathbf{t}, \mathbf{n}-\mathbf{k}) = \sum_{\mathbf{i}=0}^{\mathbf{M}_c}\sum_{\mathbf{k}=0}^{\mathbf{M}_d} a_{\mathbf{i},\mathbf{k}} \frac{\partial^{\mathbf{i}}}{\partial \mathbf{t}^{\mathbf{i}}} x(\mathbf{t}, \mathbf{n}-\mathbf{k}) \tag{2.25}$$

where, without loss of generality, $b_{\mathbf{0},\mathbf{0}} = 1$ and the input and output are $x(\mathbf{t},\mathbf{n})$ and $y(\mathbf{t},\mathbf{n})$ respectively. These systems can also be fully characterized by their impulse responses.

2.3.2 M-D Impulse Response and Convolution

The M-D mixD impulse function $\delta(\mathbf{t},\mathbf{n})$ is defined by the two properties [3]

$$\int_{-\infty}^{\infty} \delta(\mathbf{t},\mathbf{n})\,d\mathbf{t} = \begin{cases} 1, & \mathbf{n}=\mathbf{0} \\ 0, & \text{otherwise} \end{cases} \tag{2.26}$$

and

$$\delta(\mathbf{t},\mathbf{n}) = 0, \ \forall \mathbf{t} \neq \mathbf{0}. \tag{2.27}$$

If an LSI system is modeled by the operator $\Phi[\bullet]$ as $y(\mathbf{t},\mathbf{n}) = \Phi[x(\mathbf{t},\mathbf{n})]$, then the impulse response of the system is defined to be [3]

$$h(\mathbf{t},\mathbf{n}) \equiv \Phi[\delta(\mathbf{t},\mathbf{n})] \tag{2.28}$$

The zero initial condition output response of the system can then be determined from the impulse response, $h(\mathbf{t},\mathbf{n})$ and the input, $x(\mathbf{t},\mathbf{n})$, with the M-D convolution integral/sum given by [2,3]

$$y(\mathbf{t},\mathbf{n}) = \sum_{\mathbf{k}=-\infty}^{\infty} \int_{-\infty}^{\infty} x(\mathbf{r},\mathbf{k}) h(\mathbf{t}-\mathbf{r},\mathbf{n}-\mathbf{k})\, d\mathbf{r} \tag{2.29}$$

This equation is similar to that for continuous domain systems given in equation 2.7, and for discrete domain systems given in equation 2.7. If the impulse response and input are both first quadrant (or hyper-quadrant) causal, that is, if they are zero for any index negative, then the lower limits of the sum and integral may be changed to 0 and the upper limits to $\mathbf{n}$ and $\mathbf{t}$ respectively.

Once again, the convolution sum is conceptually the effect of decomposing the input into an infinite set of weighted and shifted impulses, which are applied separately to the filter. The results are then summed to give the output.

The eigenfunctions of equation 2.29 are of the form $e^{-\mathbf{s}^T\mathbf{t}}\mathbf{z}^{-\mathbf{n}}$, so it is useful to apply the Laplace/Z transform and investigate the characteristics of the system in the sz-domain.

2.3.3 Laplace/Z Transform and Transfer Functions

The M-D Laplace/Z transform can be defined by applying the Laplace transform to the continuous dimensions, with all the conditions given in section 2.1.3, and the Z transform to the discrete dimensions, with all the conditions given in section 2.2.3, resulting in [2]

$$X(\mathbf{s},\mathbf{z}) = \sum_{\mathbf{n}=-\infty}^{\infty} \int_{-\infty}^{\infty} x(\mathbf{t},\mathbf{n}) e^{-\mathbf{s}^T\mathbf{t}}\mathbf{z}^{-\mathbf{n}}\, d\mathbf{t}. \tag{2.30}$$

The region in C^m where this converges is defined as the region of convergence (ROC) of $X(\mathbf{s},\mathbf{z})$. The mixD Fourier transform can be found by substituting $j\omega_i$ for s_i and $e^{j\Omega_i}$ for $z_i, i = 1, 2, \ldots, m$ [3] if the ROC includes the imaginary axis in each s-domain and the unit circle in each z-domain. This will occur if x is absolutely sum/integrable, that is, if

$$\sum_{\mathbf{n}=-\infty}^{\infty} \int_{-\infty}^{\infty} |x(\mathbf{t},\mathbf{n})|\, d\mathbf{t} \leq S_1 < \infty \tag{2.31}$$

where S_1 is a non-infinite real number.

The inverse Laplace/Z transform is given by

$$x(\mathbf{t},\mathbf{n}) = \frac{1}{(2\pi j)^m} \oint_C \int_{\sigma-j\infty}^{\sigma+j\infty} X(\mathbf{s},\mathbf{z}) e^{\mathbf{s}^T\mathbf{t}} \mathbf{z}^{\mathbf{n}-1}\, ds\, dz \tag{2.32}$$

where the contour C is the same as for equation 2.19 and σ is chosen so that the inner integral converges.

Applying the Laplace/Z transform to equation 2.25 and assuming zero initial conditions (equation 2.10 applied for the continuous dimensions and equation 2.20 for the discrete dimensions), the transformed output can be related to the transformed input via the system transfer function given by [3]

$$H(\mathbf{s},\mathbf{z}) = \frac{Y(\mathbf{s},\mathbf{z})}{X(\mathbf{s},\mathbf{z})} = \frac{\sum_{\mathbf{i}=0}^{\mathbf{M}_c} \sum_{\mathbf{k}=0}^{\mathbf{M}_d} a_{\mathbf{i},\mathbf{k}} \mathbf{s}^{\mathbf{i}} \mathbf{z}^{-\mathbf{k}}}{\sum_{\mathbf{i}=0}^{\mathbf{N}_c} \sum_{\mathbf{k}=0}^{\mathbf{N}_d} b_{\mathbf{i},\mathbf{k}} \mathbf{s}^{\mathbf{i}} \mathbf{z}^{-\mathbf{k}}}. \tag{2.33}$$

Since the transform of a unit impulse is 1, it is apparent that the transfer function is equal to the transform of the impulse response. The input and output transforms exist and are invertible if the input is absolutely sum/integrable and the filter is stable.

2.3.4 Stability

A mixD system is BIBO stable iff

$$\sum_{\mathbf{n}=-\infty}^{\infty} \int_{-\infty}^{\infty} |h(\mathbf{t},\mathbf{n})|\, d\mathbf{t} \le S_1 < \infty \tag{2.34}$$

and PBIBO stable over a rectangular region of support iff

$$\sum_{\mathbf{n}=0}^{\mathbf{N}} \int_{0}^{\mathbf{T}} |h(\mathbf{t},\mathbf{n})|\, d\mathbf{t} \le S_1 < \infty,\ \forall k. \tag{2.35}$$

where one of N_k or T_k is unbounded and all other limits are bounded and S_1 is a non-infinite real number. The proof is given in chapter 7. It is also shown in chapter 7 that transfer functions of passive M-D circuits lead to PBIBO stable mixD circuits after the bilinear transform is applied to $m - p$ dimensions in all rectangular, and some non-rectangular, regions of support.

2.4 Linear Trajectory Signals

In video sequences the class of 3-D LT signals correspond to objects moving with constant velocity. The filters designed in this thesis make use of special properties of LT signals to enhance these moving objects in a way that is useful in many applications.

2.4.1 Continuous Spatio-Temporal Domain

A signal is defined to be linear trajectory if there exists a line in M-D space such that the signal has constant value along all lines parallel to it [9]. For these lines to exist, that is, for them to have a defined value along their length, the signal must be in the continuous domain. The unit vector $\mathbf{d}_s$ in the direction of the line of constant value is known as the signal trajectory. Alternately, this can be stated as [3]

$$\nabla x(\mathbf{t}) \bullet \mathbf{d}_s = 0, \ \forall \mathbf{t} \in \mathcal{R}^m \tag{2.36}$$

where $\nabla x(\mathbf{t}) \equiv [\frac{\partial}{\partial t_1} x(\mathbf{t}), \ldots, \frac{\partial}{\partial t_m} x(\mathbf{t})]^T$ and $\bullet$ represents the dot product operation. That is, there is a direction in which the gradient of $x(\mathbf{t})$ is zero everywhere in $\mathbf{t}$.

In a time varying spatial image this corresponds to objects moving translationally along a straight path at a specific speed. The spatial velocity of the LT signal with trajectory $\mathbf{d}_s$ is [9] $\sqrt{d_{s1}^2 + d_{s2}^2}/d_{s3}$ at the angle $\arctan(d_{s1}/d_{s2})$ in the t_1, t_2 plane, where t_3 is the temporal variable. If t_1 is the horizontal dimension and t_2 is the vertical dimension the horizontal and vertical components of the velocity are $H_s = d_{s1}/d_{s3}$ and $V_s = d_{s2}/d_{s3}$ respectively.

2.4.2 Frequency Domain

The energy content of a continuous domain 3-D LT signal is confined to a plane in the frequency domain [9]. The normal to this plane in the frequency domain corresponds to the trajectory $\mathbf{d}_s$ in the spatio-temporal domain. Thus a filter with a passband closely surrounding a plane in the frequency domain (a frequency planar passband) can selectively enhance objects with particular trajectories.

The definitions of LT signals imply continuous domain signals of unbounded extent. To examine the effects of sampling and cropping an LT signal in several dimensions, as is done with video signals, the resulting aliasing and gating can be applied in each dimension successively. Assuming that the signal is bandlimited to half the sampling frequency in each dimension before it is sampled, the baseband signal will not experience interference from the image bands; that is, the signal will be zero in the plane of interest outside of half the sampling rate before sampling, so no aliasing effects will occur. Cropping a signal is mathematically equivalent to multiplying by a gate function in the spatio-temporal domain, or convolving with a sinc function in the frequency domain, which has the effect of spreading the energy of the signal out of the plane. The width of the main lobe of the sinc function is inversely related to the width of the gate function; so, if the length of the region of support of the signal is much greater than the sampling interval, this generally results in a signal with the majority of its energy in or near the original plane in the frequency domain.

2.5 NTSC Raster Scanned Video Format

A raster scan is a method of representing a 3-D signal as a 1-D signal for broadcasting or storage. The majority of North American video cameras and monitors use the National Television Standards Committee (NTSC) raster scanned video format, which encodes spatially bounded, time-varying images as a time-varying voltage [39]. The process of converting light reflecting from objects into a time-varying voltage signal involves a number of steps, most of which are not considered in detail here.

Initially, the light from an angularly (spatially) bounded portion of the scene is focused through a lens onto a grid of phototransducers and sampled. The aspect ratio (width to height) of the sampled image is 4:3. It is assumed that the signal is averaged both spatially and temporally by the phototransducers such that this sampling does not cause noticable aliasing.[2] The grids generally contain about 480 rows and 640

[2]This is not a particularly good assumption [40], and a great deal of effort is being made at reducing the aliasing effects in high definition television.

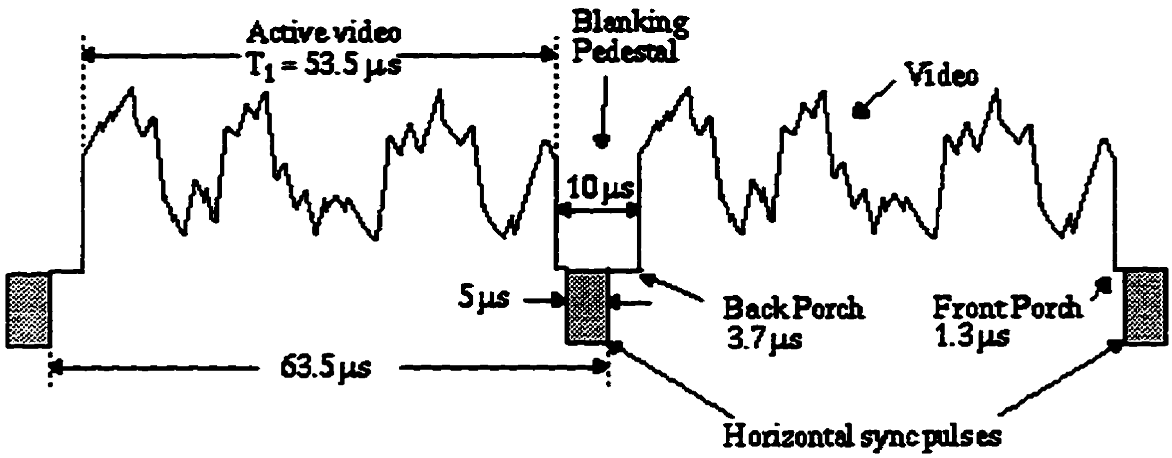

Figure 2.2: Two Lines of NTSC Raster Scanned Video(courtesy of Norm Bartley)

columns, and are sampled 30 times per second. Each still image, sampled temporally, is known as a frame.

The transformation from this 3-D discrete domain image sequence to the 1-D signal is of more interest here. Once a frame is sampled, the grid of transducers is scanned, row by row, left to right and top to bottom. The output voltage is proportional to the intensity of the image at the point being scanned, and is continuous along each row and bandlimited to 4.2MHz. Synchronization information is inserted between rows and frames in the form of "sync pulses" to enable reconstruction of the image at the monitor. Two lines of output are shown in Figure 2.2. NTSC video displayed on a screen is a mixD signal, continuous in the horizontal dimension and discrete in the vertical and temporal dimensions. Note that the fact that the scanning proceeds from the top downward implies a left handed coordinate system [3].

NTSC video is also interlaced; that is, two passes are made on each frame with odd rows being scanned on the first pass and even rows on the second. Each pass is known as a field. In this thesis the two fields are treated separately, in the order that they are sampled; that is, as a vertical concatenation. This sub-sampling by 2 results in aliasing in the vertical direction. Fortunately, most of the energy in most video sequences lies in the region $\|\mathbf{\Omega}\| < \pi/2$, known as the region of interest (ROI) [3,12], which means that there is little energy in the high frequency region to be aliased. This is equivalent to the assumption that the original image is highly oversampled vertically.

Due to the insertion of the synchronization information into the signal, the voltage only corresponds to the intensity of the image during the active video portion of each row (see Figure 2.2). Each field consists of 262.5 rows, so that there are 525 lines per frame. Of these, approximately 21 lines per field are reserved for the vertical synchronization information, during which the active video portion is constant at the back porch level. The number of active rows may vary from camera to camera.

In this thesis a **grain** is defined to be a unit of length equal to the distance between successive (odd or even) rows and between successive frames. It provides a device independent way of comparing distances and angles in 3-D image space for NTSC raster scan video signals.

2.5.1 The Raster Scan as a Transformation

The raster scanning process can be considered to be a transformation from a three dimensional mixD signal—continuous in the horizontal dimension and discrete in both vertical and temporal dimensions—into a continuous one dimensional signal. Consider the signal $x(h_1, n_2, n_3)$, where $h_1 \in \mathcal{R}$ is the horizontal position in grains, $n_2 \in \mathcal{Z}$ is the vertical position in grains (row number) and $n_3 \in \mathcal{Z}$ is the temporal position, also in grains (frame number). Then the raster scan applies the transformation of variables:

$$t = \begin{bmatrix} T_1 & T_2 & T_3 \end{bmatrix} \begin{bmatrix} h_1 \\ n_2 \\ n_3 \end{bmatrix} \tag{2.37}$$

so that

$$\begin{gathered} r(t) = r(h_1T_1 + n_2T_2 + n_3T_3) = x(h_1, n_2, n_3) \in \mathcal{R} \\ \text{for } 0 \leq h_1 \leq \tfrac{T_2}{T_1} \text{ and } 0 \leq n_2 \leq \tfrac{T_3}{T_2} \end{gathered} \tag{2.38}$$

where r is the raster scan of x; t is time in seconds; and $T_1, T_2, T_3 \in \mathcal{R}$ are the time it takes to scan one grain horizontally (167ns), one grain vertically (63.5μs) and one grain temporally (1/30 s) respectively. The fact that NTSC raster scan lines are tilted down from the horizontal by 0.15° is ignored.

The sync pulses allow the receiver to recover the row and frame numbers n_2 and n_3 and thus perform the inverse transformation[3]

$$h_1 = \frac{1}{T_1}\begin{bmatrix} 1 & -T_2 & -T_3 \end{bmatrix}\begin{bmatrix} t \\ n_2 \\ n_3 \end{bmatrix} \tag{2.39}$$

to recover $x(h_1, n_2, n_3)$ from $r(t)$.

The operation $t \rightarrow t - T_2$ corresponds exactly to $n_2 \rightarrow n_2 - 1$ and is called a row delay. Similarly, $t \rightarrow t - T_3$ corresponds to $n_3 \rightarrow n_3 - 1$ and is called a frame delay. The operation $t \rightarrow t - \tau$ where $-T_2 < \tau < T_2$ corresponds to $h_1 \rightarrow h_1 - \tau/T_1$.

In previous work on the subject [3–5] the horizontal dimension was measured in units of time and a scaling factor similar to T_1 was introduced to correct for it. Using grains as units eliminates this extra complexity and allows horizontal distances to be measured in units of length.

2.5.2 Raster Scan Transformation in the Frequency Domain

Najafi-Koopai has shown [5] that the Fourier transform of a three dimensional signal can be recovered from the Fourier transform of the raster scan of that signal and vice versa.

The frequency domain transformation corresponding to the raster scan is

$$\mathbf{\Omega} = \omega \mathbf{T} \tag{2.40}$$

where Ω_1, Ω_2 and Ω_3 in rad/grain are frequency variables associated with h_1, n_2, n_3 respectively and ω in rad/s is associated with t. Thus if $R(e^{j\omega})$ is the 1-D Fourier transform of $r(t)$ and $X(e^{j\Omega_1}, e^{j\Omega_2}, e^{j\Omega_3})$ is the 3-D Fourier transform of $x(h_1, n_2, n_3)$,

[3]The frame number is fictitious in that both t and n_3 are only defined with respect to an arbitrary constant which represents the point in time at which decoding began.

then

$$R(e^{j\omega}) = X(e^{j\omega T_1}, e^{j\omega T_2}, e^{j\omega T_3}). \tag{2.41}$$

This implies that the 1-D Fourier transform of the raster scan is equivalent in value to the 3-D Fourier transform of the original signal along a line in 3-D frequency space, known as the slicing line. Because of the periodicity of the Fourier Transform in the second and third dimensions, this line fills the 3-D frequency space sufficiently to allow $X(\mathbf{\Omega})$ to be recovered from $R(\omega)$ as an interpolation [5]. This fact is used to measure the response of the filter in section 5.4. The interpolation step can be avoided by measuring the 3-D response of the filter on a discrete grid approximately aligned with the slicing line.

Chapter 3

Design of Mixed Domain LT Filters

A number of techniques have been developed for the design of 2-D and 3-D digital filters. To avoid stability problems and to take advantage of their linear phase characteristics, which are important in most image processing applications, FIR filters have been designed by optimizing the coefficients to fit an ideal magnitude frequency response [20]. However, IIR filters promise similar performance with much lower order, though stability must be ensured by design and they generally do not have linear phase. Since the number of operations required in a filter grows much more rapidly with order in a 2-D filter than in 1-D, and more rapidly yet in 3-D, this is a very important consideration, especially when real-time operation is considered. Some IIR design techniques have involved optimizing the coefficients of the transfer function to fit an ideal magnitude and/or phase response with some constraints to ensure stability [20]. Others have optimized other parameters, which have by their form, guaranteed stability [22], low sensitivity to parameter values [29] and/or wide stability margins leading to short transients [25]. However, while very selective filters can be designed, these techniques tend to require a large number of coefficients and therefore high order (> 2) filters. Also, the optimization can take a long time to perform, inhibiting real-time steering. While filters that can be simply steered can be designed with these methods [41, 42], they are generally of even higher order. To overcome these problems, algebraic techniques, usually based on continuous domain circuit models, have been advanced, and will be used here.

A number of different filter structures have been suggested for the implementation of these designs, including direct form [2,9,43], wave digital filters [21,32,44,45], differentiator-type ladder form filters [11, 12] and differential-integral operator form filters [30]. Many of these techniques could be extended to mixD filters, but, following Bertschmann, et. al. [3,4] and Najafi-Koopai [5], I will use the ladder form technique.

The initial work in 2-D mixD filters for raster scanned signals used a direct form structure [2]. The direct form structure is attractive in that the filter parameters correspond directly to the coefficients of the transfer function to be implemented, but it suffers from the same sensitivity problems that the 1-D direct form structure has, and from a very large number of operating elements [3], such as delays, integrators and multipliers. As well, stability must be ensured analytically from the transfer function.

The ladder form structure, on the other hand, has excellent sensitivity properties [11], as shown later in chapter 6, and the minimum number of operating elements. It also lends itself to algebraic design techniques that make use of well known properties of circuits to ensure stability (see chapter 7) while eliminating the optimization step in steering the passband by relating passband properties to circuit element values and thereby to filter parameters. The structure also has highly local interconnections, which makes for short signal routes in the implementation.

The ladder form design technique begins with a continuous domain prototype circuit in a ladder form, as described in section 3.1, whose properties are understood in terms of resonance, passivity, losslessness and energy. The equations that define the circuit are modeled by a signal flow graph (SFG) in section 3.2, which can then be manipulated into a form that can be easily implemented. section 3.4 describes the application of the bilinear transform, in a modified form, to the elements associated with the discrete dimensions, resulting in a replacement of the differentiation and integration operations with discrete time operations. This can lead to non-realizable delay-free loops, so a predistortion technique [11] described in section 3.3 is used on the original prototype to avoid the situations in which delay-free loops would occur.

Some final SFG manipulation may be required before implementation.

The designs of the IDD filter and the Bowl filter that result are almost exactly the same as those given in [3, 4] and [5, 12] respectively. The only difference is the modification of the bilinear transform in section 3.4. The major contribution of this thesis is the implementation described in chapter 4 which allows real-time tuning of the filter, and the precise frequency domain characterization given in chapter 5.

Other types of velocity selective filters have been designed [13, 36, 40, 46, 47], some of which offer superior performance to the filters designed here, but they do not lend themselves as well to real time implementation or control.

A serious limitation of the designs is that the component values of the prototype must be positive for stability to be guaranteed. This effectively limits the LT filter trajectories to one octant. Fortunately, a technique discovered by Bertschmann [3] involving passband manipulation by changing the length of the delay elements allows the full range of trajectories to be covered. This will be discussed in section 3.5.

3.1 Continuous Domain Prototypes

The ladder form design technique, an M-D extension of the 1-D technique (see [48]), begins with a prototype circuit of the form shown in Figure 3.1a. The extension to M-D is simply that the shunt and series impedences are, in general, M-D elements.

3.1.1 First Order Frequency Planar Filter

The purpose of a linear trajectory filter is to enhance (or remove) objects in the image moving with a specific velocity. This corresponds to enhancing (or removing) energy in the signal that lies in a specific plane in the frequency domain. The simplest LT filter is a first order system that is resonant in a plane, and corresponds to the planar pass circuit [3, 4, 9, 11, 12, 24, 45] shown in Figure 3.2, which has the transfer function

$$H_{pp}(\mathbf{s}) = \frac{R_L}{R_L + \mathbf{s}^T\mathbf{L}}. \tag{3.1}$$

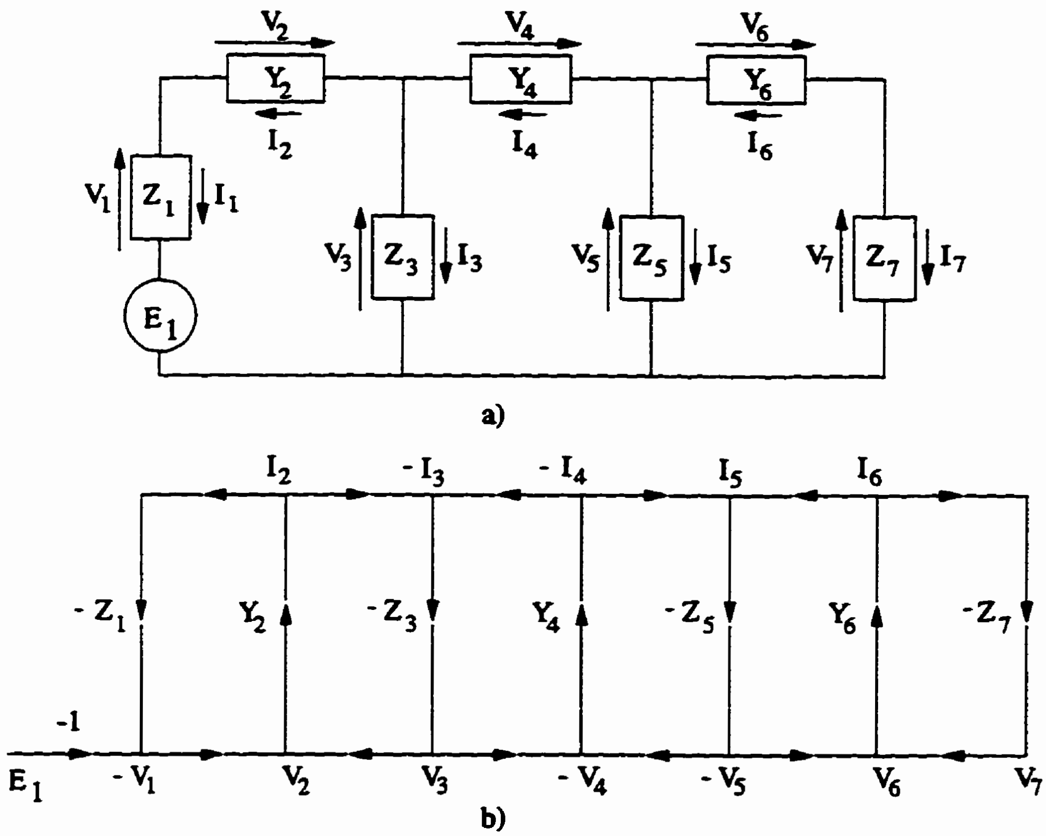

Figure 3.1: General Ladder Form Prototype Circuit

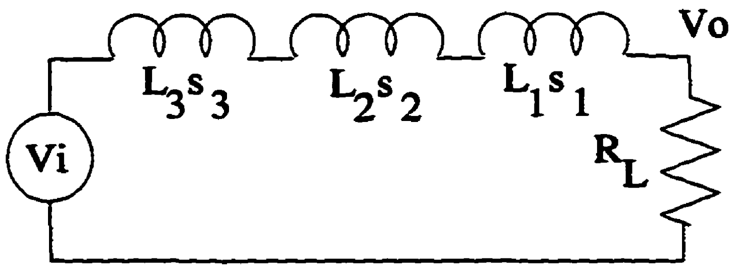

Figure 3.2: Continuous Domain First Order Frequency Planar Pass Ladder Form Filter Prototype

The 3-D inductor is resonant in the plane

$$\mathbf{L}^T \boldsymbol{\omega} = 0 \tag{3.2}$$

which has the unit normal $\mathbf{L}^T \mathbf{e}_\omega / \|\mathbf{L}\|$, where $\mathbf{e}_{\omega i}$ are the unit basis vectors in the frequency domain and $\| \bullet \|$ is the Euclidean norm. Thus it will pass LT signals with the unit trajectory $\mathbf{d}_L$, known as the filter trajectory, given by

$$\mathbf{d}_L = \frac{\mathbf{L}^T \mathbf{e}_t}{\|\mathbf{L}\|} \tag{3.3}$$

where $\mathbf{e}_t$ are the unit basis vectors in the spatio-temporal domain. Objects in the passband will then be traveling with horizontal speed, H_s, and vertical speed, V_s, given by:

$$H_s = \frac{L_1}{L_3} \qquad \text{and} \qquad V_s = \frac{L_2}{L_3}. \tag{3.4}$$

The -3dB surfaces are planes at

$$\mathbf{L}^T \boldsymbol{\omega} = \pm R_L \tag{3.5}$$

so that it has a uniform bandwidth

$$B_{pp} = \frac{2R_L}{\|\mathbf{L}\|}. \tag{3.6}$$

Now while this filter will pass undistorted any object with the filter trajectory, it has poor directional selectivity at low frequencies. That is, LT signals with signal trajectories significantly different from the filter trajectory will have some of their low

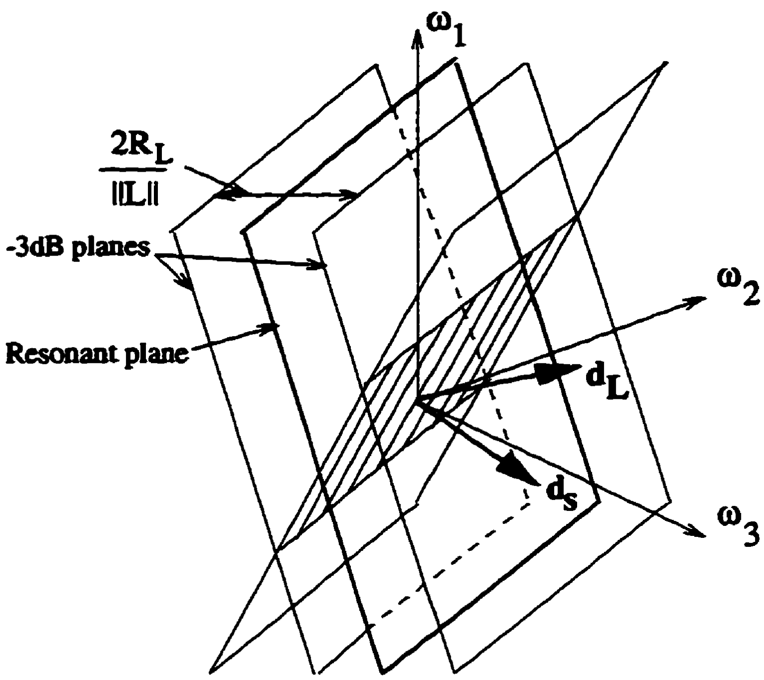

Figure 3.3: The Passband of the First Order Planar Pass Filter. Note that because the bandwidth is uniform, some of the low frequency energy in the plane shown not aligned with the filter trajectory would be in the passband.

frequency energy passed through the filter. This is shown in Figure 3.3. It is especially problematic in that a static background has all of its energy in the $\Omega_3 = 0$ plane and often consists of large regions, which have a large amount of energy in the low spatial frequencies. It has also been shown [49] that the attenuation of non-passband objects is affected by both their velocity and their shape, and that the attenuation experienced by spatially elongated objects can be quite low, even for significantly different signal and filter trajectories. The second order bowl filter is designed to address these drawbacks.

3.1.2 Second Order Frequency Bowl Filter

Ideally, an LT filter should attenuate energy in a plane uniformly throughout the plane by an amount proportional to the angle between the signal trajectory and the filter trajectory. Thus the passband would have the bowl, or cone, shape shown in Figure 3.4 [13]. This would eliminate the effect of the object's shape on the attenuation;

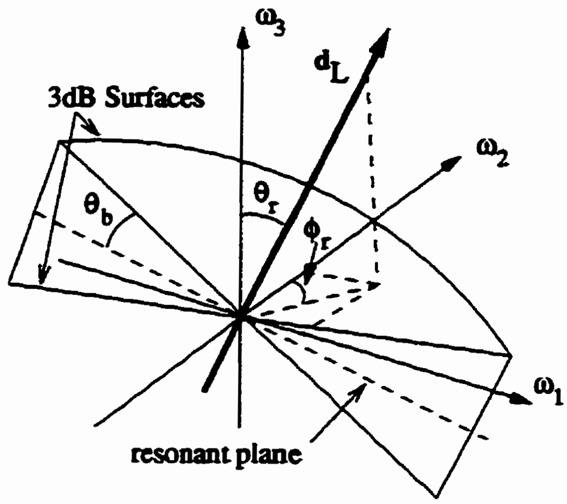

Figure 3.4: Ideal Bowl Shaped LT Passband Approximated by the Second Order Bowl Filter

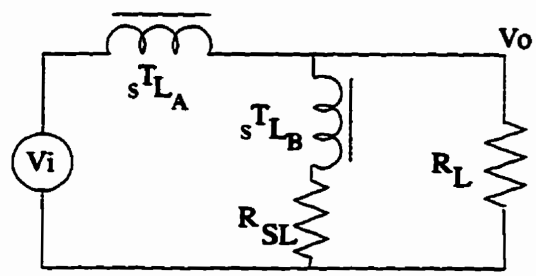

Figure 3.5: Continuous Domain Second Order Bowl Shaped Passband Ladder Form Filter Prototype

and, for filter velocities corresponding to quickly moving objects, greatly decrease the background interference.

The second order bowl filter derived from the prototype shown in Figure 3.5 approximates this shape and thus has better directional selectivity than the first order filter [12]. A discrete domain design is described in [12] and a mixD implementation in [5]. The transfer function of the prototype is

$$H_{bw}(\mathbf{s}) = \frac{R_L(R_{SL} + \mathbf{s}^T \mathbf{L}_B)}{\mathbf{s}^T \mathbf{L}_A \mathbf{s}^T \mathbf{L}_B + R_L \mathbf{s}^T \mathbf{L}_B + (R_L + R_{SL})\mathbf{s}^T \mathbf{L}_A + R_L R_{SL}}. \tag{3.7}$$

After using two different $\mathbf{s} \Rightarrow \mathbf{z}$ transforms on $\mathbf{L}_A$ and $\mathbf{L}_B$ (see section 3.4) the bandwidth in the ROI has been shown to be proportional to $\|\dot{\mathbf{\Omega}}\|$, giving the -3dB surfaces a bowl shape. Normally the inductance values are chosen such that $\mathbf{L}_B = K\mathbf{L}_A$ and the angular bandwidth is proportional to $K\|\mathbf{L}_A\|$. The constant $K > 0$ is known as the inductance ratio.

Because the input to the filter is not a continuous, unbounded domain signal, having been sampled and spatially bounded in the raster scanning process, it is not a perfectly linear trajectory signal. Also, the original input may only be piecewise linear trajectory. Thus the energy of the signal that is to be passed does not lie exactly within the central passband plane. To allow for this, the resistor R_{SL} increases the bandwidth near the origin where the bandwidth of the ideal bowl response shown in Figure 3.4 approaches zero. R_{SL} is usually set to a low value so the minimum bandwidth is

$$B_{bw\,\mathrm{min}} \approx \frac{2R_{SL}}{\|\mathbf{L}_A\|(1+K)}. \tag{3.8}$$

If R_{SL} were zero there would be a non-essential singularity of the second kind (NSSK) at $\Omega = 0$ in the mixD transfer function following bilinear transformation [12]. NSSKs have been of interest in terms of stability for some time, and have a number of useful properties [50, 51] but are generally problematic, especially if the filter coefficients are not exact [16, 18, 26, 52–54]. In this case, the stability margin is directly related to R_{SL} and therefore to the minimum bandwidth. In practice this means that the filter responses are less robust with low minimum bandwidth settings.

The resistor R_L limits the bandwidth at high frequencies, so the maximum bandwidth (outside the ROI) is

$$B_{bw\,\mathrm{max}} \approx \frac{2R_L}{\|\mathbf{L}_A\|}. \tag{3.9}$$

Note that the bandwidth of the continuous domain filter is uniform, but the two different transformations used cause the bowl shape [12].

3.2 Signal Flow Graphs

Signal flow graphs have been widely used in both 1-D and M-D digital filter design [3–5,11,21,30,48,55] and provide a powerful tool for developing easily implementable filter structures. Two general types of SFGs can be derived from the generic ladder form circuit prototype shown in Figure 3.1a [3,48]. The type I SFG shown in Figure 3.1b results from modeling the series elements as admittances and the shunt elements as impedences, and is used for both designs. The type II SFG results from the opposite choice.

The SFG treats both voltages and currents in the prototype as values to be operated on by transmittances. These transmittances exhibit a one to one correspondence to the prototype element values. The equivalence of the SFG to the circuit can be shown by applying KCL, KVL and Ohm's law to create equations relating voltages and currents in the circuit and then comparing them to the corresponding equations implied by the SFG.

The advantage of creating a SFG from the circuit is that while the characteristics of the passive circuit are maintained with the equations, the limitations of form are not. A SFG can be manipulated into a form that is simple to implement using active components and discrete delay elements.

3.2.1 First Order Frequency Planar Filter

Bertschmann showed [3] that modeling the inductor associated with the continuous dimension in Figure 3.2 as an impedence leads to continuous domain differentiators in the mixD filter, which are highly sensitive to noise and therefore undesirable. However, discrete domain differentiators resulting from modeling the other inductors as impedences can be constructed easily and robustly, so L_2 and L_3 are designated as Z_1, L_1 as Y_2 and R_L as Z_3, resulting in the SFG shown in Figure 3.6.

3.2.2 Second Order Frequency Bowl Filter

Similar arguments applied to the second order prototype of Figure 3.5 lead to choosing L_{A2} and L_{A3} as Z_1, L_{A1} as Y_2, R_L as Z_3, L_{B1} and R_{SL} as Y_4 and L_{B2} and L_{B3} as Z_5, as shown in Figure 3.7. Note that the output signal is V_3, rather than V_5 as is usual.

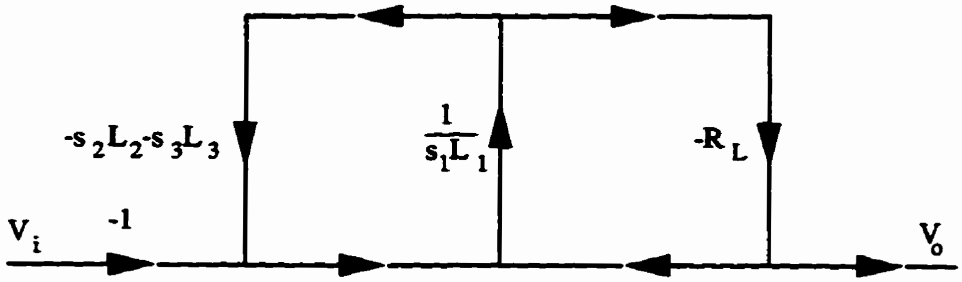

Figure 3.6: Signal Flow Graph Corresponding to the First Order Planar Pass Filter Prototype in Figure 3.2

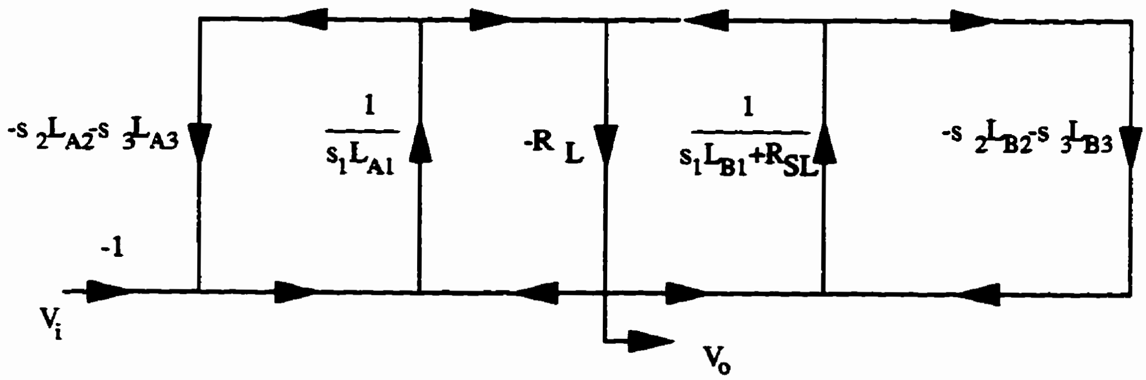

Figure 3.7: Signal Flow Graph Corresponding to the Second Order Bowl Filter Prototype in Figure 3.5

3.3 Predistortion

If the bilinear transform is applied to a branch in the continuous domain SFG with a transmittance that is proportional to s, the corresponding branch in the mixD SFG will contain a delay-free forward path. In some cases this will result in a delay-free loop (or a very short delay loop) which is not robustly implementable [3]. A technique for eliminating these loops, called predistortion, that was developed for the discrete domain case [11] has been adapted to the mixD case in [3].

The predistortion is applied to the circuit prototype before it is modeled by the SFG. A negative series resistance is associated with each inductor to be transformed and a positive series resistance of the same value with the rest of the circuit, with no net effect on the transfer function. Note that the inductors associated with the continuous dimension do not need to be predistorted. If the circuit contains capacitors, a negative and a positive conductance are inserted in parallel with the capacitor. The negative one is associated with the capacitor during bilinear transformation and the positive one with the rest of the circuit. If the correct value is chosen for the predistorting resistor, applying the bilinear transform will no longer result in delay-free forward paths in those branches.

A useful side effect of the technique is to reduce the large gains in some branches and the attenuations in others. This keeps the signal levels fairly even throughout the implementation, as well as easing the requirement for high gain-bandwidth product amplifiers.

3.3.1 First Order Frequency Planar Filter

For the first order frequency planar filter, $-r_2$ and $-r_3$ are associated with L_2 and L_3 as shown in Figure 3.8. The positive resistances r_2 and r_3 can be combined with the terminating resistor so that the number of branches in the SFG does not increase. This increases the overall gain, but does not alter the shape of the response.

The type I SFG of the predistorted circuit has the same form as Figure 3.6, but

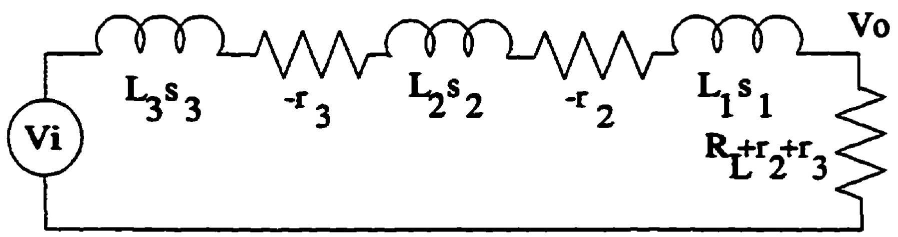

Figure 3.8: Predistorted Frequency Planar Filter Prototype

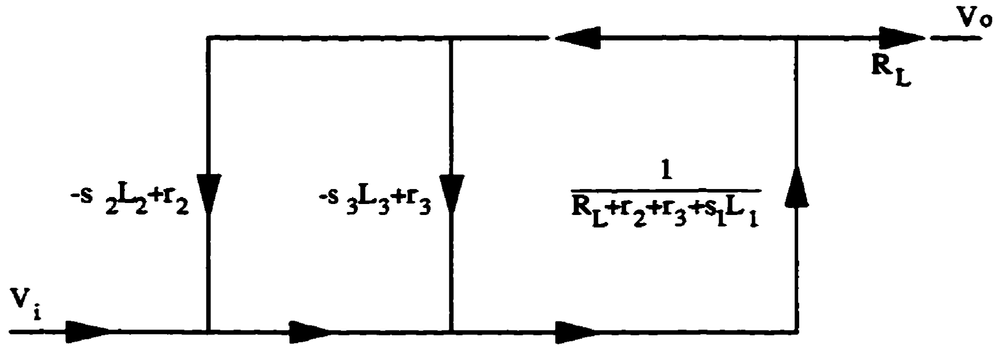

Figure 3.9: Signal Flow Graph of the Predistorted Frequency Planar Filter

the predistortion resistances change the transmittances slightly. Combining the center and right branches, and splitting the left branch in two results in the SFG shown in Figure 3.9. The right branch is suitable for implementation as a continuous domain lossy integrator, and the left two are suitable for bilinear transformation.

3.3.2 Second Order Frequency Bowl Filter

The predistorted second order frequency bowl filter is shown in Figure 3.10, with the separation into ladder form elements indicated by the dashed boxes. The circuit has been redrawn as in [5] so that the implementation is similar to two first order frequency planar sections. The SFG is shown in Figure 3.11. Breaking the branch marked "x" and routing it separately to the three summing points it leads to marked "o" results in the SFG shown in Figure 3.12. The halves on either side of the dashed line are similar in form to the first order frequency planar filter (see Figure 3.6), with either an extra input or an extra output. This will lead to an implementation in which the two filters share hardware as shown in [5].

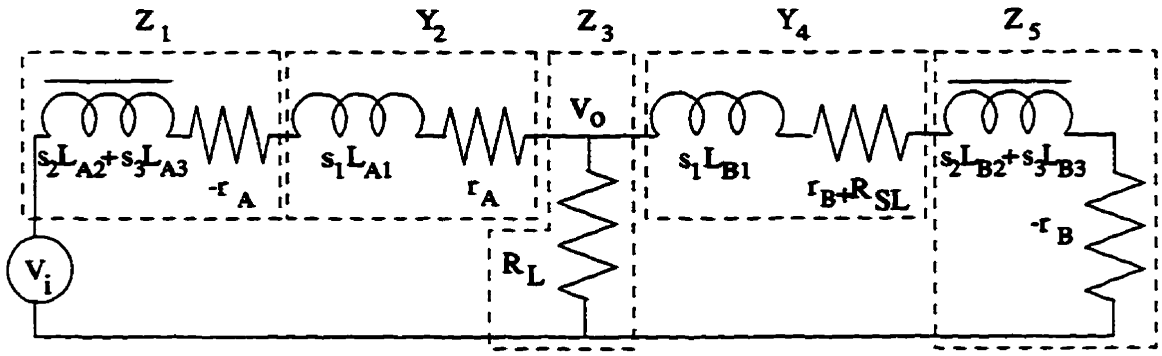

Figure 3.10: Predistorted Bowl Filter Prototype

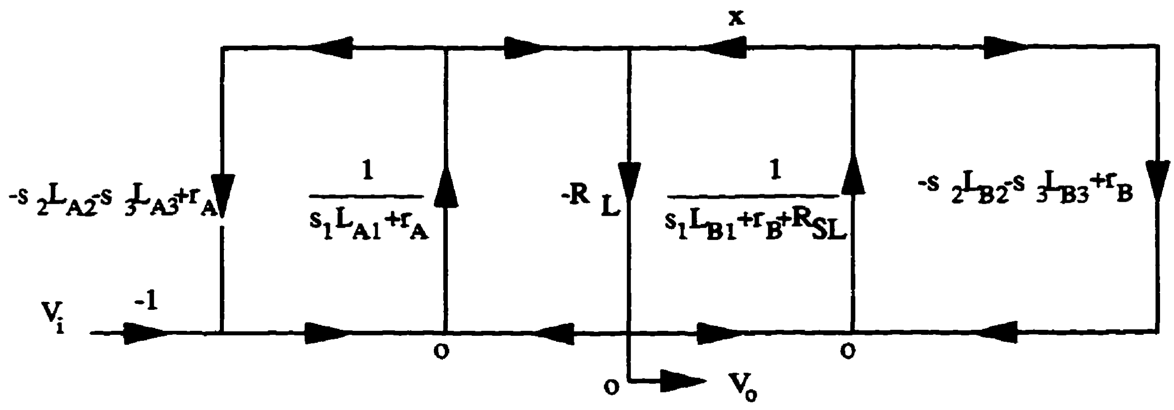

Figure 3.11: Signal Flow Graph of the Predistorted Bowl Filter

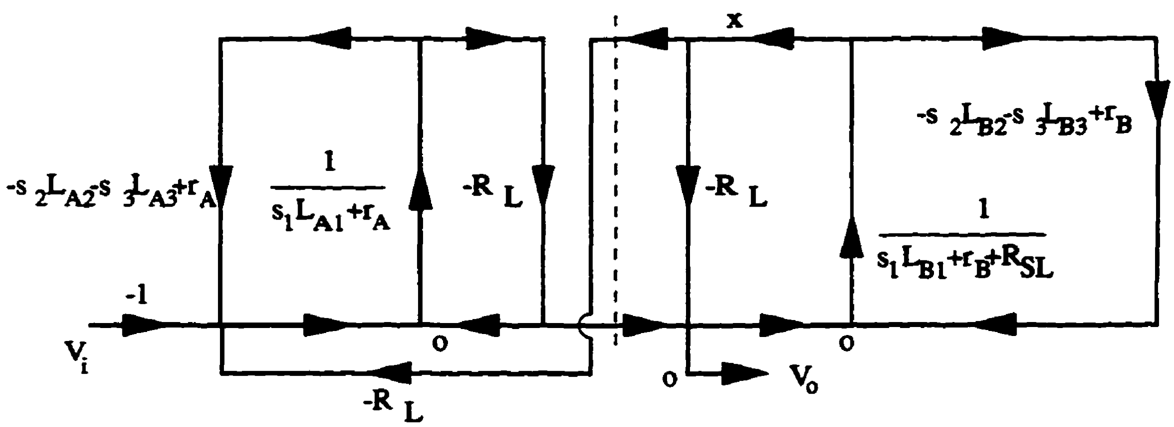

Figure 3.12: Manipulated Signal Flow Graph of the Predistorted Bowl Filter. It has been manipulated into two sections similar to the first order frequency planar sections

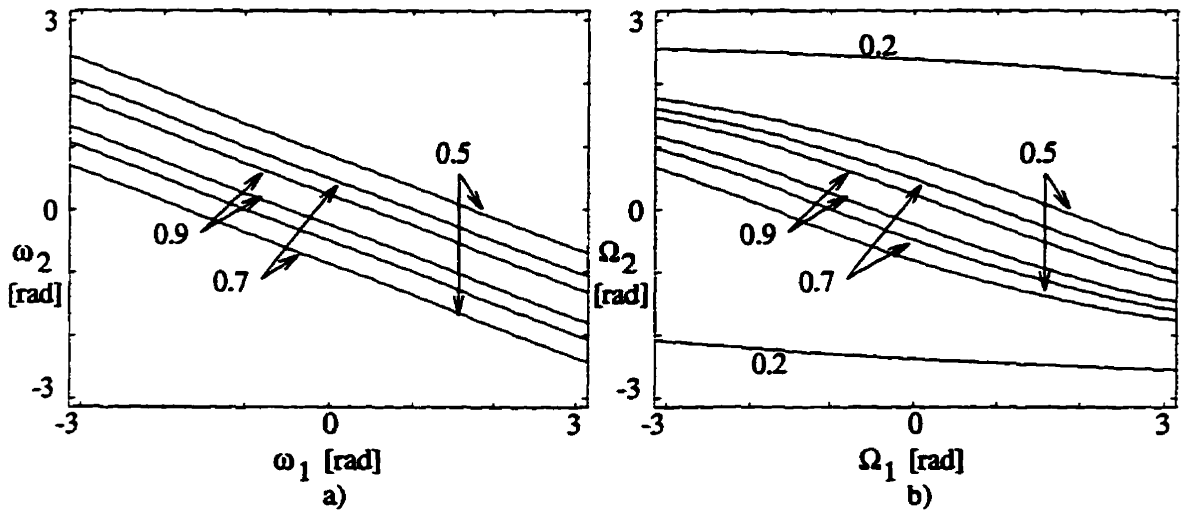

Figure 3.13: Frequency Domain Warping Effect of the Bilinear Transform for the First Order Frequency Planar Filter. Contour diagram of the magnitude frequency response for $\mathbf{L} = [1, 2, 2]$. a) The continuous domain prototype for $\omega_3 = 0$. b) The mixD prototype for $\Omega_3 = 0$.

3.4 Modified Bilinear Transform

A linear transformation is applied to the continuous domain prototype signal flow graph to produce the mixD signal flow graph. The bilinear transform, equation 2.22, is commonly used in discrete domain design and maps the s_i-plane imaginary axis onto the z_i-plane unit circle. This mapping maintains the magnitude and phase response values, but warps the frequencies at which they occur. In the M-D discrete domain case, the warping is applied to all the dimensions and is thus symmetrical, but in the mixD case it is not.

In the case of LT filters, the frequency domain warping effect causes the central passband surface to deviate from a plane; however, the warping is at its minimum within the ROI. This is illustrated in Figure 3.13, where a slice through the continuous domain magnitude response in a) is compared to the same slice through the mixD response in b). The distortion of the frequency response of both filters is discussed in depth in [12] and a method is given for compensating for the linear part of the central passband deviation by accounting for the warping in the selection of the inductor

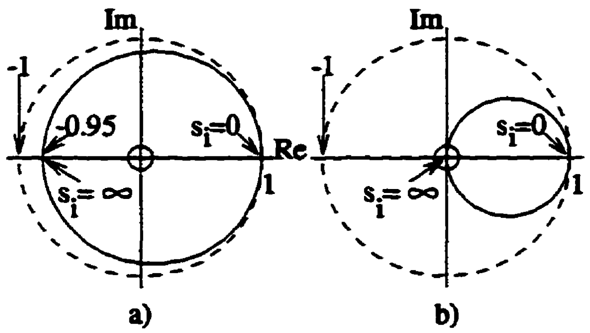

Figure 3.14: The Mapping of the Imaginary Axis in the s-plane to the z-plane of the Modified Bilinear Transform. a) $D = 0.95$ b) $D = 0$.

values. The second order bowl filter has less warping of the central passband plane within the ROI than the first order filter.

For these designs the bilinear transform is modified as

$$s_i \Rightarrow \frac{1+D}{T_i}\frac{z_i - 1}{z_i + D}, \; i = 2, 3. \tag{3.10}$$

D is used both to increase the stability margin [25], and to shape the passband in the case of the bowl filter. The increase in stability margin is important for an analogue implementation because coefficients of the transfer function are related to component values, which have tolerances that may lead to unacceptable inaccuracy in the coefficients and potential instability. For example, the filters contain discrete differentiators (see Figure 3.16) which would become unstable if the gain through a delay block or the feedback was greater than its ideal value of one. Decreasing the value of D indirectly corresponds to reducing the gain through the delay block, resulting in a more robust filter.

This transform maps the left half s_i-plane into a circle in the z_i-plane passing through the points -D+j0 and 1+j0, as shown in Figure 3.14. The polar surfaces of a passive continuous domain circuit are guaranteed to lie entirely in the non-right hand s_i-plane [25]; so, with D less than unity, all polar surfaces of the transfer function are guaranteed to be within the unit circle (except at $z = 1$). This decreases the likelihood of instability due to imprecise component values and slight non-linearities. However, because the s_i-plane imaginary axis is not mapped onto the unit circle in the z_i-plane,

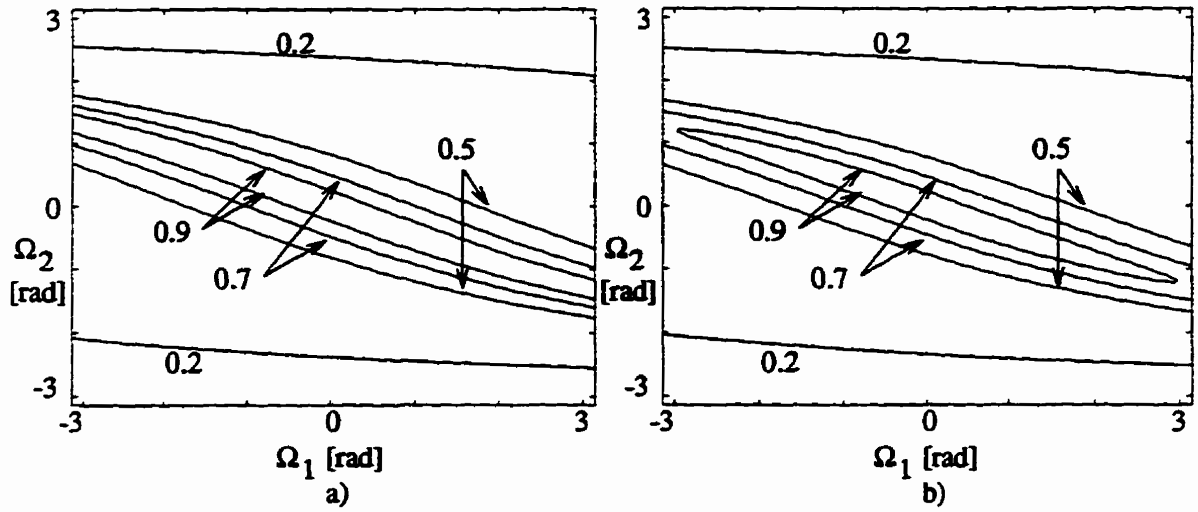

Figure 3.15: The Effect of the Modified Bilinear Transform on the Frequency Planar Filter Response. A contour diagram of a slice through the magnitude frequency response with $\mathbf{L} = [1, 2, 2]$ at $\Omega_3 = 0$. a) With $D = 1$ (unmodified). b) With $D = 0.9$.

the magnitude and phase responses are not only warped in frequency by the modified bilinear transform, but take on slightly different values. The high frequency region of the mapping in Figure 3.14 is furthest from the unit circle, so this portion of the response is most affected.

In the LT filters, decreasing D from unity causes the high frequency portion of the magnitude response to decrease, but preserves passband bandwidth and orientation, as shown in Figure 3.15. Fortunately, the effect is minimal within the ROI. As the gain of the passband at higher frequencies has been reduced from the upper bound of 1 guaranteed by passivity [29], the sensitivity of the response in the passband is no longer necessarily zero.

3.4.1 First Order Frequency Planar Filter

The modified bilinear transform is applied to the branches of the frequency planar filter SFG containing s_2 and s_3 in Figure 3.9 with $D \approx 1$ resulting in the mixD SFG shown in Figure 3.16. This is very similar to the design in [4], with a continuous domain lossy integrator (I) and two discrete domain differentiators (D) in a double loop (D),

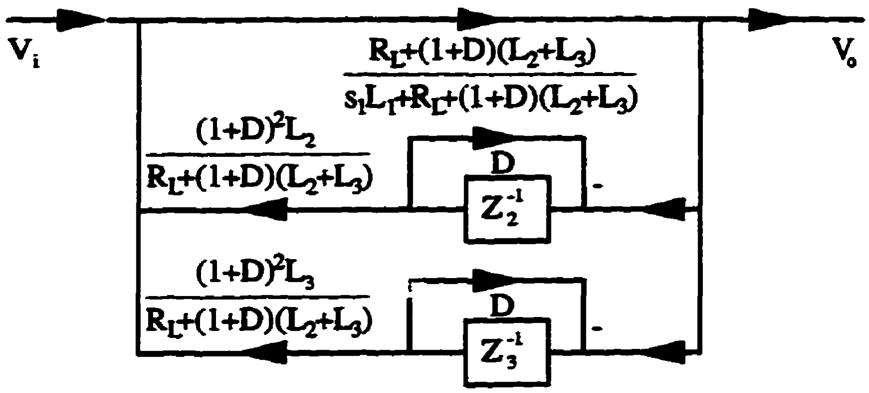

Figure 3.16: Mixed Domain Signal Flow Graph of the First Order Frequency Planar Filter, known as the IDD Filter

so it is known as an IDD filter. It contains the minimum number of delay elements and integrating elements, which is important for a low cost, low power implementation. Also, the multiplier values are very simply related to the filter trajectory, making steering simple.

3.4.2 Second Order Frequency Bowl Filter

Setting D to 0 for the transformation used on the reactances associated with L_{B2} and L_{B3} in Figure 3.12 and $D \approx 1$ for L_{A2} and L_{A3} shapes the passband of the second order filter into a bowl [12]. Careful manipulation of the SFG results in the final form shown in Figure 3.17 [5]. Again, the SFG contains the minimum number of delay elements for a second order system and the multiplier values are simply related to the filter trajectory.

3.5 Passband Manipulation by Delay Changes

A major limitation of the filters as described so far is the requirement that the inductor values be positive for stability to be guaranteed by passivity. This effectively bounds the filter trajectory to the first octant. In discrete domain implementations this is not a problem, as the input image can simply be reoriented before and after filtering so that the appropriate filter trajectory lies in the first quadrant [12, 43, 55]. However,

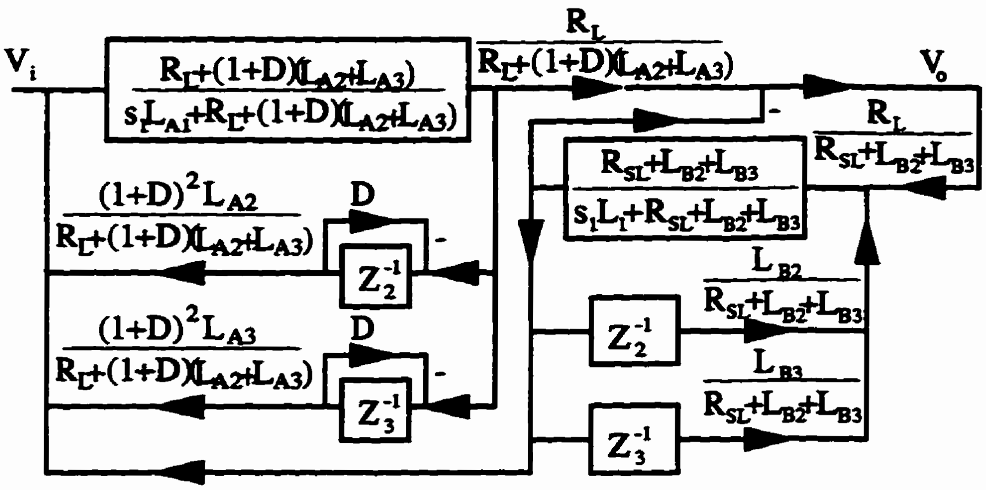

Figure 3.17: Mixed Domain Signal Flow Graph of the Second Order Frequency Bowl Filter

this is not possible in the case of a mixD implementation. A technique reported by Bertschmann [3] to manipulate the trajectory by changing the length of the delays overcomes this limitation.

The input signal to the filters is a time varying voltage—a 1-D signal—that is interpreted to be a 3-D image according to the raster scanning technique used. However, this same signal could represent a different 3-D image acquired by a different scanning technique or displayed by a different inverse scanning. In the 2-D case, using a row delay for filtering and reconstruction that is shorter than that used to scan the original image effectively shifts the rows relative to one another, as shown in Figure 3.18b. However, if a smaller delay is used in the filter but the boundary conditions are maintained, the image from the filter's point of view is as shown in Figure 3.18c. Once it has been filtered, it can be reconstructed using the original delay length.

This skew in the spatio-temporal domain corresponds to a skew in the frequency domain that allows the full range of trajectories to be covered. The skew shown in Figure 3.18 corresponds to a skew in the frequency domain as shown in Figure 3.19. In [3] and [5] the effect is described as a rotation of the passband. While this is the first order effect, treating it as a skew is both simpler and more complete.

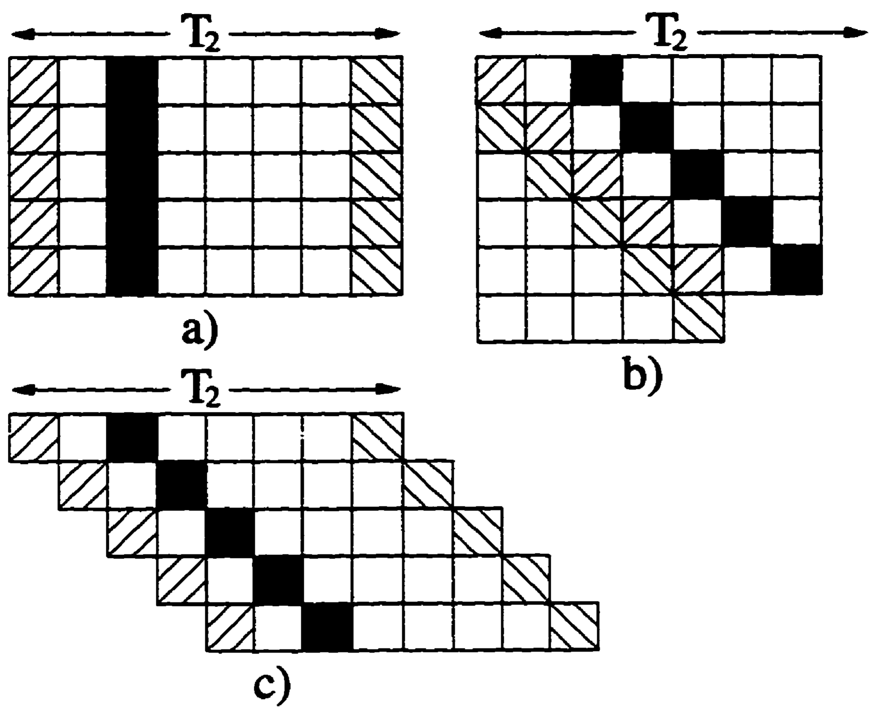

Figure 3.18: Effect of Delay Manipulation. a) with delay of T_2 (the original) b) with delay less than T_2 c) with delay less than T_2 but the original boundary conditions.

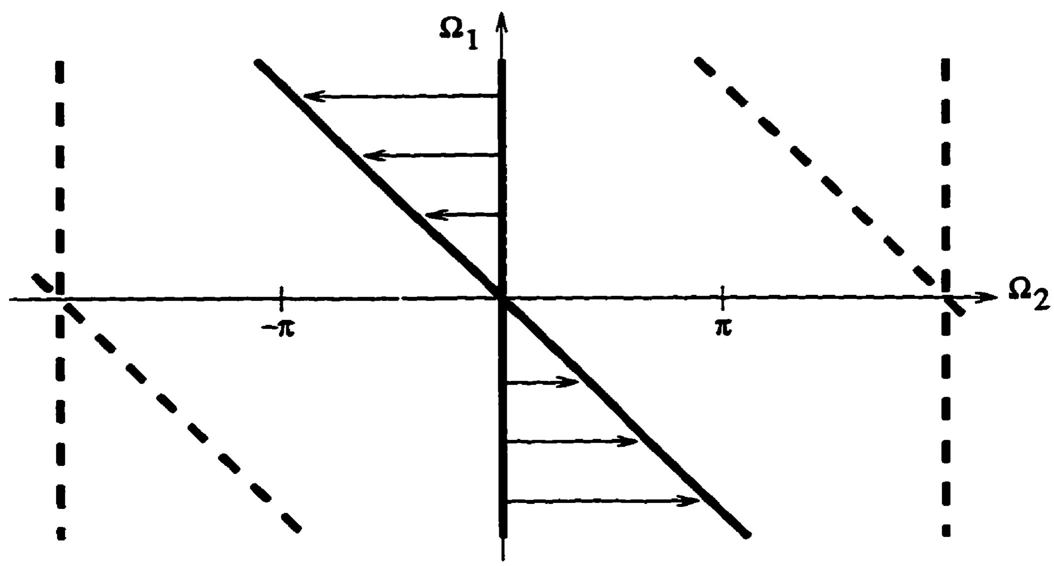

Figure 3.19: Skew in the Frequency Domain Corresponding to the Skew in the Spatio-Temporal Domain shown in Figure 3.18. Note: not to scale.

In a 2-D LT filter, the passband direction can be described in terms of slope. That is, in terms of horizontal displacement of the filter trajectory for a unit vertical displacement. Reducing the row delay T_2 by $\Delta T_2 = m_1 T_1$ $(m_1 \in \mathcal{R})$ will reduce the effective horizontal displacement of the filter passband by m_1 grains per vertical grain. This corresponds to increasing the effective slope of all input signals by the same amount.

In a 3-D LT filter, the relationship can be expressed as:

$$\begin{aligned} \Delta T_3 &= m_1 T_1 + m_2 T_2 \\ \Delta H_s &= m_1,\ m_1 \in \mathcal{R} \\ \Delta V_s &= m_2,\ m_2 \in \mathcal{Z} \end{aligned} \tag{3.11}$$

so increasing the frame delay T_3 by $\Delta T_3 = m_1 T_1 + m_2 T_2$ will increase the effective horizontal speed associated with the filter trajectory by m_1 grains/frame and the effective vertical speed[1] by m_2 grains/frame. In the frequency domain, the mapping $\Omega_3 \Rightarrow \Omega_3 + m_1\omega_1 + m_2\Omega_2$ is applied. If m_2 is zero, this corresponds to shifting the planes of constant ω_1 by $m_1\omega_1$ in the Ω_3 direction. If m_1 is zero, it corresponds to shifting the planes of constant Ω_2 by $m_2\Omega_2$ in the Ω_3 direction. A passband plane will be tilted, and stretched, around the opposite axis, though the action is not a rotation, as a passband that lies entirely in the $\omega_1 = 0$ plane is not affected by the first case and vice versa. If neither m_1 nor m_2 is zero, the mapping corresponds to shifting lines of constant ω_1 and Ω_2 by $m_1\omega_1 + m_2\Omega_2$ in the Ω_3 direction. Thus a passband plane will be tilted about the line $m_1\omega_1 + m_2\Omega_2 = 0$.

Changing the length of the row delay will have the same effect as in the 2-D case, except that planes of constant ω_1 will be shifted in the Ω_2 direction rather than lines. This corresponds to changing the shape of the input signals, rather than changing the direction of motion, and may be of use in filtering elongated objects [49].

Because the filter response is periodic in the second and third dimensions, the skew in these dimensions will eventually lead to aliasing of the passband, as can be seen from Figure 3.19. The passband from the next period of the response will be

[1]Positive vertical speed is downward in NTSC raster scan video.

skewed into the baseband response. Fortunately, for LT filters this aliasing occurs mostly outside the ROI, and is therefore assumed to have little effect. An examination of this is still necessary.

Another concern that has not been addressed before is whether the filter will remain stable for a given change in the delay lengths. This issue is addressed in chapter 7.

Chapter 4

Hardware Implementation of the Signal Flow Graphs

The advantages and disadvantages of the mixD approach to multidimensional filtering compared to the digital approach are largely in the implementation. One of the primary purposes in the development of the mixD approach is a reduction in the size and cost of the hardware. While this has been achieved, the mixD implementation is not nearly as flexible as the digital implementations can be. One major advantage that digital filters have over the mixD filters previously developed is in the control of the filter parameters, such as multiplier coefficients corresponding to transmittances in the SFG.

The goal of the research work described in this chapter was to develop a hardware platform implementing two 3-D mixD filters for raster scanned video signals in which the filter parameters are controllable interactively and precisely while the filter is running. The filters were to be more robust than those previously developed, as well as easy to use and flexible enough to be used in further research. Also, the hardware platform was to be designed to be expandible, specifically for the addition of real-time adaptive control. The two filter structures that have been implemented are the first order LT, or IDD, filter shown in Figure 3.16 and the second order Bowl filter shown in Figure 3.17.

The filters, which together with the control circuitry and software are known as the *Endevour Analogue Video Filter*, are constructed from readily available integrated circuits and discrete components such as resistors and capacitors mounted on printed circuit boards. Also mounted on these boards are interfaces to a standard NTSC format video camera, as a signal source, an NTSC format video monitor, as a signal sink, and a workstation, for control through a graphical user interface (GUI). This implementation, which meets the goals outlined in the previous paragraph, is a principal contribution of this thesis.

The only appropriate comparison between a fully digital implementation and a mixD implementation of a controllable real-time 3-D video filter, of those reported in the literature, is between the Challenger real-time video processor [1, 55] and the Endeavour analogue video filter described in this thesis. Challenger is a fully digital, discrete domain implementation of a general first order 3-D transfer function, developed recently by C. Kulach, et. al. [1] in the Micronet MDDSP research group at the University of Calgary. It is fully controllable in real time through the same kind of GUI as the Endeavour and has built in video extraction and reconstruction A/D and D/A capability, enabling it to input and output NTSC raster scan video. It contains 634 integrated circuits (ICs) on 17 8" x 10" printed circuit boards (PCBs) and draws an average of about 15A of current at 5V [1]. Endeavour is a mixD implementation of the first and second order LT filters designed in chapter 3, using digital signal paths only in the delay elements (see section 4.5). It is also fully controllable in real-time and has NTSC raster scan video inputs and outputs. It contains 279 ICs on 6 8" x 10" PCBs and draws a total of approximately 8A at $\pm$5V and $\pm$12V. While the comparison is not between two implementations of the same filter, the reduction in size, cost and power consumption due to the mixD technique is readily apparent.

It is also appropriate to mention that the non-controllable IDD filter constructed by K. Bertschmann and N. Bartley [4] uses only 96 ICs and draws about 2A at 5V. Much of the increase between this filter and Endeavour is the control interface.

Other issues in the comparison include precision, non-linearities and calibration. The internal precision, or dynamic range, of digital systems can be improved by increasing the internal wordlength, to a limit imposed by the speed of multipliers of that wordlength. High quality analogue hardware, however, is limited by noise and slew rates to a dynamic range comparable to a 9 bit wordlength in a digital system. This has proven adequate in many filters, but some very narrow bandwidth filters may require 13 bits precision [55] and are not, therefore, achievable in mixD. Both implementations experience overflow non-linearities and either finite precision effects or non-linear component effects. One consideration that appears in the control of mixD filters that is not a problem in digital filters is calibration. Because the filter parameters are dependent on component values, which are not exact, the designer must measure the response of the circuit to determine these values exactly and thereby calibrate the control system. These effects are discussed further in section 5.3.

4.1 Top Level Layout

The physical layout of the hardware is modular, both to ease testability of the various functions and for future expansion. The 5 modules are connected by a 150 conductor backplane which has space for up to 3 more modules, and the entire system is mounted in a vented, fan-cooled enclosure. The five modules are: the analogue board, the row delay, the frame delay, the controller and the power supply. The analogue board contains video extraction and reconstruction sections, described in section 4.2; the actual analogue filter block, described in section 4.3; and the coefficient D/A subsection, described in section 4.4. The row delay and frame delay are described in section 4.5 and provide z_2^{-1} and z_3^{-1} operations respectively. The controller board interfaces with a workstation to provide GUI control of various parameters including multiplier coefficients and filter structure. It is described in section 4.6. The power supply provides +5V at 7A and -5V, +12V and -12V at 1.5A. A block diagram of the analogue signal path and the elements that operate on it is shown in Figure 4.1.

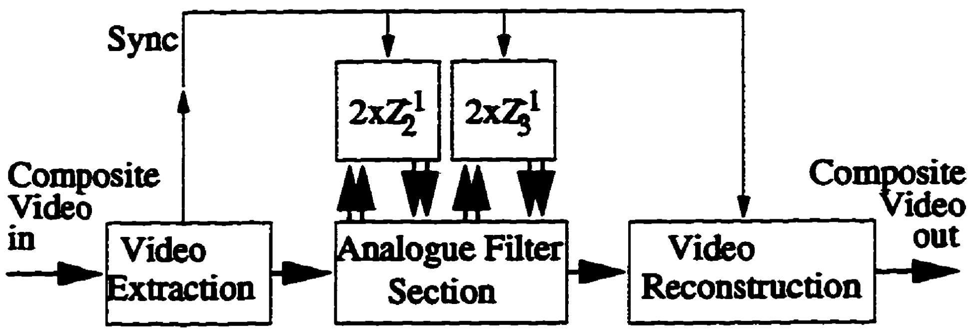

Figure 4.1: Top Level Block Diagram of the Analogue Signal Path

4.2 Video Extraction and Reconstruction

The video extraction circuit used on the analogue board (see Figure 4.1) is nearly identical to that presented in [3] and is shown in Figure 4.2. It provides three basic functions: DC level restoration of the AC coupled composite input signal, sync pulse removal to provide a clean active video signal and sync pulse detection and extraction. The LM1881 video sync detector [15] extracts the composite sync signal and passes it to the delay elements and the video reconstruction circuit. It also provides a burst signal during the back porch portion of the scan line (see Figure 2.2) closing the 4066 analogue switch and clamping the zero level of the op-amp to the back porch level. The emitter follower circuit then eliminates any traversal above zero, effectively eliminating the sync pulses. The active video is then applied directly to the analogue filter block.

The option exists to use a separate external synchronization signal, which is then used for all synchronization in the system. The composite sync signal, odd/even field signal and vertical sync signals are made available externally. It is also possible to bypass the extraction block and apply an input signal directly to the analogue filter block. This is useful in testing, especially for applying an artificial input signal as in the frequency response measurement (see chapter 5).

The video extraction block has a dominant pole response with a cutoff frequency of 2.3 MHz, which is necessary to remove impulse noise which can cause the analogue multipliers or the delay elements to overflow. 2.3 MHz corresponds to

$\omega_1 > 2.02$ rad/grain, which is well outside of the ROI. The problem of impulse noise is discussed further in section 5.2.

Once the active video signal has been filtered, the composite video output signal is created by the video reconstruction block, shown in Figure 4.3. The active video signal is added to the synchronization signal only during the active video portion of the line, providing clean sync pulses. The 2N3906 transistor acts as an AND gate for the Control line and the inverted sync signal and provides the increase in voltage swing from 0–+5V to -5V–+5V needed to control the 4066 analogue switch. The Control line must be high and the inverted sync signal low for the switch to be on. The Control line is one bit in a register that is set by the control board which selects the filter structure. There are actually two control lines, transistors, switches and active video sources: one each for the output of the IDD/Bowl filter or the temporal highpass filter discussed in section 4.7. If the Add Sync switch is moved to the off position, the sync pulses are not added to the output and the 4066 switch remains closed through the sync interval. This is useful for testing purposes, as described in chapter 5.

4.3 Analogue Filter Block

The SFGs of the IDD and Bowl filters shown in Figures 3.16 and 3.17 contain three types of transmittances. Three branches contain s_1 in the form $R_L/(R_L + s_1L_1)$, and are called lossy integrators. An example is the top most branch in Figure 3.16. Four branches, including the other two branches in Figure 3.16, contain a discrete differentiator, which is a delay element with negative feedback around it. The other branches are simply multiplications, some with a delay; and all three types require summation at either input or output. The SFGs are repeated in Figures 4.4 and 4.5 where the transmittances are marked I for lossy integrator, D for discrete differentiator and M for multiplier.

Previous implementations of mixD filters [3–5] have used operational amplifier

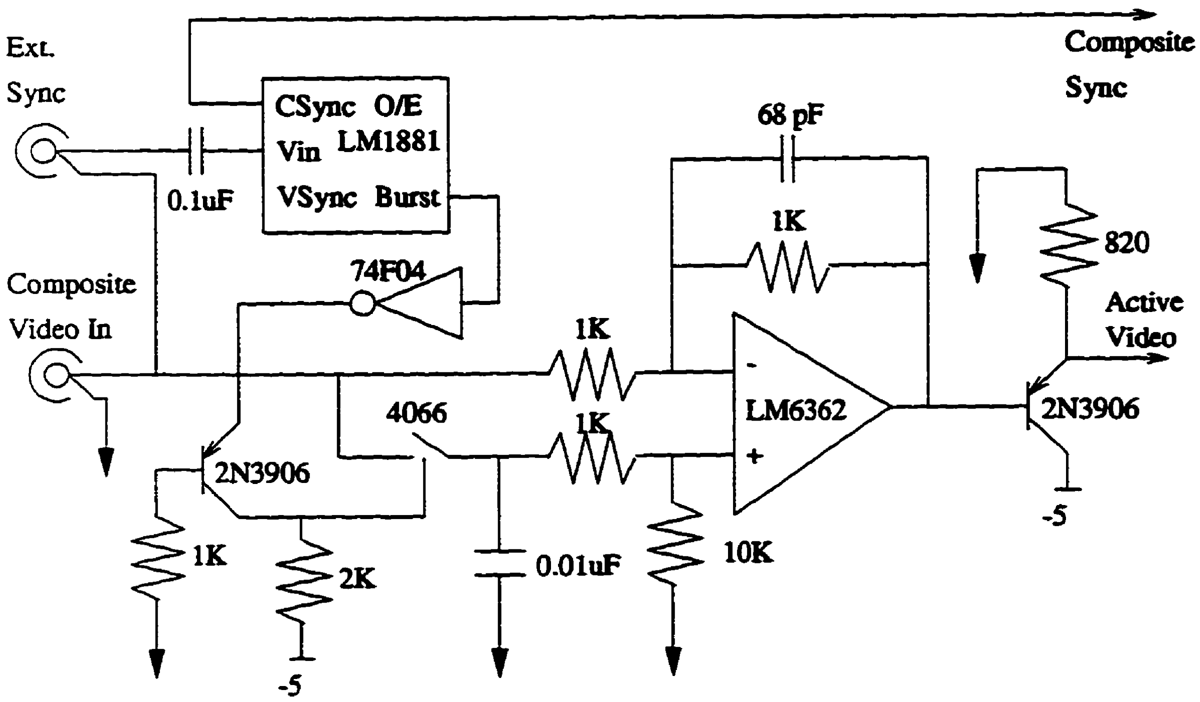

Figure 4.2: Video Extraction Circuit

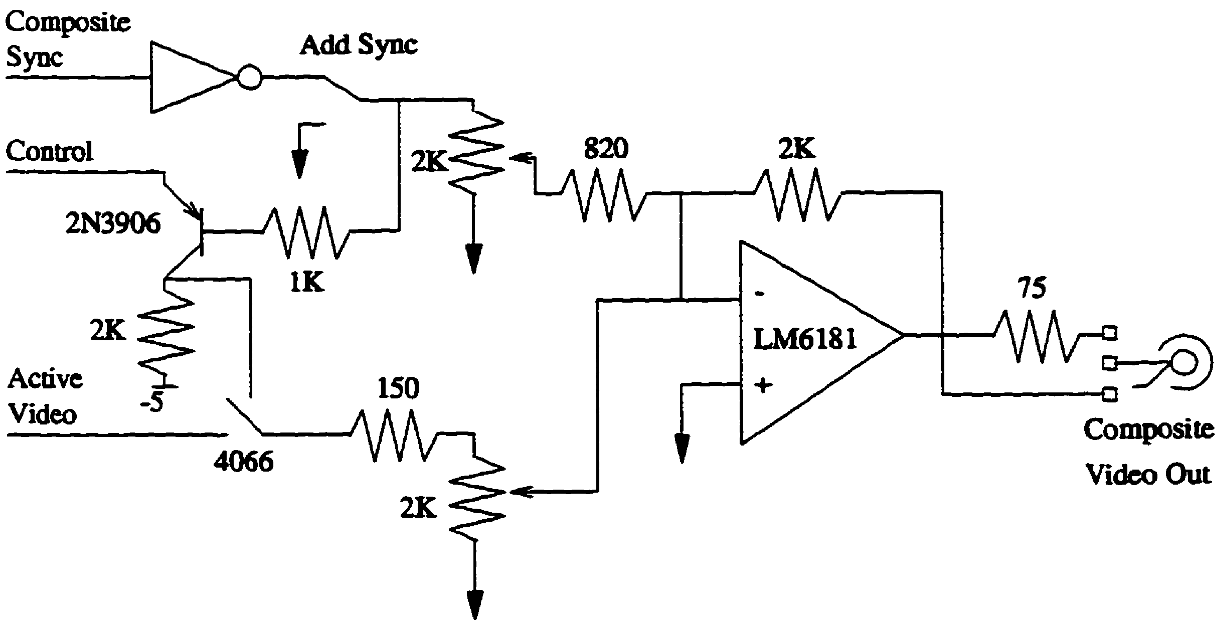

Figure 4.3: Video Reconstruction Circuit

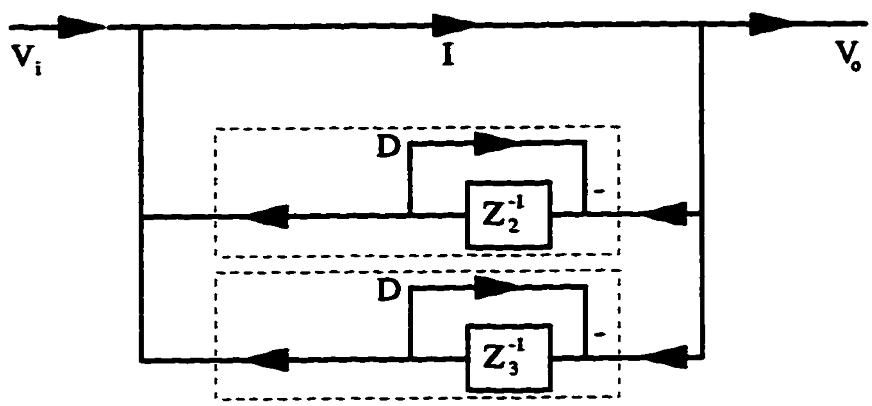

Figure 4.4: Signal Flow Graph of the IDD Filter Reduced to Components: (I) Lossy Integrators, (D) Discrete Differentiators, (M) Multipliers.

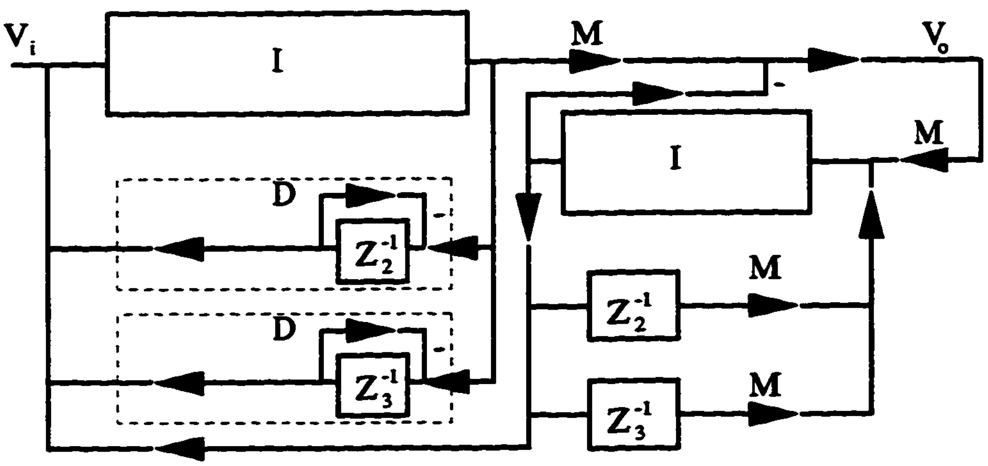

Figure 4.5: Signal Flow Graph of the Bowl Filter Reduced to Components: (I) Lossy Integrators, (D) Discrete Differentiators, (M) Multipliers.

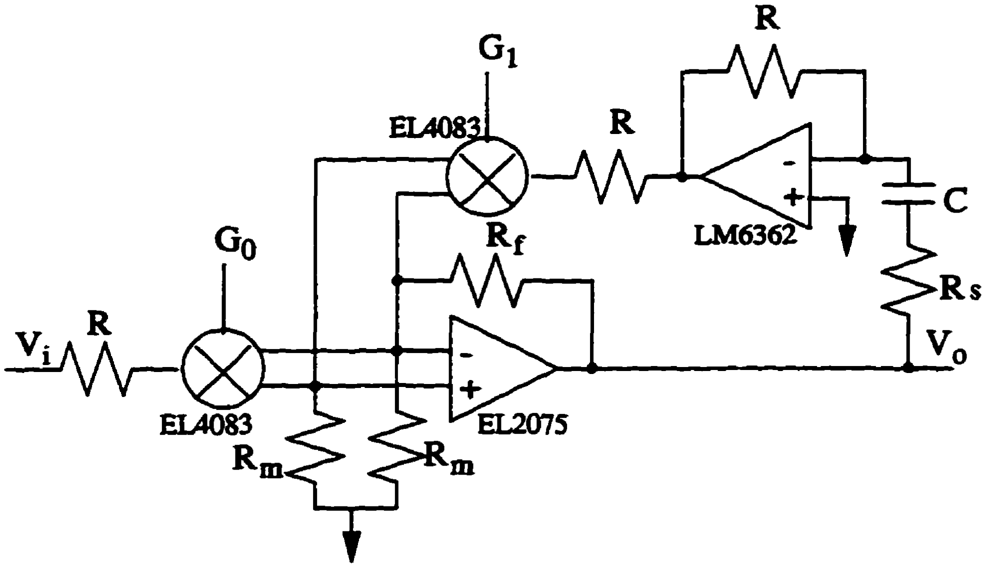

Figure 4.6: A Variable Multiplier and Variable Lossy Integrator

circuits to compute the multiplications, summations and lossy continuous domain integrations. The transmittance values have been proportional to ratios of resistors and capacitors, which could be changed using potentiometers. The circuits are simple and work well, but are laborious to tune.

In order to make the transmittances proportional to currents, which can be controlled externally, the multiplications, summations and lossy continuous domain integrations are implemented with high bandwidth, current mode analogue multipliers and high gain bandwidth product, voltage mode operational amplifiers. The bandwidth of the lossy integrator is controlled through an analogue multiplier in the feedback path with the capacitor, as shown in Figure 4.6. Resistor R_s is small and serves to stabilize the circuit. Currents G_0 and G_1 are supplied by D/A converters in the coefficient control block (see section 4.4).

Comparing the structure in Figure 4.6 to the SFGs in figures 4.4 and 4.5, the two op-amps and the feedback multiplier implement the lossy integrations in the form $R_L/(R_L+s_1L_1)$, where L_1 is proportional to G_1R_fC. The input gain is proportional to G_0 and the summing point is differential over the two resistors R_m. More multiplier outputs can be summed by connecting their outputs directly to these points, and

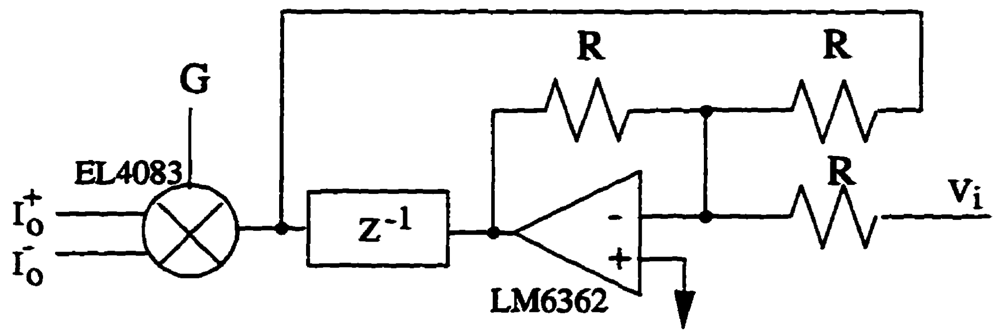

Figure 4.7: A Discrete Domain Differentiator

inversion is performed simply by exchanging the positive and negative outputs.

The op-amp in the feedback path is required because the analogue multiplier is unstable with a capacitor connected to its input. However, this introduces a delay in the feedback path leading to a second (non-dominant) pole. The second pole limits the range of cutoff frequencies for which the response is a good approximation of a first order lossy integrator. This effectively limits the range of L_1.

The discrete domain differentiation is computed with a circuit of the form shown in Figure 4.7, where the block labeled z^{-1} indicates an analogue row or frame delay element. The input to this block is a voltage, V_i, as output by the summing op-amp in Figure 4.6, and the output is a differential current, $I_o^{\pm}$, suitable for summation. A small compensation capacitor is added in parallel to each resistor to stabilize the op-amp.

The parameter D in the modified bilinear transform, equation 3.10, is actually provided by a variation of the gain through the delay block, which is set by potentiometers. It is then compensated for in the calibration of the multiplier control values. The exact gain through the delay element is important because if it is set too high the filter will be unstable but if it is set too low the filter will have a poor high frequency response. By trial and error, the best trade off was found to be a gain of approximately 0.9.

The final circuit design is constructed from the two basic blocks in figures 4.6 and 4.7 according to the connections defined by the SFGs. An overall block diagram is

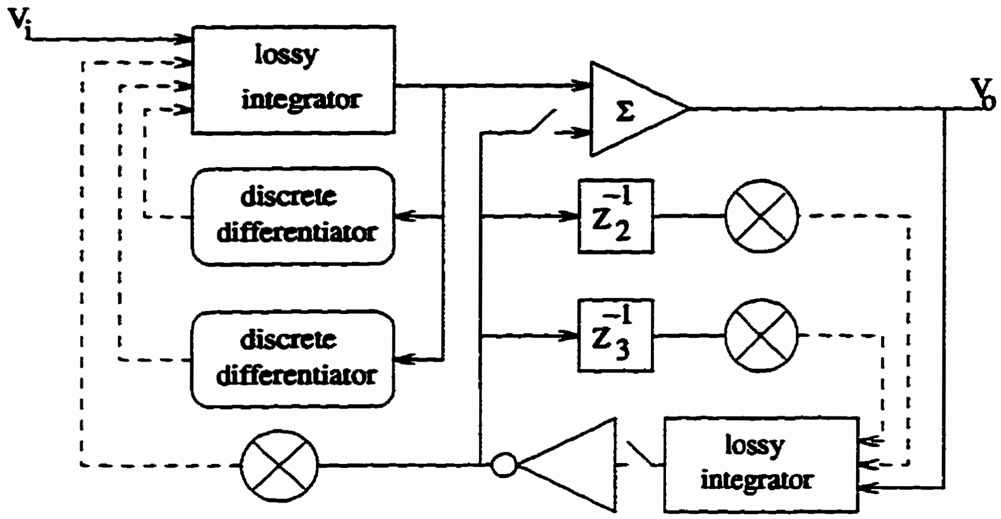

Figure 4.8: Block Diagram of the Complete Analogue Filter Block Implementing both the IDD Filter and the Bowl Filter

shown in Figure 4.8 where voltage signal paths are shown as solid lines and differential current signal paths as dashed lines. Besides the aforementioned blocks, the circuit contains one summing op-amp and one inverting op-amp. The inverting op-amp could be avoided topologically, but is used as a buffer, as the signals going to the delay elements travel across the backplane. The IDD filter structure is selected when the analogue switches are open and the Bowl filter structure when the switches are closed, allowing most of the components of the IDD filter to be shared by the Bowl filter. The state of the switches is controlled by registers set by the control block and accessible from the GUI. In total, the analogue filter block contains 8 op-amps and 9 analogue multipliers and accesses 2 row delays and 2 frame delays.

4.4 Coefficient Control Block

The filter coefficients are supplied as a current to one input of each multiplier. A D/A converter supplies a buffered voltage which is converted to a current by a resistor connected to the multiplier input, which is a virtual ground. The value of the voltage is directly related to the binary sequence stored in the D/A converter by the control

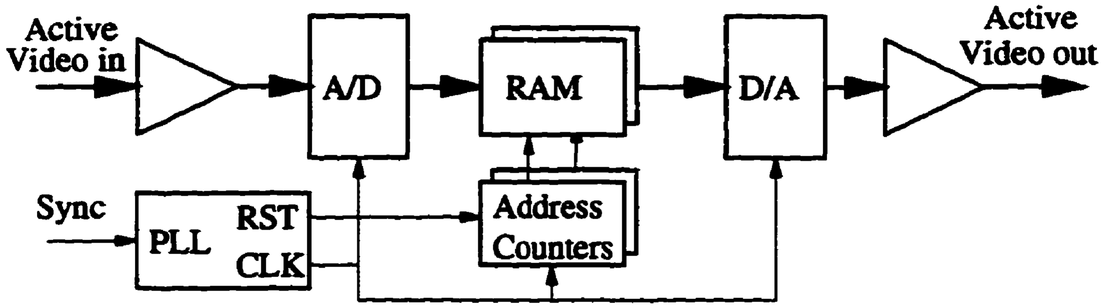

Figure 4.9: Functional Diagram of Delay Elements

block, and is directly accessible from the GUI. While the multipliers are capable of four quadrant multiplication, the requirement that the inductor values be positive means that only positive currents need be supplied. Thus the output voltage from the D/A converters ranges from 0 to 1V, though the maximum current supplied differs between multipliers according to the resistor value used for conversion.

The video extraction and reconstruction block, the analogue filter block and the coefficient control block are laid out on one 8"x10" PCB.

4.5 Delay Elements

The delay elements z_2^{-1} and z_3^{-1} would, ideally, provide an analogue, continuous domain continuously controllable delay of nominally one row length T_2 or one frame length T_3 respectively. Because the length of the delay affects the shape of the passband, and especially the filter trajectory, the length of the delays must be precisely determined; and, because this effect is to be used to tune the filter, they must be controllable in real-time.

The current implementation, based on a design by N. Bartley [56] shown in Figure 4.9, uses a high sample rate A/D converter, digital storage and D/A converter to approximate this. The input and output of the delays are analogue, but the internal signal path is digital at 8 bits quantization, which is nearly the dynamic range of the other analogue components in the system. The input and output of the delays are also continuous time, while the digital signal path is discrete time. Sampling is done at 768 samples per row during the active video for the row delay and 384 samples per row during the entire row length for the frame delay, for a sample rate of 2.4 and 1 sample per grain respectively. The Nyquist frequency is then π rad/grain for the frame delay, which is well above the ROI limit of $\pi/2$. The digital storage uses a double buffer approach with separate counters for reading and writing, allowing delay skews of any number of samples. The clock is generated by a phase locked loop referenced to the synchronization signal. The frame delay length is determined from the vertical sync pulse and the row delay length from the horizontal sync pulse, which matches the filter delays exactly to the timing of the camera applying the raster scan. A total of 178 ICs on 4 8" x 10" printed circuit boards are used to create 2 row delays and 2 frame delays. This is currently the largest component of the system and the major cost, so improvements and alternate implementations are of interest. They are further discussed in section 4.7.

4.6 Controller and User Interface

To make the control of the filter parameters easy a communication system and graphical user interface has been developed for the filter. A functional diagram of the system, which has been adapted from a similar system developed for the Challenger RTVP [55] with the assistance of its designer, C. Kulach, is shown in Figure 4.10. The coefficients, structure settings and delay lengths are stored in registers in the appropriate filter modules, which are set via the backplane by the controller module. The microprocessor on the controller module communicates via a RS-232 serial connection with a Unix daemon running on a workstation, which in turn communicates with a GUI via Unix interprocess communication. The entire system is designed to be transparent to the user.

The filter module contains 10 8 bit D/A converters which supply currents to the analogue multipliers that represent the filter coefficients. Each D/A has a register and a 4Kx8 bit look up table associated with it. Each register or look up table can be

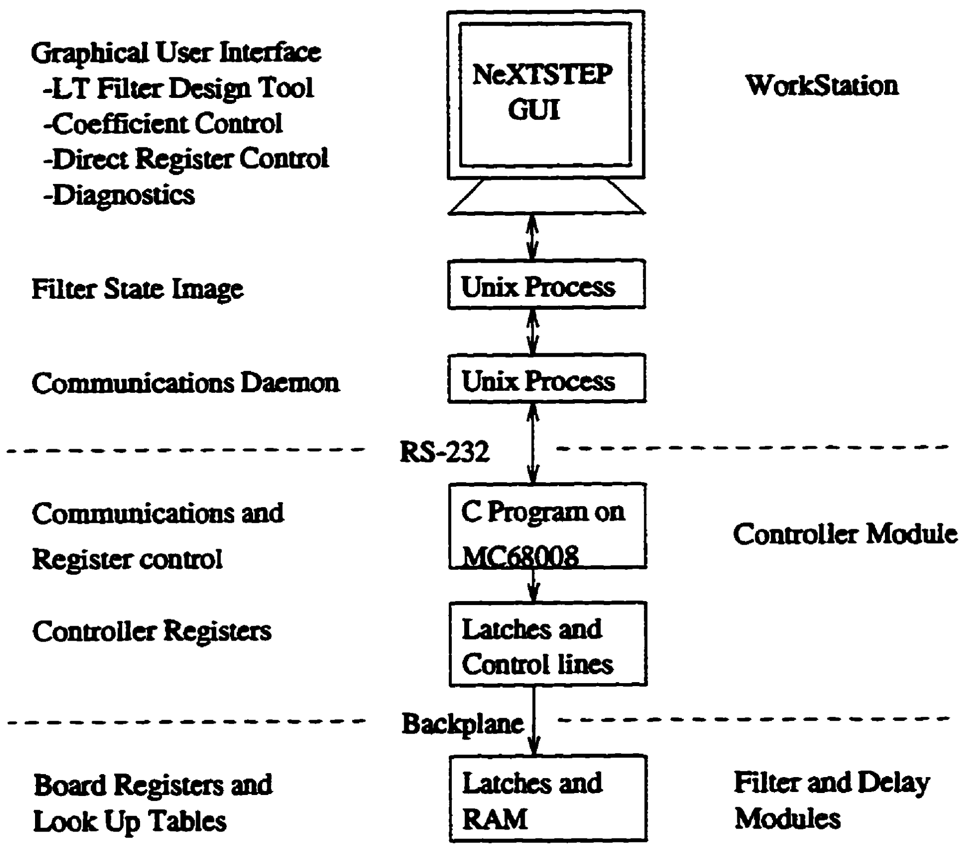

Figure 4.10: The Communication and User Interface for the Endeavour Analogue Video Filter (adapted from the Challenger RTVP)

loaded individually from the controller or all 10 registers can be loaded in parallel from their look up table at a common address. This allows many different filter settings to be stored and applied quickly, as in a tracking application. The filter module also uses 3 bits of a control register to set the filter structure.

The delay modules contain either two row delays or two frame delays. The starting addresses for the two address counters (see Figure 4.9) are stored in two 16 bit registers, each of which is associated with a 4Kx16 bit look up table. The difference between the start addresses for reading and writing compensates for latency through the A/D and D/A converters and allows the length of the delay to be controlled. Since the two delays on each board share address counters, the two different delay elements are of the same length. If they were of different lengths the analysis of the delay change effects in section 3.5 would be much more difficult. Also two bits of a control register can be used to put the module into a diagnostic mode in which either one or the other buffer is continually read out, or else the signal bypasses the storage entirely, running straight through with only the converter latency as a delay.

The controller module has been adopted complete from the Challenger RTVP system [55]. It is based on a Motorola 68008 microprocessor, with a RS-232 serial data interface and a number of registers and control lines driving lines on the backplane. The RS-232 communications part of the program running on the microprocessor, written entirely in C and assembly with no library functions, has also been adapted from the Challenger project, but the code that interacts with the filter and delay modules was all developed for this project.

Communication between the controller and the other modules occurs over several shared busses on the backplane. Each module other than the controller has a unique 8 bit address. When the controller receives a request to load a register on a module it puts that module's address on the board address bus, to which the module replies by identifying its type on the 3 bit identity bus. Then the controller puts a 16 bit value on the coefficient bus and clocks it into either an address register for the look up tables or the module's control register, which selects the next data register to be loaded. Finally it puts the data on the coefficient bus and clocks it into the appropriate data register. Besides responding to individual requests for state changes (which take precedence) the controller has a batch processing ability. A command file containing a list of operations to be done, with an optional pause and/or loop, can be loaded from the GUI. This has been useful for testing the operation of inter board communication, registers and even board functionality. The controller can also operate independently of the workstation using interrupts or an 8 bit bus that can be connected to buttons to load predefined filters.

To avoid blocking the GUI while communication on the RS-232 serial line occurs, a Unix daemon acts as a buffer on the serial port, running asynchronously with the GUI and communicating with it via Unix sockets. This daemon has been adopted complete from the Challenger RTVP system [55].

The GUI is based on an image of the filter states. The main controlling object is responsible for maintaining this record of the state of the filter, updating the filter by forwarding requests from panel objects to the communications daemon and informing the various panel objects of changes made in other panels. It is also capable of saving the state to a file for use in another session. A number of different panels offer different views on this filter state and different ways to manipulate it.

The direct register control panels allow the user to manipulate the bits in the registers directly, without any interpretation. The filter module control panel is shown in Figure 4.11 and a delay module control panel is shown in Figure 4.12 Note that the number of samples per row (Counts Per Row) and the number of grains over which those samples are taken (Pixels Per Row) are set by switches and potentiometers rather than via registers and are therefore only information for the calculation engines and not controllable parameters.

The 3-D transfer function panel shown in Figure 4.13 translates register settings into coefficient values and vice versa. Multiplier values can be adjusted via the sliders, or entered explicitly as a desired value. The panel object performs quantization into register settings in the indicated manner. Also, the configuration can be chosen, or

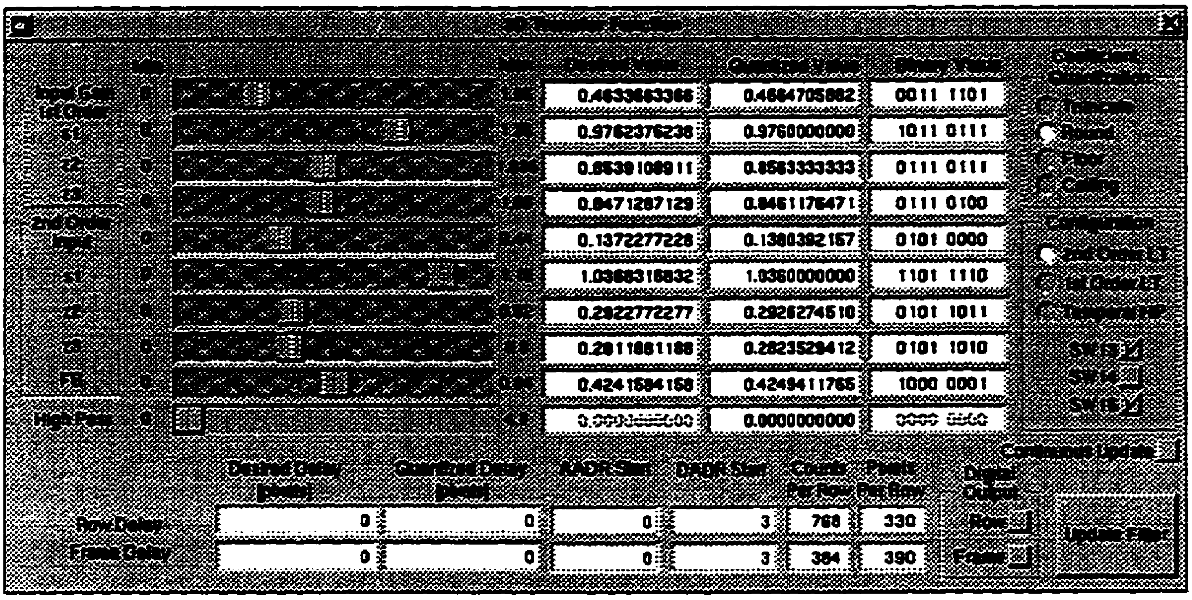

Figure 4.13: Coefficient, Structure and Delay Control Panel

the analogue switches controlled directly. Finally, the delay adjustments can be made, either in grains, or in samples, with quantization applied as for the coefficients. If the Continuous Update button is checked, requests for changes are made to the controlling object as the changes are made to the various fields, otherwise they are made in a batch when the Update Filter button is pressed.

The LT filter design tool shown in Figure 4.14 takes the abstraction one step further and represents the filter state in terms of the design parameters: filter trajectory, bandwidth, etc. When changed, these parameters are translated into filter coefficients via a simple optimization routine which selects delay lengths that minimize the difference in the quantized coefficients from a nominal "best" set. After quantization into register values, they are translated back into the design parameters as the quantized values. When the first order configuration is chosen, only the top five fields are active. When the second order configuration is chosen there are two modes. In tracking mode $\mathbf{L}_B$ is proportional to $\mathbf{L}_A$ and parameters are indicated in terms of inductance ratio, minimum and maximum bandwidth and R_{SL} as described in section 3.1.2. In independent mode the proportionality is not enforced and the second order section is parameterized as another first order section. The temporal

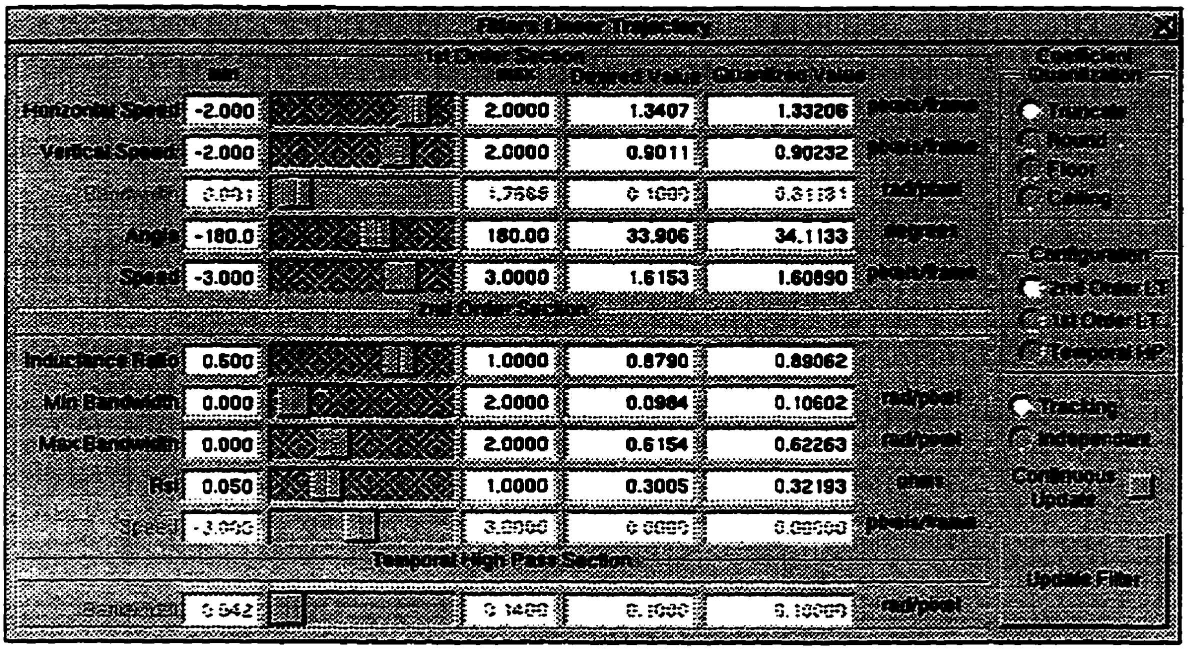

Figure 4.14: Linear Trajectory Filter Design Tool

HP configuration is discussed in section 4.7.

The fourth major control panel is the diagnostics panel shown in Figure 4.15. Sequences can be applied to each register or function of the modules and the function verified with logic probes. Note that since the diagnostic sequences generally use a batch process, the filter state image will not be accurate after diagnostics are applied.

4.7 Future Improvements

During the construction and testing of this prototype a number of issues have arisen which suggest improvements to the system.

In the current implementation a large majority of the components are required for the delay modules. Other delay implementations may be available in the near future to reduce this requirement. Bertschmann [3] showed that the available CCD based analogue delay elements introduce too much clock noise, or distortion, to be useful. Random access analogue memories may be available soon [57], which would eliminate

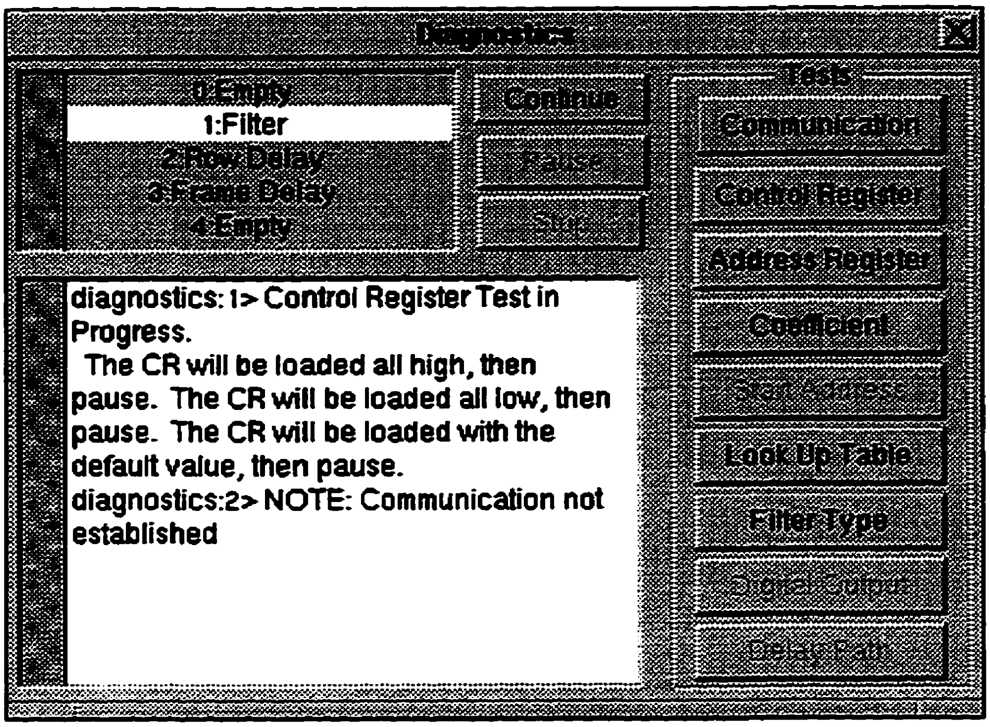

Figure 4.15: Diagnostics Control Panel

the need for an A/D and a D/A. They would also provide a continuous range delay, as opposed to the present digital system, but currently do not have enough cells or persistence for a frame delay. Many of them are also being designed as first in first out (FIFO) memories, which would reduce the addressing logic required. Digital FIFO memories of sufficient size and speed have recently become available which would eliminate the need for double buffering and much of the external control logic. The resulting design would be much smaller than the present implementation. A feasibility design indicates that about 34 ICs would be required for an implementation with only negative delay skews and that it would fit on one 8" x 10" PCB, compared to the present 178 ICs on 4 8" x 10" PCBs.

The original specifications for the system include the idea of an adapter module, to be developed in the future. The adapter module would control the system and implement a higher level algorithm such as tracking [10] or classification [14] using the filter as a building block. For this purpose jumpers exist so that the filter coefficient currents can be supplied through the backplane from the adapter module position. Also, the digital delays can output the delayed version of the signal as a digital stream to the backplane. The controller can yield control of the busses to the adapter and all the software has been developed with such expansion in mind.

Usually, most of the signal energy in a scene is associated with a static, or nearly static background, which contains large low spatial frequency components. The first order LT filter will pass some of this low frequency energy, which is usually not desired. A solution to this problem would be a temporal highpass filter in cascade, which would remove the static portion of all images. A temporal highpass filter was designed and implemented as part of Endeavour using the same frame delay element as the second order section of the bowl filter. Because the design required a large gain in one feedback loop, the current implementation is unstable and therefore no results are available.

A primary motivation in the development of mixD systems is a reduction, in comparison with digital systems, in size and power consumption. A very large scale integration (VLSI) implementation of this type of system would be the obvious next step in this direction. The analogue components and control would fit easily on one IC, with the FIFO memories external.

Chapter 5

Characterization of the Filter Response

Detailed measurement of the response of 2 and 3-D mixD filters in both the spatio-temporal and frequency domains has been a problem in the past, largely due to the difficulty in creating 2 and 3-D test images. Previously, the input signals were created from either printed 2-D sinusoids or computer generated movies of 3-D sinusoids with a video camera. The output of the camera was then applied to the filter and the input and output magnitudes measured with an oscilloscope [3,4]. In addition to the large amount of labour involved in each measurement, the method introduces inaccuracies from lighting, printing or display, the camera response and human measurement.

Najafi-Koopai [5] has recently derived a method to transform the 1-D frequency response of a raster scan based filter, which is relatively easy to measure from the 1-D input and output signals, into the 3-D frequency response. This serves to simultaneously reduce the labour involved in each measurement and to increase the accuracy. The application of this method to the Endeavour Analogue Video Filter for a number of filter settings forms the bulk of this chapter.

The steady state frequency response is only meaningful if the spatio-temporal transient response dies away quickly, and if all the components operate in the linear region. This is discussed qualitatively with examples in section 5.1. Non-linear effects due to signal values outside the linear region of some components can cause the filter to become unstable. This overflow-induced instability and the problem of calibration are discussed in sections 5.2 and 5.3.

The 1-D to 3-D frequency response transformation is described in section 5.4, followed by the measurement technique and test setup in section 5.5. Measurements of the frequency responses of several IDD and Bowl filters are given in section 5.6, along with a suggestion for a simple 1-D highpass post filter. The chapter concludes with a summary of the results.

5.1 Spatio-Temporal Response

The response of the filter to a variety of video sequences is acceptable for a wide range of settings in terms of clearly visible passband objects and visibly attenuated and smeared stopband objects. The difficulty in creating test sequences containing objects moving with exactly the same trajectory the filter is tuned to remains; but, with practice, you can move your hand in front of the camera in approximately the right direction and speed. Even with narrow bandwidth filters your hand is clearly visible in the output. A slight blurring of the lines and edges attests to the expected high frequency drop off due to the modified bilinear transform. If you move your hand in a different direction or especially at a higher speed the output image is smeared and attenuated.

In an attempt to make a more objective visual measurement, three cards with checkerboard patterns on them were attached to a wheel which was rotated at different speeds in front of a black backdrop. The scene was then evenly lighted and shot with a video camera. An input frame is shown in Figure 5.1a. If the wheel is rotated such that the speed of the cards is equal to the magnitude of the spatial velocity associated with the filter trajectory (see section 2.4.1), then, at one point in the rotation, one of the cards will be moving along the passband trajectory while the other two cards will be in the stopband. An output frame from the IDD filter with one card (at the top left) nearly in the passband is shown in Figure 5.1b. The contrast has been enhanced so that the printed output image looks similar to that seen on the monitor. Input and output frames for the bowl filter are shown in Figure 5.2, where the card at the top right is nearly in the passband. In both cases the checkerboard pattern is smeared somewhat due to the fact that the response of the filter is lower at high spatial frequencies than at low spatial frequencies (see figures 5.5 to 5.11) and because the card only has the correct trajectory momentarily as it rotates. It can be seen that the bowl filter eliminates more of the low frequency background energy than the IDD filter, especially in removing the static image of the tire.

5.2 Experimental Observations of the Transient Response and Overflow Effects

The filter is stable for a wide range of filter settings, with the transient effects confined to an area near the edge of the screen. However, it is necessary to vary the settings slowly and through stable ranges to maintain stability. If they are changed suddenly, or a poor setting is chosen, the transient effects may cover a large area of the screen for an extended period, or even lead to instability due to overflow. This is especially true for narrow bandwidth filters, which have long transients and small stability margins. Choosing effective filter settings, and varying them appropriately, is an important issue if closed loop adaptive algorithms are to be implemented.

Both filter structures are prone to instability due to overflow non-linearities. The overflow can be caused by spikes on the input, as mentioned in section 4.2, or by sudden large changes in coefficient values, especially when the filter is set to a narrow bandwidth. When an overflow occurs in a delay element or multiplier at some point in the image, the effect spreads quickly through the image and eventually all of the signals in the filter take on extremum values. A low pass filter has been incorporated into the video extraction circuit to reduce the incidence of spikes on the input, virtually eliminating that cause; but the coefficient change problem has not been addressed.

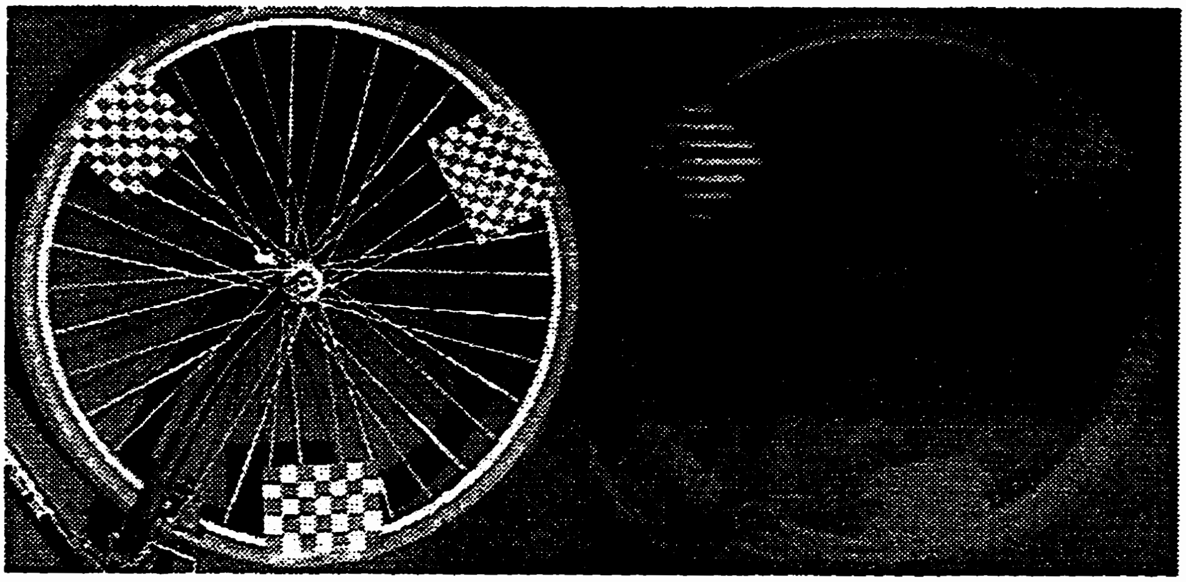

a) b)

Figure 5.1: IDD Filter Spatio-Temporal Response. a) Input Frame. b) Output Frame. The top left card is nearly in the passband.

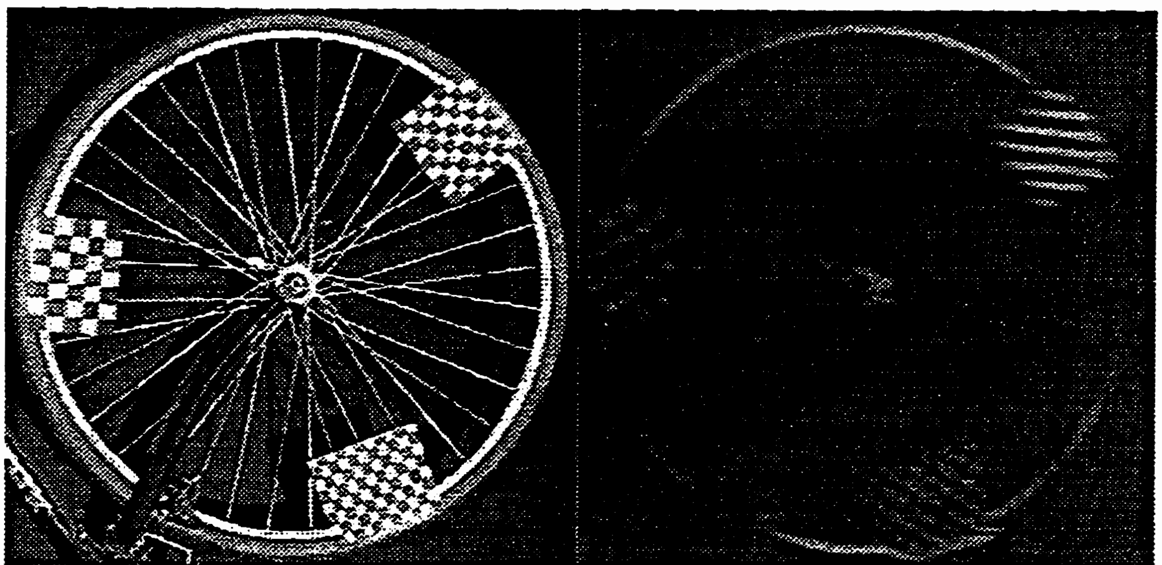

a) b)

Figure 5.2: Bowl Filter Spatio-Temporal Response. a) Input Frame. b) Output Frame. The top right card is nearly in the passband.

Possibly the GUI could enforce slow changes. It may even be possible to alter the design so that the extremum would not be an equilibrium point.

Once the filter enters the extremum state, it is not possible to reset it without changing the coefficient values with the current implementation. One solution would be to apply zero initial conditions by zeroing the input to each delay element and integrator for at least one frame. This might even be done automatically using the overflow detection signal from the A/D converters in the delay circuits. The filter could be made more robust by forcing the outputs of the delay elements to be zero during the entire synchronization interval [11].

5.3 Calibration

One of the more difficult aspects of tuning an analogue filter is calibration. Because filter parameters depend on component values, the exact relationship between the binary value stored in the D/A converter and the multiplier coefficient associated with it must be measured to be known. It must also be measured in the closed loop situation to be accurate, as the loading on the op-amps, and therefore the response, will change slightly if the loops are broken. A solution to the problem is to model the measured frequency response of the system with the appropriate transfer function (IDD or Bowl) and to compare the coefficients of the modelling transfer function with the filter settings. To get an accurate relationship between the filter settings and the coefficients of the transfer function, measurements must be taken for widely different values of each parameter. The D/A converters and multiplier gains should be nearly linear; so, once calibrated, the system should be fully predictable.

5.4 1-D to 3-D Frequency Response Transformation

To simplify the task of measuring the frequency response of raster scan based filters it is useful to investigate the effect of the raster scan transformation in the frequency domain. It may be shown that, along the line in 3-D frequency space given by [5]

$$\begin{aligned} \omega_1 &= \omega T_1 \\ \Omega_2 &= \omega T_2 \qquad \text{or} \qquad \frac{\omega_1}{T_1} = \frac{\Omega_2}{T_2} = \frac{\Omega_3}{T_3} \\ \Omega_3 &= \omega T_3 \end{aligned} \tag{5.1}$$

known as the slicing line, the raster scan of a 3-D sinusoid with frequency $(\omega_1, \Omega_2, \Omega_3)$ is a 1-D sinusoid with frequency ω:

$$\sin(\omega_1 h_1 + \Omega_2 n_2 + \Omega_3 n_3 + \phi) = \sin(\omega T_1 h_1 + \omega T_2 n_2 + \omega T_3 n_3 + \phi) = \sin(\omega t + \phi). \tag{5.2}$$

Thus the response of the filter to a 1-D sinusoidal input will be exactly the same as its response to the raster scan of the 3-D sinusoidal input with frequency given by equation 5.1. Using the raster scan transformation given in equation 2.38, it is also possible to show that the Fourier transform of the raster scan of a signal is equivalent to the Fourier transform of the 3-D signal along the slicing line [5]:

$$\begin{aligned} R(e^{j\omega}) &= \int_{-\infty}^{\infty} r(t) e^{-j\omega t}\, dt \\ &= \sum_{n_3=-\infty}^{\infty} \sum_{n_2=0}^{T_3/T_2} \int_0^{T_2/T_1} r(h_1 T_1 + n_2 T_2 + n_3 T_3) e^{-j\omega(h_1 T_1 + n_2 T_2 + n_3 T_3)}\, dh_1 \\ &= \sum_{n_3=-\infty}^{\infty} \sum_{n_2=0}^{T_3/T_2} \int_0^{T_2/T_1} x(h_1, n_2, n_3) e^{-j\omega h_1 T_1} e^{-j\omega n_2 T_2} e^{-j\omega n_3 T_3}\, dh_1 \\ &= X(e^{j\omega T_1}, e^{j\omega T_2}, e^{j\omega T_3}) \end{aligned} \tag{5.3}$$

Because the Fourier transform of a raster scanned signal is periodic in the discrete dimensions, the slicing line can be drawn modulus 2π in these dimensions as in Figure 5.3. This slicing line fills the 3-D frequency space such that the entire 3-D frequency response can be recovered from the 1-D frequency response as an interpolation [5]. However, for NTSC raster scanning this grid of lines fills the 3-D frequency space quite tightly, and for the purposes of characterizing the filter frequency response, measuring the response on a set of closely spaced points aligned with this grid is more than sufficient. The angles that the lines make with the Ω_3-axis are very small ($< 0.1°$), and may be ignored. The lines are spaced 0.0120 rad/grain ($2\pi T_2/T_3$) apart in the Ω_2 direction and 0.0294 rad/grain ($2\pi T_1/T_2$) apart in the ω_1 direction.

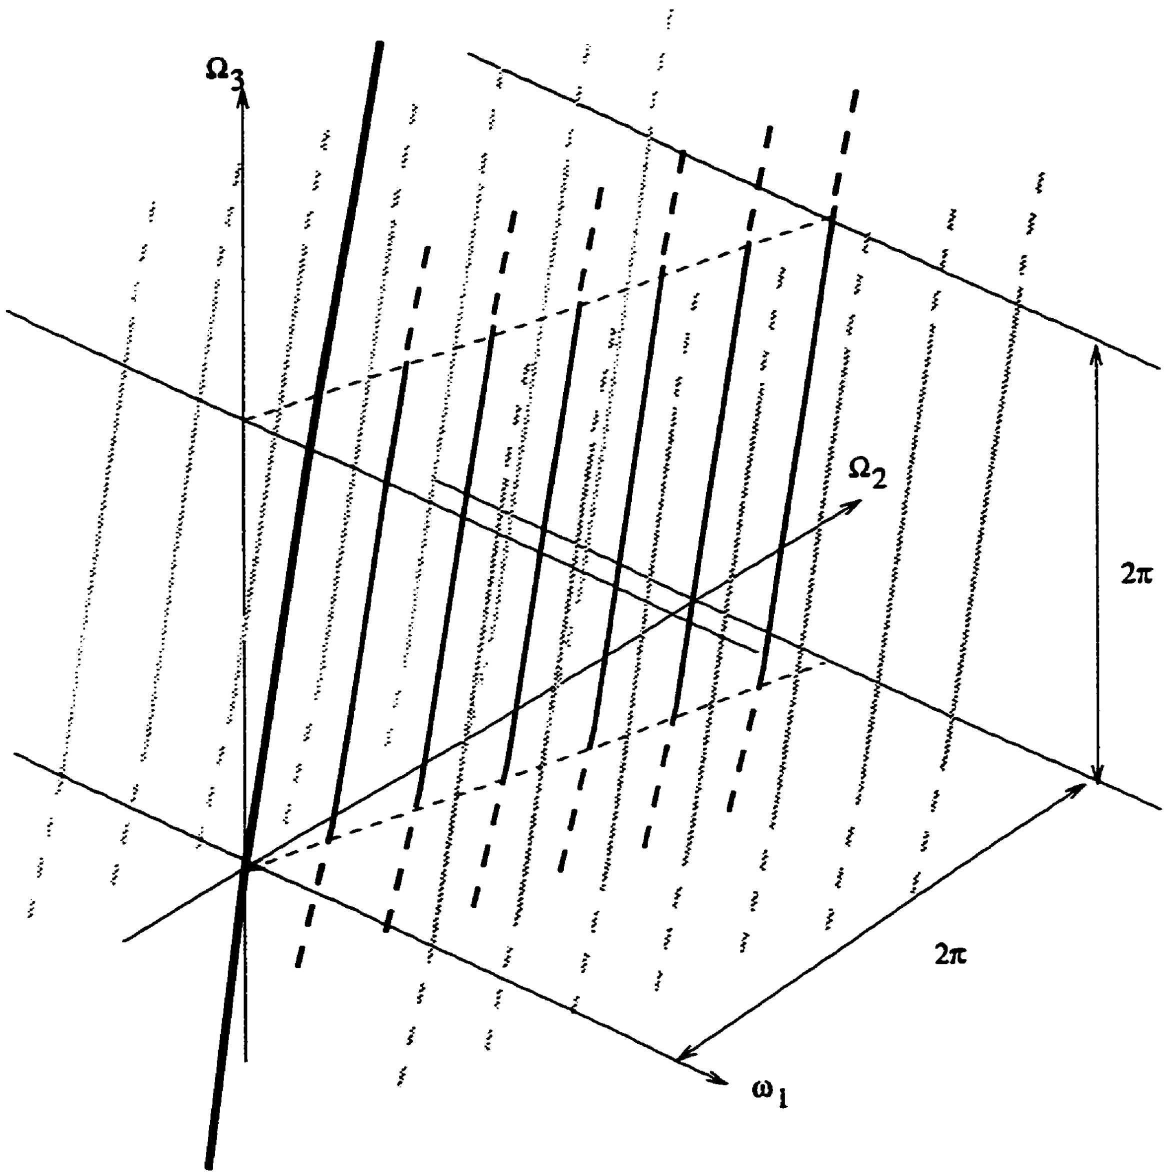

Figure 5.3: A Slicing Line Drawn modulus 2π in Ω_2 and Ω_3. Note: the slicing line for the NTSC raster scan fills the space much more densely than is shown.

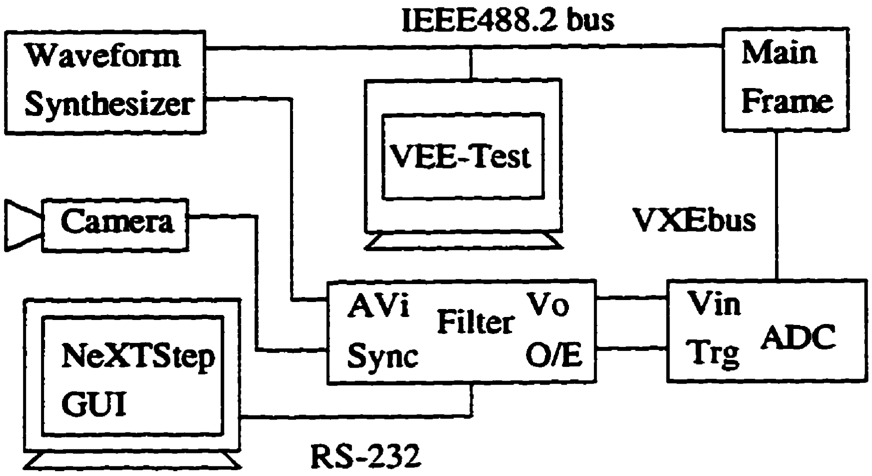

Figure 5.4: Test Setup

A simple approximation to the 3-D magnitude response at a set of closely spaced points may then be made by measuring the magnitude response to a 1-D sine wave input at a frequency, f, determined by

$$\begin{aligned} f &= \frac{k_1}{T_2} + \frac{k_2}{T_3} + \frac{f_3}{T_3} \\ f_1 &= \frac{k_1 T_1}{T_2},\ k_1 \in \mathcal{Z} \\ f_2 &= \frac{k_2 T_2}{T_3},\ k_2 \in \mathcal{Z} \end{aligned} \tag{5.4}$$

where f_1, f_2 and f_3 correspond to ω_1, Ω_2 and Ω_3 respectively, $-1/2 < f_2 \leq 1/2$ and $-1/2 < f_3 \leq 1/2$.

5.5 Measurement Technique and Test Setup

The test setup shown in Figure 5.4 uses the IEEE488.2 standard instrument communication protocol and HP VEE-Test software to control an HP 3325A waveform synthesizer and an HP E1406 Main Frame VXI controller with an HP E1430C 10 MHz 23 bit analogue to digital converter and storage (ADC). The filter is configured to use the video camera as an external synchronization source and to bypass the video extraction circuit, using the output of the waveform synthesizer directly as the input to the filter block. The video reconstruction module is also configured not to add the sync pulse to the output.

For each point $\{k_1, k_2, f_3\}$ in the requested set in three dimensional frequency space the, test routine sets the frequency of the signal source according to equation 5.4, pauses for 1 second to arrive at steady state in the temporal dimension, and then arms the ADC. The ADC uses the odd/even field sync signal to trigger and collects the active video portion of lines 65–68, which are far enough from the top for the system to be in steady state in the vertical dimension. Then the first $16.5\mu s$ of each line is discarded to arrive at steady state in the horizontal dimension. Finally a portion of each line equal in length to an integer number of cycles of the input is taken, the mean subtracted and the rms value calculated as the output magnitude. This value is stored in a file and the next iteration begins.

The row length T_2 used by the camera. and therefore by the delay elements, must be measured very accurately for the frequency response measurements to be meaningful. In equation 5.4 the term k_1/T_2 will quickly dominate the f_3/T_3 term with increasing f_1. Also, since f_3 is the only continuous frequency variable, any error in the calculation of f between the actual 3-D frequency created and the calculated frequency will appear in f_3. Thus a very small error in T_2 will indirectly cause a large error in f_3. For example, at $f_1 = 1/4 cycle/grain$, an error in T_2 of 0.002% will result in an error in f_3 of 1 cycle/grain. To address this, the period of the horizontal sync pulses was measured to within $\pm 10^{-7}$ with a frequency counter that was calibrated to the waveform synthesizer.

The frequency produced by the waveform synthesizer must be very precise also, as $\Delta f = 30$ Hz $\Leftrightarrow \Delta f_3 = 1$ cycle/grain. These sources of error primarily affect the orientation of the measured passband. rather than the bandwidth, and can be compensated, but this is unnecessary with the present equipment. The output frequency of the HP 3325A waveform synthesizer is accurate to $\pm 10^{-7}$ which, combined with the tolerance in T_2, results in an accuracy in Ω_3 of better than ± 0.01 rad/grain throughout the measurement range.

5.6 Filter Responses

5.6.1 IDD Filter

The IDD filter has four important controllable parameters:[1] the gain, G_1, associated with the lossy integrator; two gains, G_2 and G_3, associated with discrete differentiators; and the variation in the length of the frame delay, ΔT_3. The gains correspond to transmittances shown in Figure 3.16, and are related to the filter design parameters by the following equations.

$$\begin{aligned} G_1 &= \frac{L_1}{R_L + (1 + D)(L_2 + L_3)} \\ G_2 &= \frac{(1 + D)^2 L_2}{R_L + (1 + D)(L_2 + L_3)} \\ G_3 &= \frac{(1 + D)^2 L_3}{R_L + (1 + D)(L_2 + L_3)} \end{aligned} \tag{5.5}$$

The input gain is set to make the passband gain approximately one and the variation in the row delay length changes the response to spatial shapes rather than to motion and so is not used here.

Four IDD filters, labelled A to D, were characterized. The filter design parameters for each, which were found by modelling the measured response with a calculated frequency response as suggested in section 5.3, are given in table 5.1. The delay length was kept constant for all four tests. Filters B, C and D are identical to A except that the gain G_1, G_2 or G_3 respectively was reduced slightly. Equation 5.6 reveals that G_1 only affects L_1, but while G_2 primarily affects L_2 it has secondary effects on L_1 and L_3. Similarly, G_3 is primarily related to L_3 but is also related to L_1 and L_2. The horizontal and vertical speeds associated with the filter trajectory, as calculated from equations 3.4 and 3.11, are also given in table 5.1 along with the bandwidth from equation 3.6. The effect of the delay change on the bandwidth is not taken into account as it is only an approximate measure.

While the video extraction module, with its lowpass characteristic, is bypassed for the tests, the filter exhibits a 1-D dominant pole response in cascade with the fre-

[1]The delay element gain, D, is set by potentiometers and so is not a controllable parameter.

Table 5.1: Filter Parameters for Frequency Response Measurement of IDD Filter

Parameter	Unit	A	B	C	D
L_1	Ωgrain	12	9	8	8
L_2	Ωgrain	3	3	2.7	2.8
L_3	Ωgrain	3	3	2.9	2.5
R_L	Ω	1	1	1	1
ΔT_3	s	$-1.0T_1$	$-1.0T_1$	$-1.0T_1$	$-1.0T_1$
Horizontal Speed	grains/ frame	3.0	2.0	1.8	2.2
Vertical Speed	grains/ frame	-1.0	-1.0	-0.93	-1.12
Bandwidth	rad/grain	0.08	0.10	0.11	0.11

quency planar response. The pole is at approximately 530 kHz, or $\omega_1 \approx 0.46$ rad/grain and does affect the responses within the ROI. This pole is taken into account in the calculated filter responses and its effect is discussed further in section 5.6.3.

Contour plots of slices through the three dimensional magnitude frequency response of the filters at $\omega_1 = 0.032$ rad/grain, $\omega_2 = 0$ and $\Omega_3 = 0$ are shown in figures 5.5–5.8 part a, along with the calculated magnitude response in figures 5.5–5.8 part b. The responses are normalised to a maximum of 1 and contours are shown for 0.9, 0.7, 0.5 and 0.2 times the maximum. Note that the regions separating isolated contours of the same level are mostly sampling artifacts and the areas would be connected if measurements were made at more points. The samples are separated by 0.048 rad/grain in ω_1, 0.31 and 0.14 rad/grain in Ω_2 for the constant ω_1 case and the $\Omega_3 = 0$ case respectively and by 0.21 rad/grain in Ω_3.

The shapes of the filter responses are very similar to the calculated responses, confirming the operation of the filter. The measured response generally drops off more quickly with increasing ω_1 than the calculated response, which may indicate that the single dominant pole mentioned above is not sufficient. The planar shape of the passband is apparent in all the measurements.

The slight decrease in G_1 between filters A and B causes an anti-clockwise rotation in the $\Omega_2 = 0$ and $\Omega_3 = 0$ planes and an increase in bandwidth, but does not

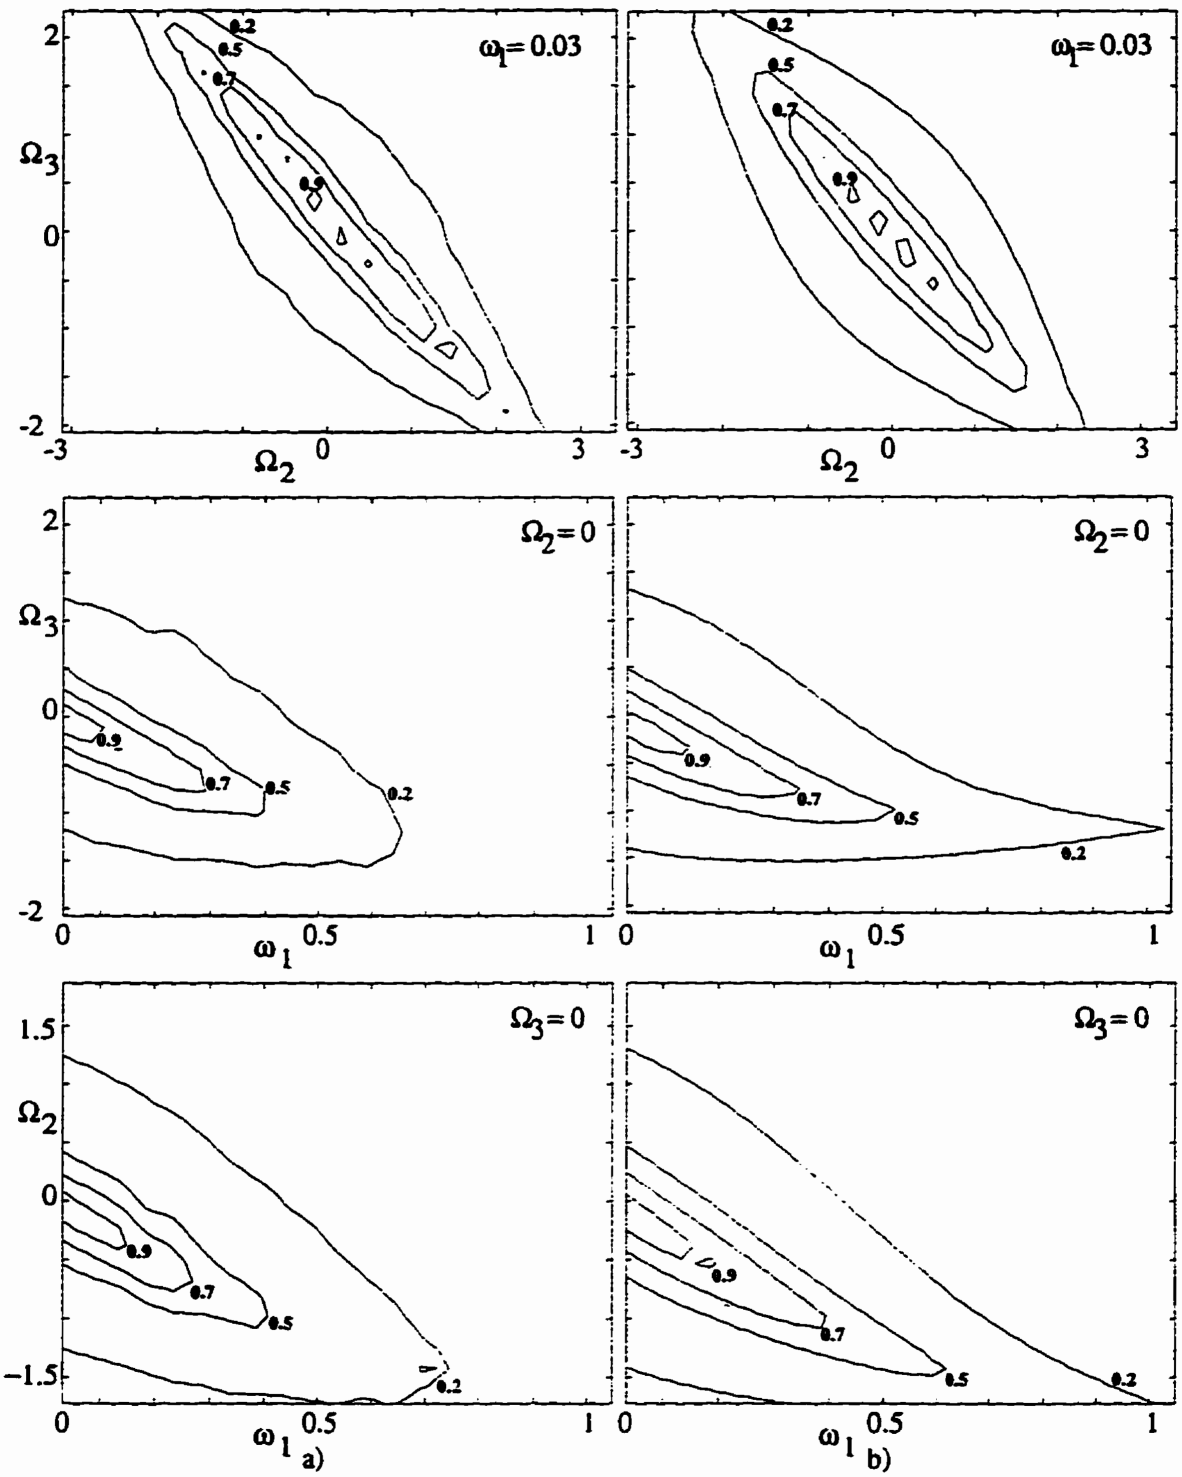

Figure 5.5: IDD Filter Response A. a) Measured b) Calculated. Frequencies in rad/grain.

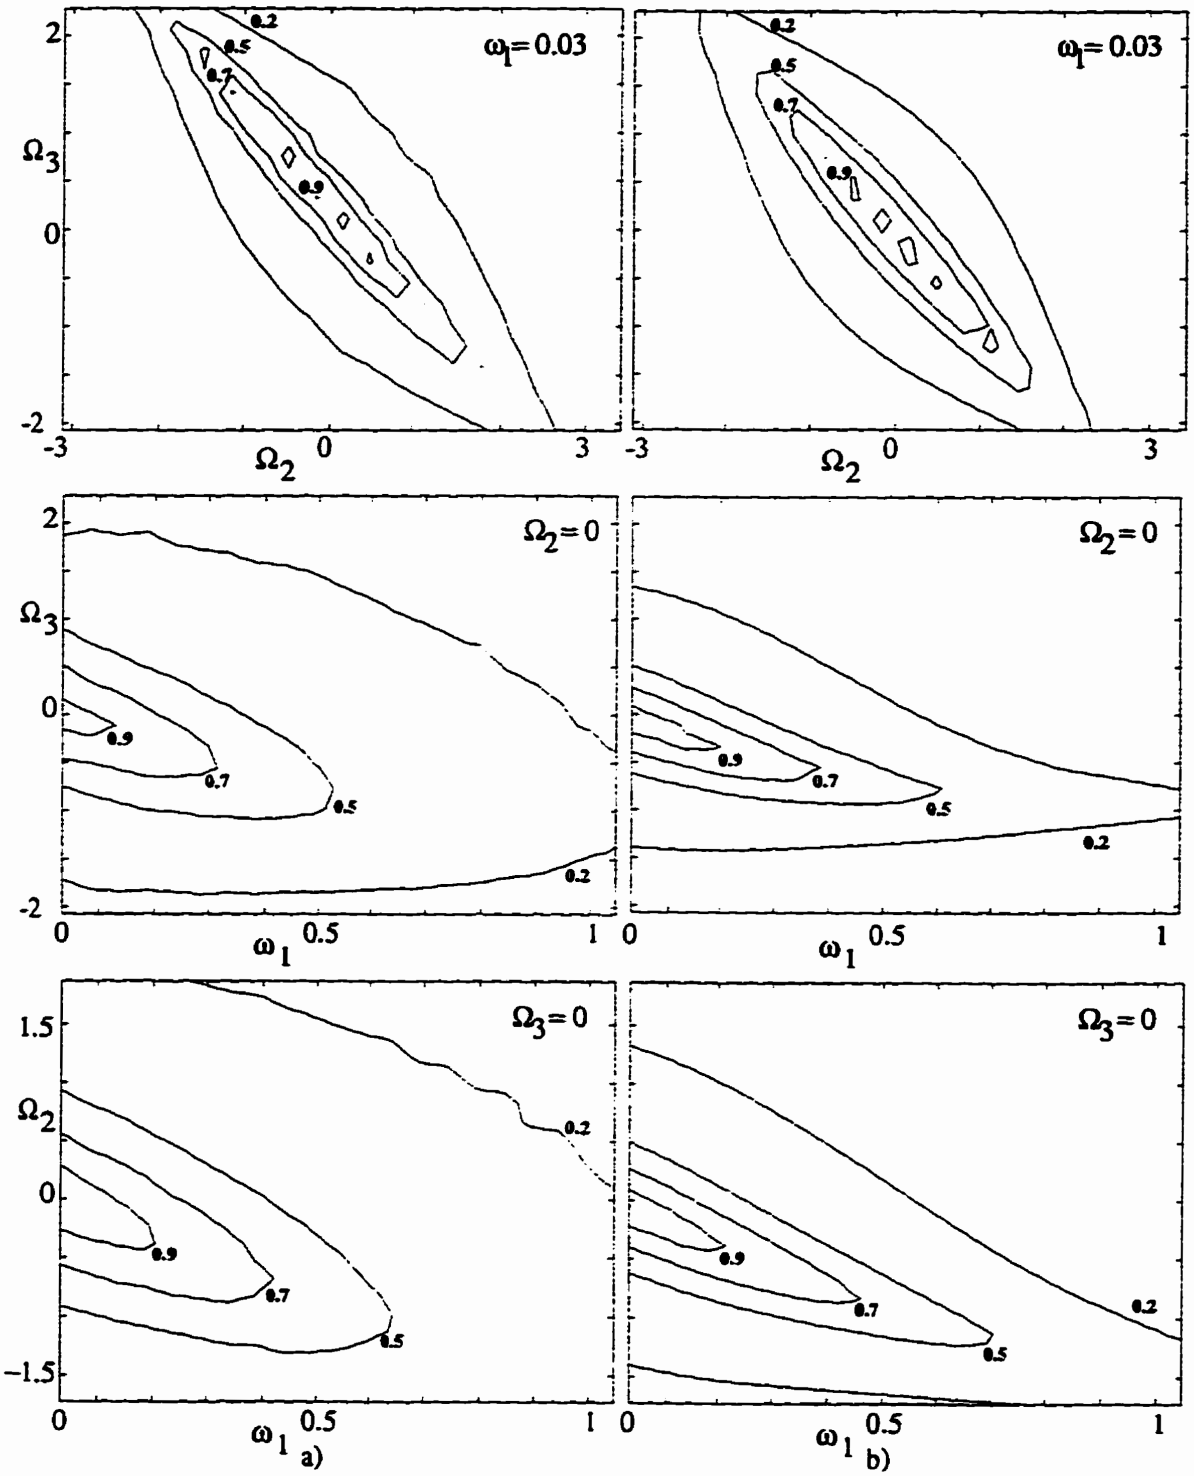

Figure 5.6: IDD Filter Response B. a) Measured b) Calculated. Frequencies in rad/grain.

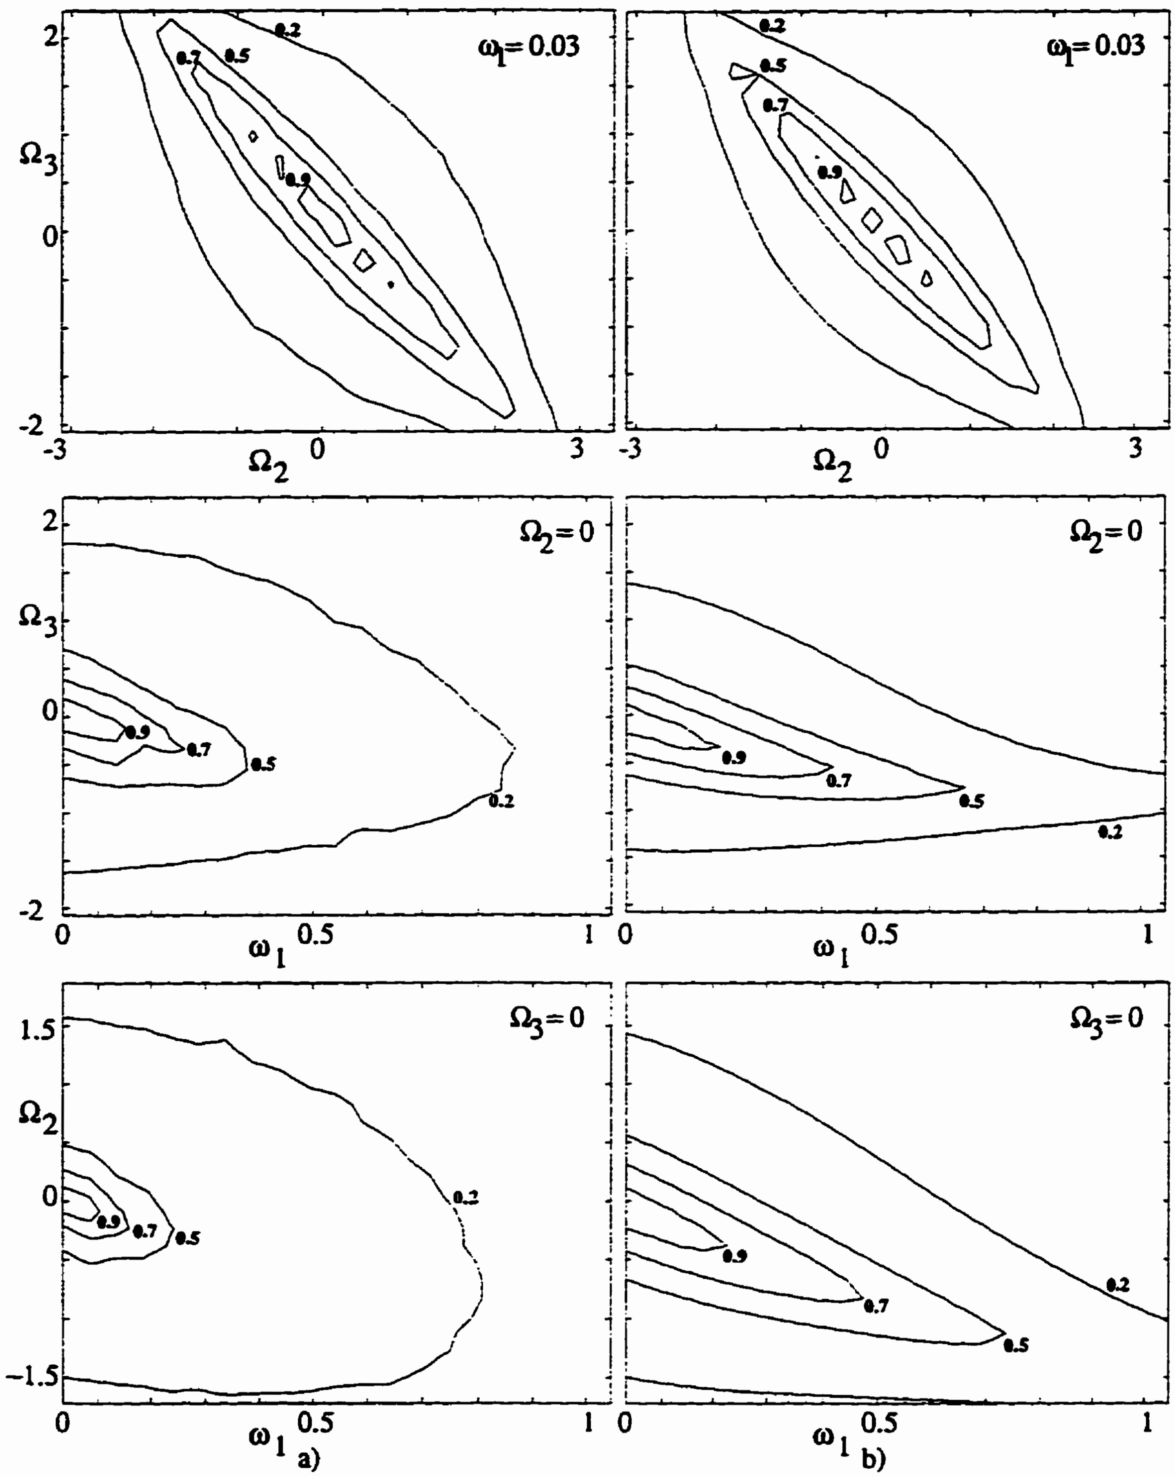

Figure 5.7: IDD Filter Response C. a) Measured b) Calculated. Frequencies in rad/grain.

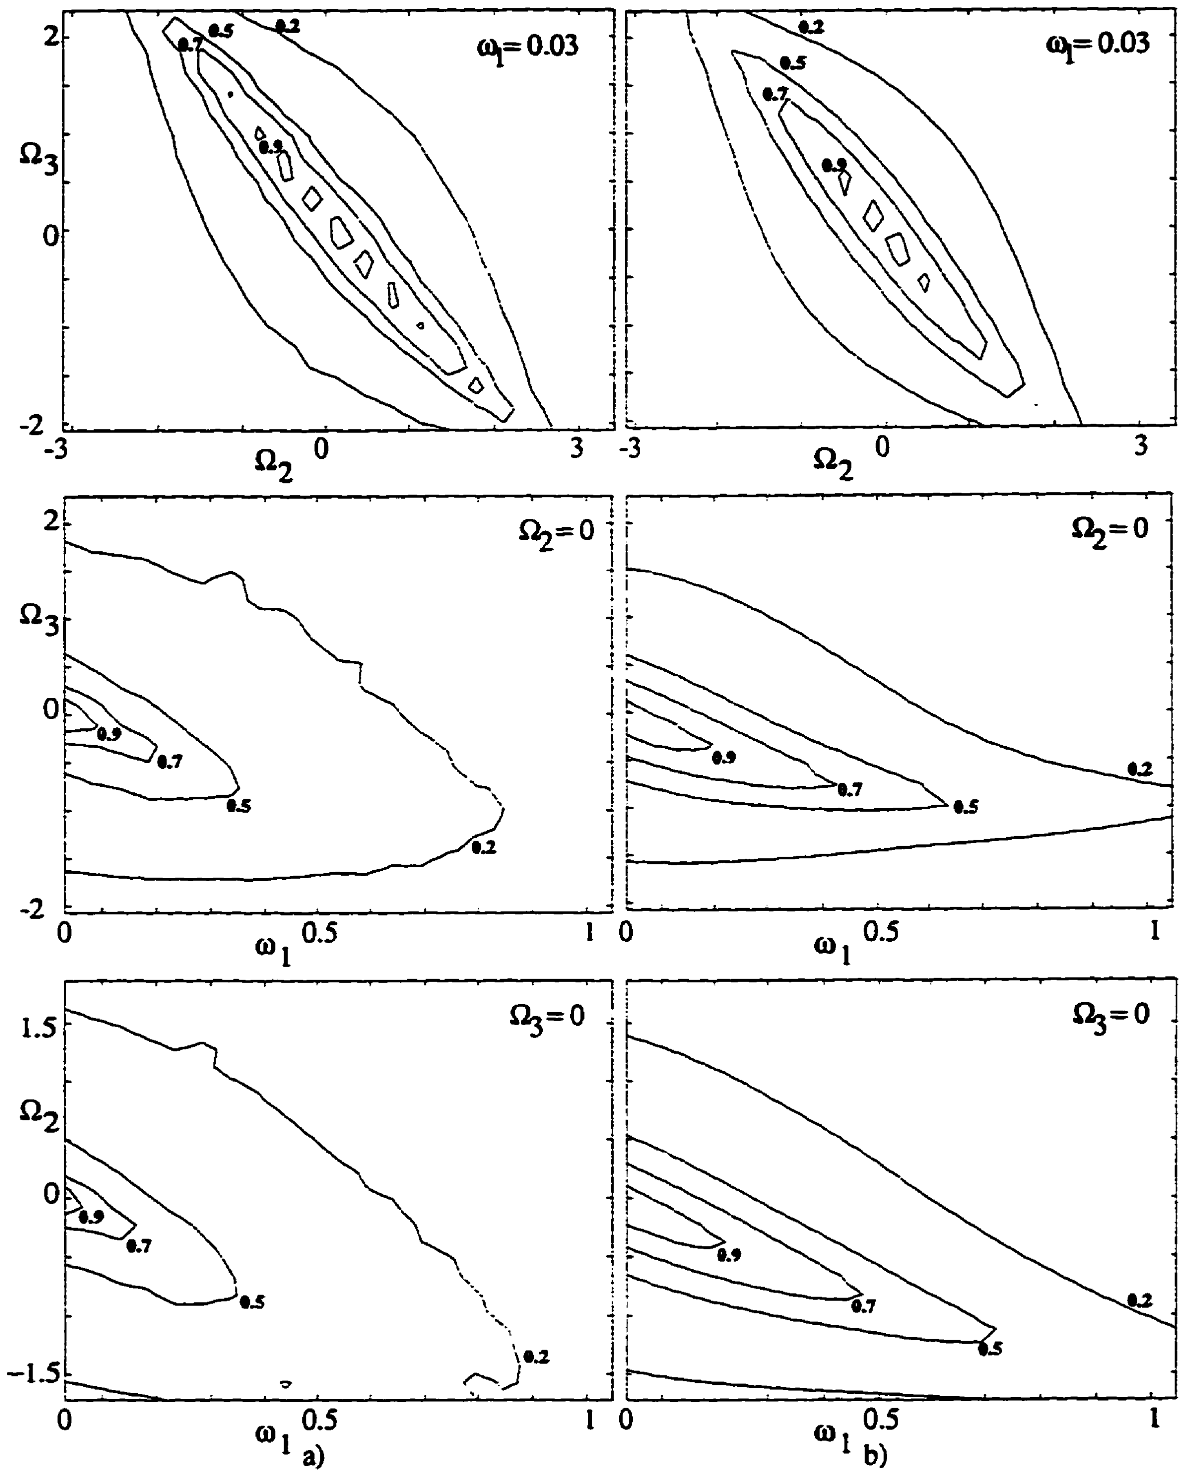

Figure 5.8: IDD Filter Response D. a) Measured b) Calculated. Frequencies in rad/grain.

Table 5.2: Filter Parameters for Frequency Response Measurement of Bowl Filter

Parameter	Unit	E	F	G
L_1	Ωgrain	1.1	1.1	1.1
L_2	Ωgrain	0.6	0.6	0.6
L_3	Ωgrain	1.2	1.2	1.2
R_L	Ω	1	1	1
K		0.879	0.879	0.879
R_{SL}	Ω	2	2	0.5
ΔT_3	s	0	$-3.05T_1$	$-1.0T_2$
Horizontal Speed	grains/ frame	0.92	-2.13	0.92
Vertical Speed	grains/ frame	-0.5	-0.5	+0.5
Maximum Bandwidth	rad/ grain	1.15	1.15	1.15
Minimum Bandwidth	rad/ grain	1.23	1.23	0.31

affect the constant ω_1 plane. This is consistent with a decrease in L_1. Changing G_2 between filters A and C causes an anti-clockwise rotation and an increase in bandwidth in all three, though the rotation in the constant ω_1 plane is barely noticable. The decrease in G_3 between filters A and D has a similar effect, except that the rotation in the constant ω_1 plane is clockwise. The fact that the passband plane rotates in a manner consistent with the change in settings of the filter indicates that it is working as expected in terms of steering.

5.6.2 Bowl Filter

The bowl filter has many more independent parameters than the IDD filter, so an exhaustive exploration of the effects of individual settings is too extensive for inclusion here. Instead, the measurements presented in figures 5.9 to 5.11 are of three filters labelled E to G that are identical except in the length of the frame delay. The design parameters for each filter, as determined by comparison with the calculated response, are given in table 5.2. While no changes were made to the coefficient settings between filters E and G, the apparent minimum bandwidth, determined by R_{SL}, did change. The reason for this is not known, but R_{SL} is proportional to $1 - G_{B2} - G_{B3}$, where G_{B2} and G_{B3} are gains associated with L_{B2} and L_{B3} respectively. As the gains are both approximately 1/2 in this filter, small changes will result in large changes in R_{SL}.

Contour plots of slices through the three dimensional magnitude frequency response of the filters at $\omega_1 = 0.032$ rad/grain, $\omega_2 = 0$ and $\Omega_3 = 0$ are shown in figures 5.9–5.11 part a, along with the calculated magnitude response in figures 5.9–5.11 part b, with the same resolution as in section 5.6.1. The dominant pole is taken into account in the calculations. The responses are normalised to a maximum of one and contours are shown for 0.9, 0.7, 0.5 and 0.2 times the maximum. The measured responses correspond very closely to the calculated responses, more closely, in fact, than the IDD responses. The one exception is in the $\Omega_3 = 0$ plane for filter G where the passband angles differ by about 30°. This may be related to the change in R_{SL}.

These responses show the frequency response skewing effect of the delay change. The reduction in the length of the frame delay by $3.05T_1$ skews the lines of constant ω_1 in the $\Omega_2 = 0$ plane upwards by $3.05\omega_1$ between figures 5.9 and 5.10, while not affecting the $\omega_1 = 0.03$ plane. Similarly, the reduction in the length of the frame delay by $1T_2$ skews the lines of constant ω_2 in the $\Omega_1 = 0.03$ plane by $1\Omega_2$, while not affecting the $\Omega_2 = 0$ plane.

While the measured response of the bowl filter matches the expected response very closely, the desired bowl shape is not achieved due to large minimum bandwidths related to large values of R_{SL}. The passband of filter G (Figure 5.11) is narrower near the origin than that of the others, but low values of R_{SL} corresponding to bowl-like shapes have not yet been possible. The problem may be that the value of R_{SL} is very sensitive to G_{B2} and G_{B3} and that the stability margin is proportional to R_{SL}. R_{SL} is also inversely proportional to the input gain of lossy integrator B, G_{B1}, and so can be lowered by increasing G_{B1}.

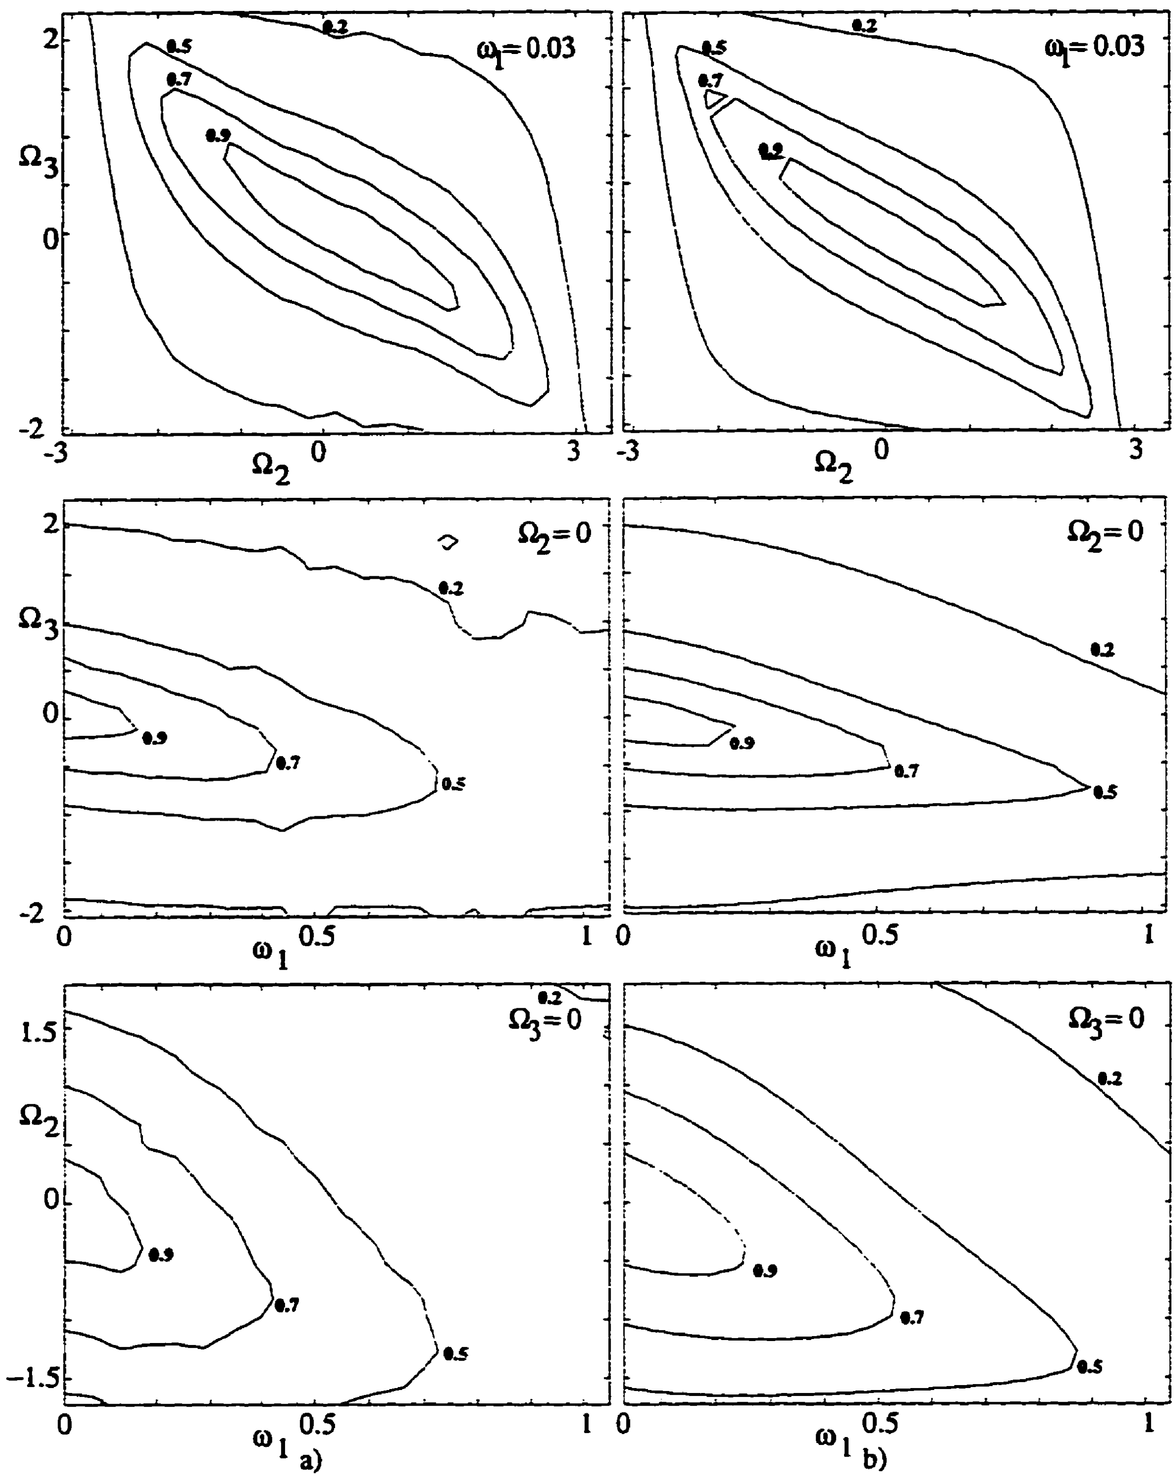

Figure 5.9: Bowl Filter Response E. a) Measured b) Calculated. Frequencies in rad/grain.

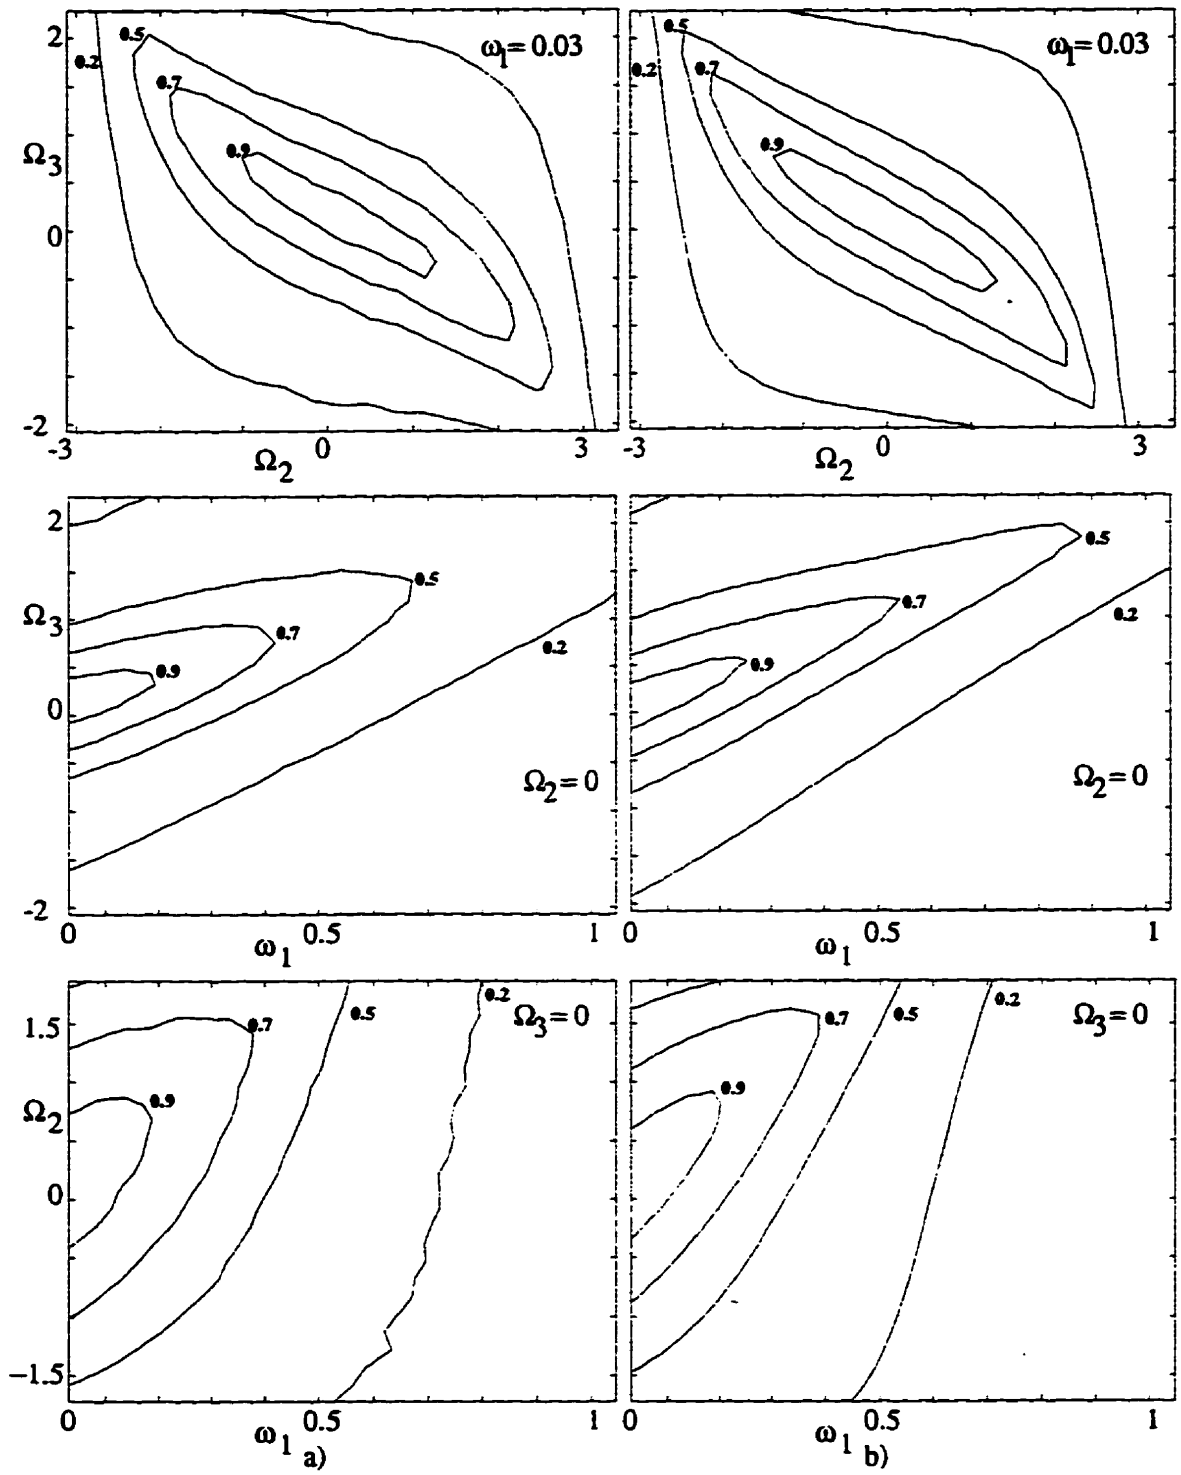

Figure 5.10: Bowl Filter Response F. a) Measured b) Calculated. Frequencies in rad/grain.

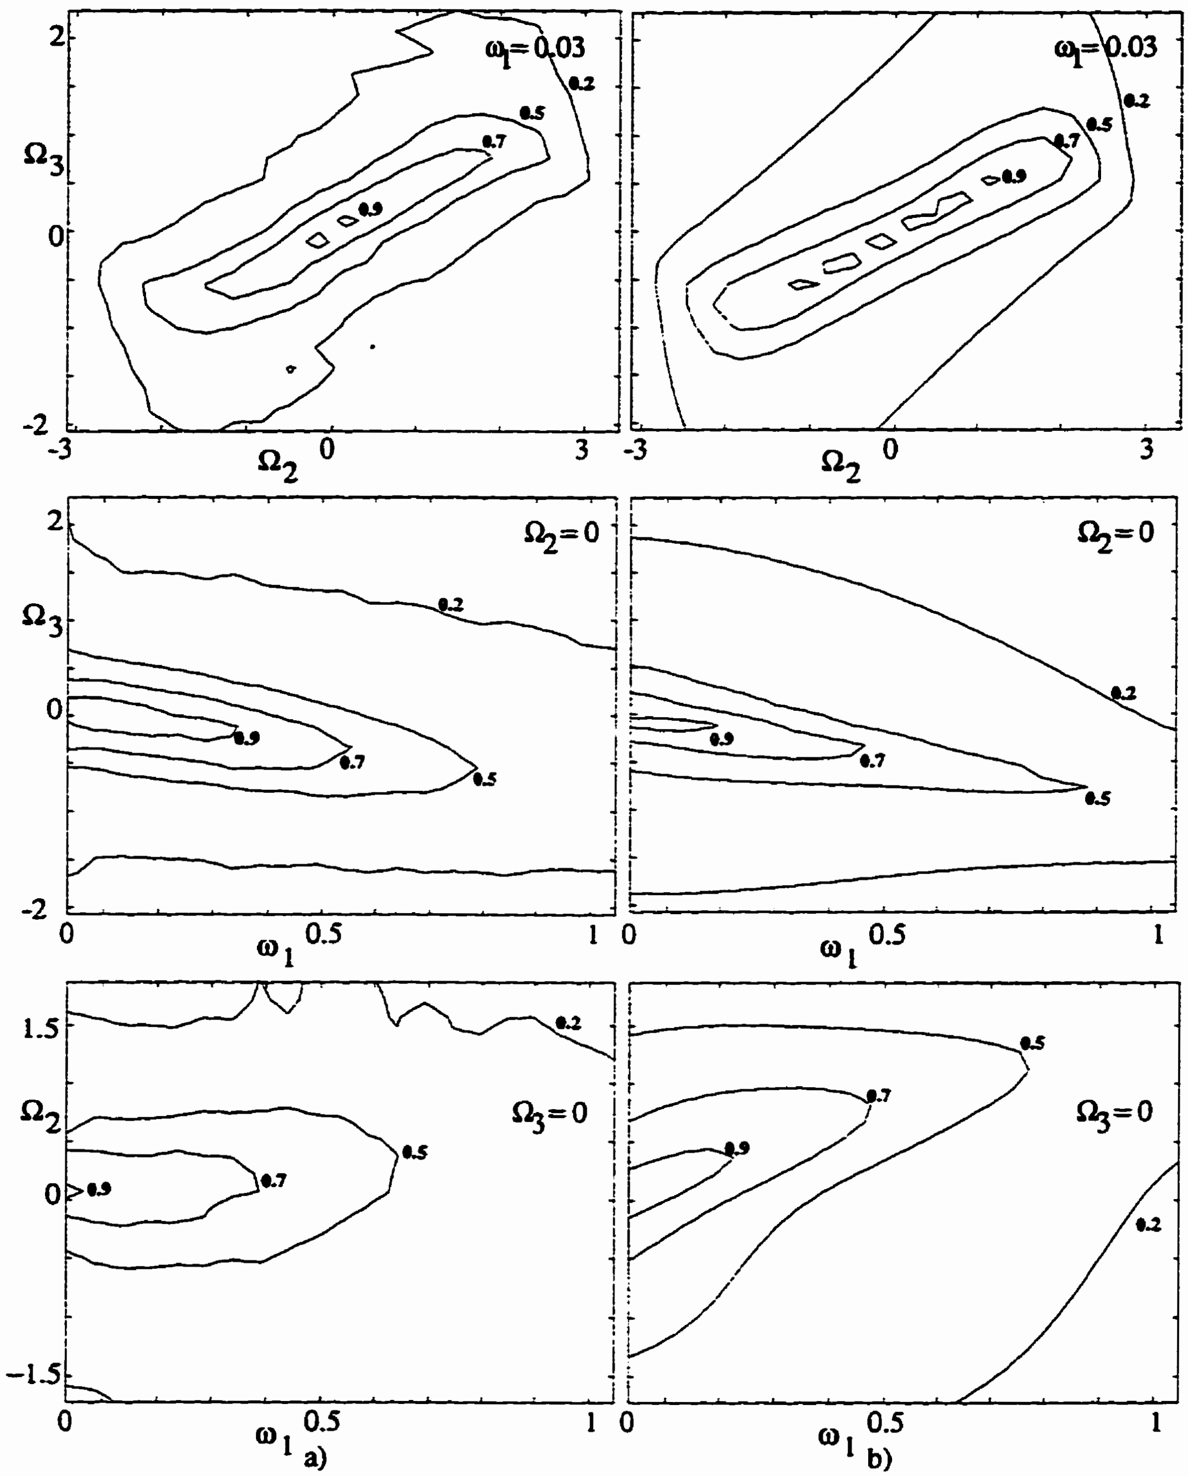

Figure 5.11: Bowl Filter Response G. a) Measured b) Calculated. Frequencies in rad/grain.

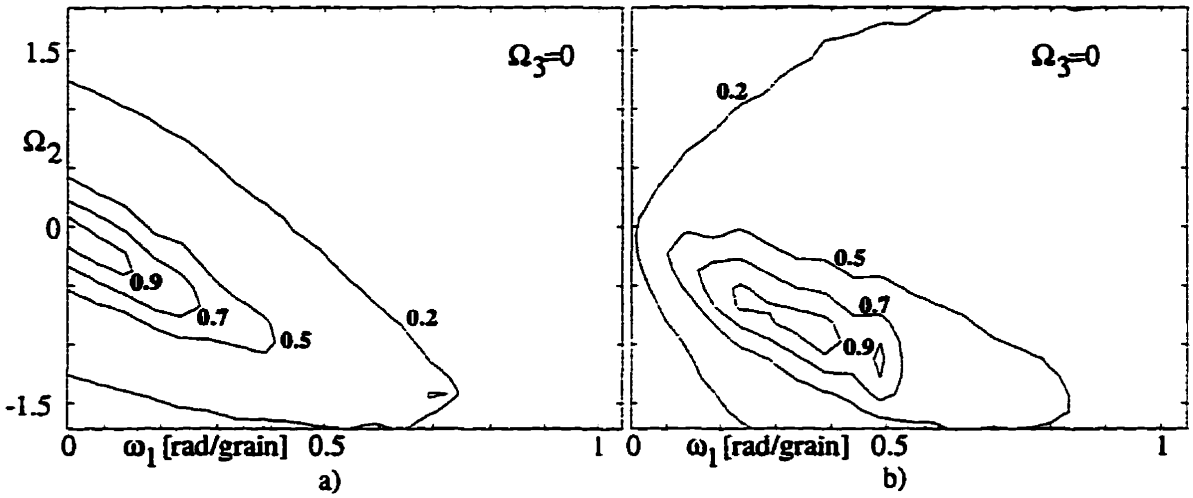

Figure 5.12: The Effect of the Highpass Postfilter on the Measured Magnitude Response of Filter A. a) The original measured response. b) The response after highpass filtering.

5.6.3 Highpass Postfilter

The reduction in the magnitude response of all the filters at high frequencies modelled by a dominant pole lowpass characteristic in cascade with the LT filter characteristics is a significant deviation from the ideal LT response. However, because it is apparently a purely 1-D characteristic, it can be compensated for easily with a 1-D highpass or bandpass filter. A 1-D highpass post filter has been suggested previously to remove the static background [58]. The use of a highpass post filter with a cut off frequency less than that of the dominant pole may actually create a semi-bowl shape from the IDD response. That is, the response near the origin and in the $\omega_1 = 0$ plane will be reduced, but the overall passband will simply be planar with an intersecting planar region removed.

The measured response of filter A in the $\Omega_3 = 0$ plane was multiplied by the first order 1-D highpass magnitude characteristic

$$M_{hp}(\omega) = \left|\frac{j\omega RC}{1 + j\omega RC}\right| \tag{5.6}$$

with $RC = 2 * 10^{-7}$s. The result is compared with the original in Figure 5.12. Note how the passband narrows near the origin giving it a fan shape. This response is better in terms of velocity selectivity than any of the bowl filter responses, but the phase distortion introduced by the highpass filter may be problematic, smearing the image visually.

The lowpass characteristic is possibly due to the delay through the feedback path in the variable lossy integrator. In future implementations it may be possible to avoid this by using a fixed bandwidth lossy integrator—a simple lowpass op-amp circuit as used in [3]. If it were set to a low cutoff frequency, corresponding to large values of L_1, variation of the effective H_s could be made with ΔT_3 and G_3. Either the bandwidth or the trajectory would then be variable only in discrete steps, but that may not be a serious problem. A much lower bandwidth, and therefore less costly and more stable, op-amp could be used;[2] and it would also simplify the control structure.

5.7 Summary

The results presented in this chapter indicate that the Endeavour Analogue Video Filter works quite well, both visually and in terms of the measured 3-D frequency responses. Problems with overflow instability, calibration and non-ideal 1-D lowpass characteristics are discussed and solutions are suggested. A technique for measuring the 3-D frequency response precisely and completely via the 1-D frequency response has been applied for the first time and the results are presented for 7 different filters. The measured responses differ from the ideal responses mostly by a 1-D lowpass characteristic, which can be compensated for by a 1-D highpass postfilter. It is also suggested that the circuit could be simplified by using a fixed bandwidth lossy integrator in place of the variable integrator in Figure 4.6. This may improve the response considerably by eliminating the delay in the feedback path which results in an extra lowpass characteristic.

[2]The EL2075 op-amp is not suitable for use with a capacitor in the feedback path.

Chapter 6

Sensitivity

The fact that implementations of mixD filters depend on non-ideal physical components such as resistors and capacitors means that the response of an implementation may differ from the ideal designed response. It is useful to quantify the effects of the variations in component values on the filter response. This enables the designer to determine what tolerances in component values must be met to guarantee that the filter response will be within some allowable deviation from the ideal. It also enables the designer to compare different implementations in terms of their probable performance before constructing them. One well known way to do this is through sensitivity measures.

The first order sensitivity of a general function F to one of its real parameters x is defined as [59]

$$S_x^F \equiv \frac{\partial \ln(F)}{\partial \ln(x)} \tag{6.1}$$

and relates the fractional change in the value of F at a point due to a small fractional change in x. Thus, if the value of the parameter x, which is related to some physical component values, is bounded to within a fractional portion its nominal value, F will be bounded to within some fractional portion of its nominal value given by [59]

$$\left|\frac{\Delta F}{F}\right| \approx \left|S_x^F \frac{\Delta x}{x}\right|. \tag{6.2}$$

While first order sensitivity can be applied to any one parameter, the filter implementations have many parameters that depend on component values. Two measures used to characterize a structure's sensitivity to multiple parameter variations and their interactions are worst case sensitivity $WS_{\mathbf{x}}^{F}$ and Schoeffler, or mean-squared, sensitivity $\Psi_{\mathbf{x}}^{F}$. Worst case sensitivity is given by [59]

$$WS_{\mathbf{x}}^{F} \equiv \sum_{i=1}^{N} \left| S_{x_i}^{F} \right| \tag{6.3}$$

where N is the total number of parameters, and indicates the maximum possible deviation in F which could occur if all parameters x_i took on their extremum values in the appropriate direction. Schoeffler sensitivity is defined as [59]

$$\Psi_{\mathbf{x}}^{F} \equiv \frac{1}{N} \sum_{i=1}^{N} \left| S_{x_i}^{F} \right|^2 . \tag{6.4}$$

and gives a measure of the distribution of the sensitivities to individual parameters. If the filter is much more sensitive to one parameter than to the others it will have a high Schoeffler sensitivity; and, statistically, the deviation from the ideal response of these filters will be larger than that of filters with a lower Schoeffler sensitivity.

6.1 Lower Bound Worst Case Sensitivity to Delay Element Errors

While deviation of the transfer function from an ideal value due to errors in the circuit elements of a filter is undesirable, the transfer function must depend in some way on some parameters if it is to be a useful shape. The design challenge is to choose the set of parameters and structure of the implementation such that the sensitivity is minimized. In this, it is useful to have a lower bound on the worst case and Schoeffler sensitivities of any structure implementing a given transfer function. These lower bounds can then act as benchmarks for the evaluation of the sensitivity properties of different filter structures. If the sensitivity of a structure approaches the lower bound then it is a good choice, otherwise a different structure may be needed. However, there is no guarantee that a structure exists that has sensitivities equal to the lower bounds.

Lower bounds on worst case and Schoeffler sensitivities to gain and phase errors of analogue delay elements in 1-D filters have been derived in [59]. These results will be extended here, with examples, to the mixD case.

Given that the transfer function can be represented as

$$H(\mathbf{s},\mathbf{z}) = M(\mathbf{s},\mathbf{z})e^{j\phi(\mathbf{s},\mathbf{z})} \tag{6.5}$$

where $M(\mathbf{s},\mathbf{z})$ is the real magnitude response and $\phi(\mathbf{s},\mathbf{z})$ is the real phase response, and defining group delay as

$$\tau_i(\mathbf{s},\mathbf{z}) \equiv \frac{\partial\phi(\mathbf{s},\mathbf{z})}{\partial\Omega_i} \tag{6.6}$$

then

$$S_{\Omega_i}^{H(\mathbf{s},\mathbf{z})} = S_{\Omega_i}^{M(\mathbf{s},\mathbf{z})} + j\Omega_i\tau_i(\mathbf{s},\mathbf{z}). \tag{6.7}$$

Also, it is useful to show that for any real parameter, x,

$$\begin{aligned} S_x^{H(\mathbf{s},\mathbf{z})} &= \frac{\partial \ln M(\mathbf{s},\mathbf{z})}{\partial \ln x} + j\frac{\partial\phi(\mathbf{s},\mathbf{z})}{\partial \ln x} \\ &= S_x^{M(\mathbf{s},\mathbf{z})} + j\phi(\mathbf{s},\mathbf{z})S_x^{\phi(\mathbf{s},\mathbf{z})} \end{aligned} \tag{6.8}$$

If all the delay elements in each dimension in the system have the same gain and phase errors, then, for the steady-state frequency response calculation, the delay element operation can be written as

$$\acute{z}_i^{-1} = B_i e^{-j\omega T_i\Theta_i} \tag{6.9}$$

where B_i represents a gain error and Θ_i represents a phase error. With these delay elements and ignoring any other errors in the system, its response is $H(\mathbf{s},\acute{\mathbf{z}}) = H(\mathbf{s},\mathbf{z})\mid_{z_i=\acute{z}_i}$. The sensitivity of this transfer function to gain errors is

$$\begin{aligned} S_{B_i}^{H(\mathbf{s},\acute{\mathbf{z}})} &= S_{\acute{z}_i}^{H(\mathbf{s},\acute{\mathbf{z}})} S_{B_i}^{\acute{z}_i} \\ &= -S_{\Omega_i}^{H(\mathbf{s},\acute{\mathbf{z}})} S_{z_i}^{\Omega_i} S_{\acute{z}_i}^{z_i} \\ &= \left[S_{\Omega_i}^{M(\mathbf{s},\acute{\mathbf{z}})} + j\Omega_i\tau_i(\mathbf{s},\acute{\mathbf{z}})\right]\frac{j}{\Omega_i T_i}. \end{aligned} \tag{6.10}$$

Similarly, the sensitivity of this transfer function to phase errors is

$$\begin{aligned}
S_{\Theta_i}^{H(s,\acute{z})} &= S_{\acute{z}_i}^{H(s,\acute{z})} S_{\Theta_i}^{\acute{z}_i} \\
&= S_{\Omega_i}^{H(s,\acute{z})} S_{z_i}^{\Omega_i} S_{\acute{z}_i}^{z_i} j\omega T_i \Theta_i \\
&= \left[S_{\Omega_i}^{M(s,\acute{z})} + j\Omega_i \tau_i(s,\acute{z}) \right] \frac{j\omega T_i \Theta_i}{j\Omega_i T_i} \\
&= \Theta_i \left[S_{\Omega_i}^{M(s,\acute{z})} + j\Omega_i \tau_i(s,\acute{z}) \right]
\end{aligned} \tag{6.11}$$

The term T_i associated with ω is in seconds, while the term T_i associated with Ω_i is in grains, so $\omega T_i = \Omega_i T_i$ in equation 6.11. From equations 6.10 and 6.11 and using equation 6.8 we can show that the sensitivity of the magnitude and phase of $H(s,\acute{z})$ to gain and phase errors are:

$$\begin{array}{ll}
S_{B_i}^{M(s,\acute{z})} = \dfrac{-\tau_i(s,z)}{T_i} & S_{B_i}^{\phi(s,\acute{z})} = \dfrac{S_{\Omega_i}^{M(s,z)}}{\phi(s,z)\Omega_i T_i} \\
S_{\Theta_i}^{M(s,\acute{z})} = \Theta_i S_{\Omega_i}^{M(s,z)} & S_{\Theta_i}^{\phi(s,\acute{z})} = \dfrac{\Theta_i \Omega_i \tau_i(s,z)}{\phi(s,z)}
\end{array} \tag{6.12}$$

These sensitivities are properties of the transfer function itself, rather than of any particular implementation or structure. In fact, any implementation of a transfer function will have exactly these sensitivities to *uniform* variations in delay element gain or phase [59]. However, the worst case sensitivity to non-uniform delay element errors cannot be less than this, as the uniform case is always a possibility. Thus, from equations 6.3 and 6.4 lower bounds on the worst case sensitivities are given by the magnitudes of equation 6.12, and lower bounds on the Schoeffler sensitivities by the squares of equation 6.12. For example

$$LWS_{B_i}^{M(s,\acute{z})} = \left| \frac{\tau_i(s,z)}{T_i} \right| \qquad \text{and} \qquad L\Psi_{B_i}^{M(s,\acute{z})} = \left| \frac{\tau_i(s,z)}{T_i} \right|^2 \tag{6.13}$$

where L indicates the lower bound.

6.2 Comparison of Direct Form and Ladder Form

The lower bounds derived in section 6.1 can be used to compare the direct form and ladder form structures implementing the first order planar pass filter. In this case

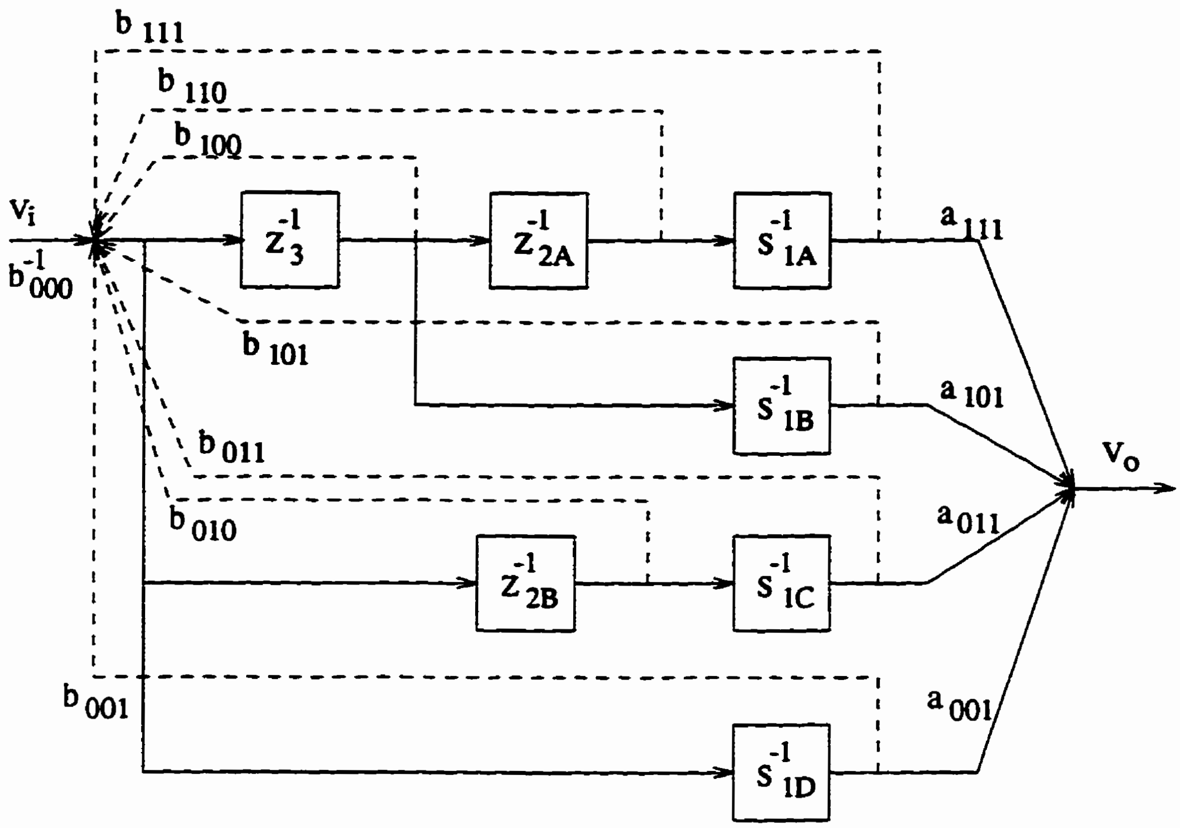

Figure 6.1: Direct Form Structure for the First Order Planar Pass Filter

the worst case sensitivity of the magnitude response to delay element gain errors are compared to the lower bound.

The direct form structure with the fewest possible delay and integrating elements is shown in Figure 6.1. It is assumed that frame delays are more costly than row delays and that both are more costly than integrators. Feedforward paths are shown as solid lines and feedback paths as dotted. This structure implements the transfer function

$$H_{df}(\mathbf{s},\mathbf{z}) = \frac{s_1^{-1} \sum_{i_2=0}^{1} \sum_{i_3=0}^{1} a_{i_3 i_2} z_3^{-i_3} z_2^{-i_2}}{\sum_{i_1=0}^{1} \sum_{i_2=0}^{1} \sum_{i_3=0}^{1} b_{i_3 i_2 i_1} z_3^{-i_3} z_2^{-i_2} s_1^{-i_1}}. \tag{6.14}$$

The coefficients for the planar pass filter are given in table 6.1 and come from equation 3.1 after application of the bilinear transform and collection of terms, with $R_L = 1$.

Table 6.1: Coefficients of the Direct Form Planar Pass Filter

		b_{000}	L_1
a_{00}	1	b_{001}	$1+2L_2+2L_3$
		b_{010}	L_1
a_{01}	1	b_{011}	$1-2L_2+2L_3$
		b_{100}	L_1
a_{10}	1	b_{101}	$1+2L_2-2L_3$
		b_{110}	L_1
a_{11}	1	b_{111}	$1-2L_2-2L_3$

This structure has two row delay elements labelled z_{2A}^{-1} and z_{2B}^{-1} in Figure 6.1, so the worst case sensitivity of the magnitude response of the direct form implementation to gain errors in the delay elements in the second dimension is

$$WS_{B_2}^{M_{df}(s,z)} = \left|S_{B_{2A}}^{M_{df}(s,z)}\right| + \left|S_{B_{2B}}^{M_{df}(s,z)}\right| \tag{6.15}$$

where B_{2A} and B_{2B} are the gains associated with z_{2A}^{-1} and z_{2B}^{-1} respectively.

Because the ladder form (IDD) filter has only one row delay element the variations are necessarily uniform, so the IDD filter has worst case sensitivities and Schoeffler sensitivities to delay variations that are equal to the lower bound.

A surface plot of the lower bound $LWS_{B_2}^{M_{PP}(s,z)}$ for $L_1 = 12, L_2 = 3, L_3 = 3$ and unity delay gains is shown in Figure 6.2a. These are the same parameters as for Filter A tested in chapter 5 except for the unity gains. A surface plot of $WS_{B_2}^{M_{df}(s,z)}$ for the direct form filter is shown in Figure 6.2b. The direct form structure is up to an order of magnitude more sensitive to these errors than the IDD filter.

Since both the direct form and the IDD filters have only one frame delay element they should have the same worst case magnitude sensitivity to delay gain errors equal to the lower bound. Deriving the sensitivity of each structure confirms this. A surface plot of the lower bound worst case sensitivity of the magnitude response to frame delay element gain errors, $LWS_{B_3}^{M_{PP}(s,z)}$, for the same filter as above is shown in Figure 6.3.

The decrease in the delay element gains that results from the modified bilinear

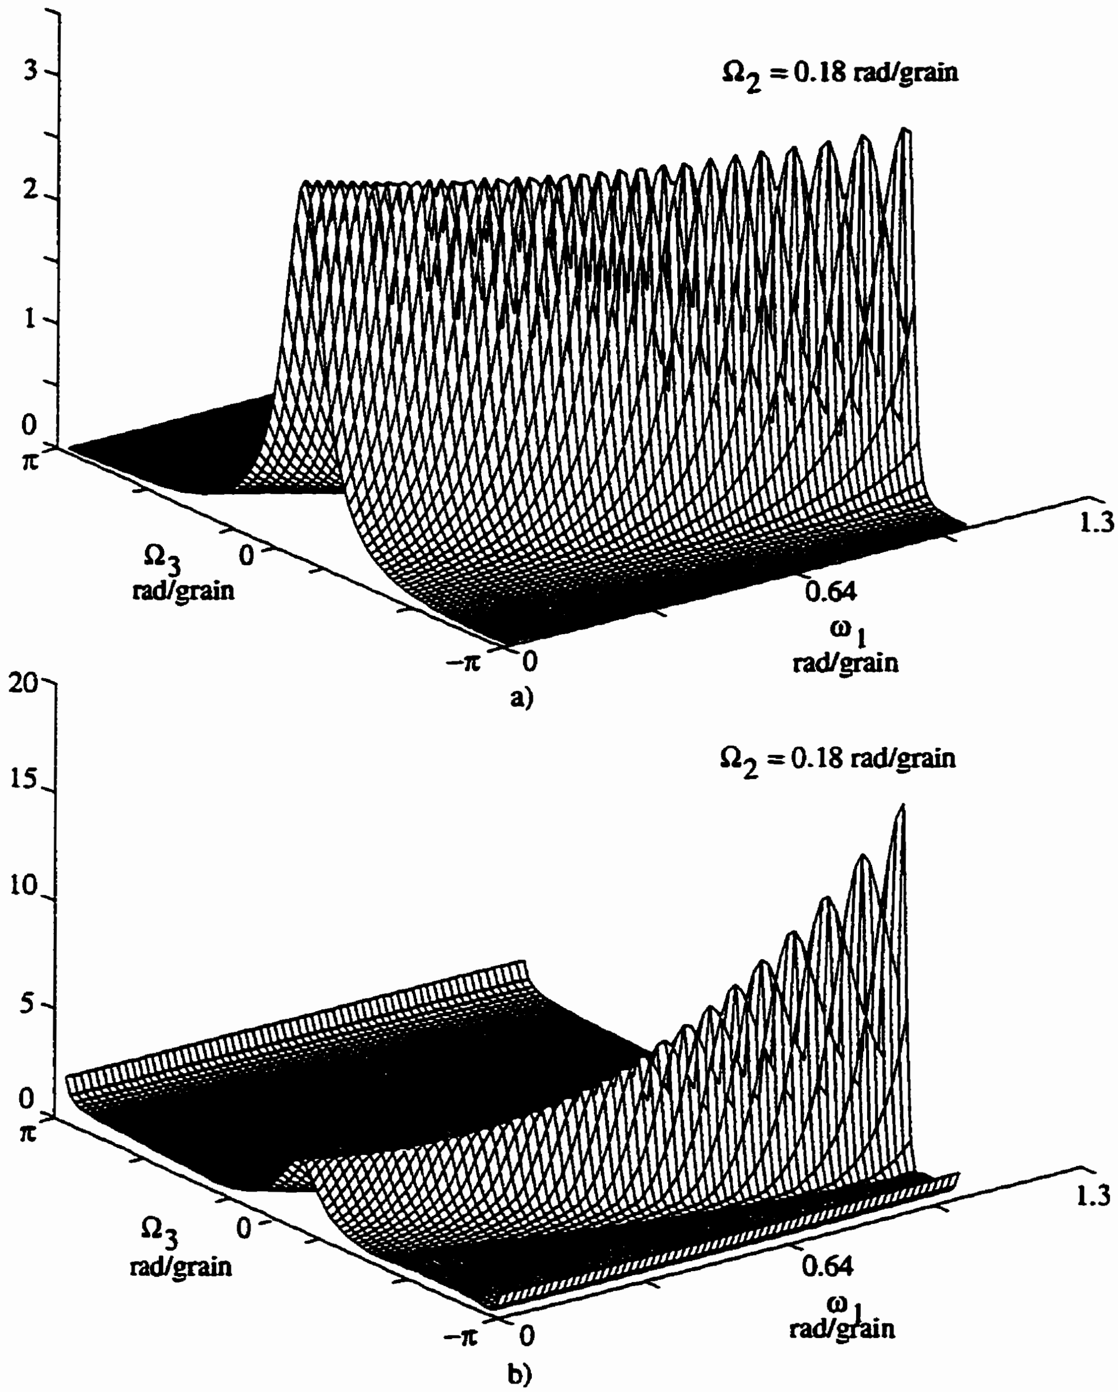

Figure 6.2: Worst Case Magnitude Sensitities to Row Delay Element Gain Errors. a) Lower Bound (IDD Filter's is equivalent) b) Direct Form Structure.

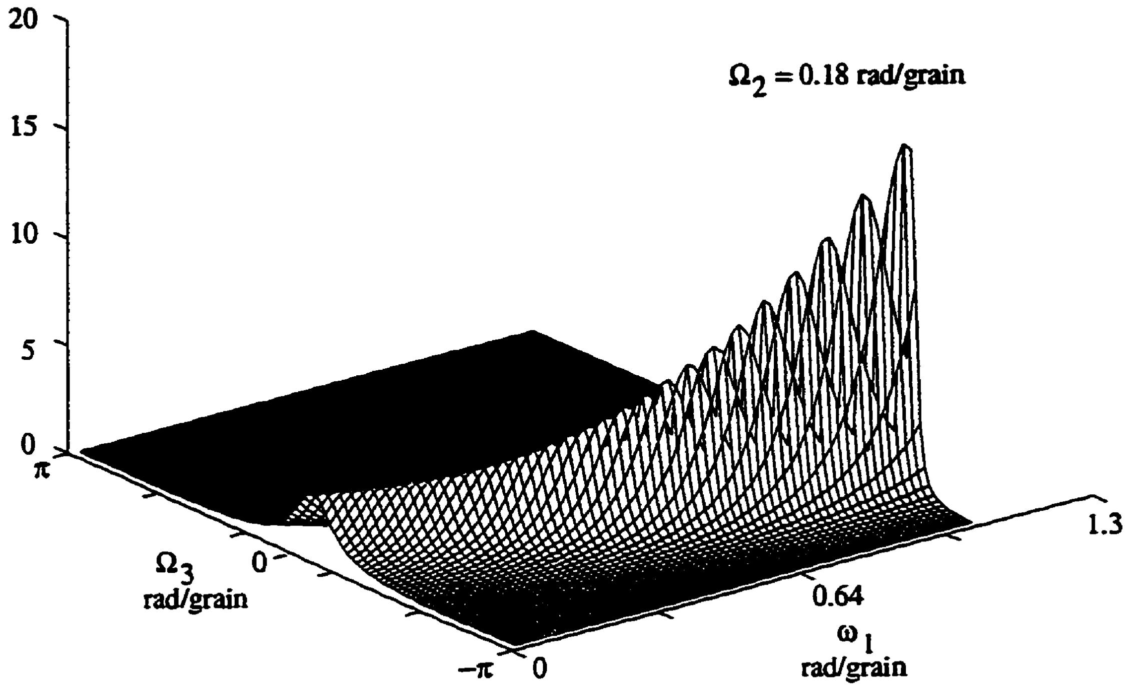

Figure 6.3: Lower Bound Worst Case Magnitude Sensitivity to Frame Delay Element Gain Errors

transform increases the stability margin of the filter. A useful side effect is that it also decreases the sensitivity to delay element gains. The lower bound worst case sensitivities of the magnitude response to errors in the row and frame delay element gains, $LWS_{B_2}^{M_{PP}(s,z)}$ and $LWS_{B_3}^{M_{PP}(s,z)}$, for the same inductance values as above, but with nominal gains of 0.9, are shown in Figure 6.4. The IDD filter worst case sensitivities are equal to the lower bound. Near $\Omega_3 = \pm\pi$ the sensitivity to frame delay gain errors actually increases significantly compared to the unity gain case, but since the magnitude response is very small there, it is not particularly a problem. The surface plots shown here are slices through the 3-D sensitivity functions, so many of the features of the functions are not shown. The slices were chosen to be as reprentative of the overall functions as possible.

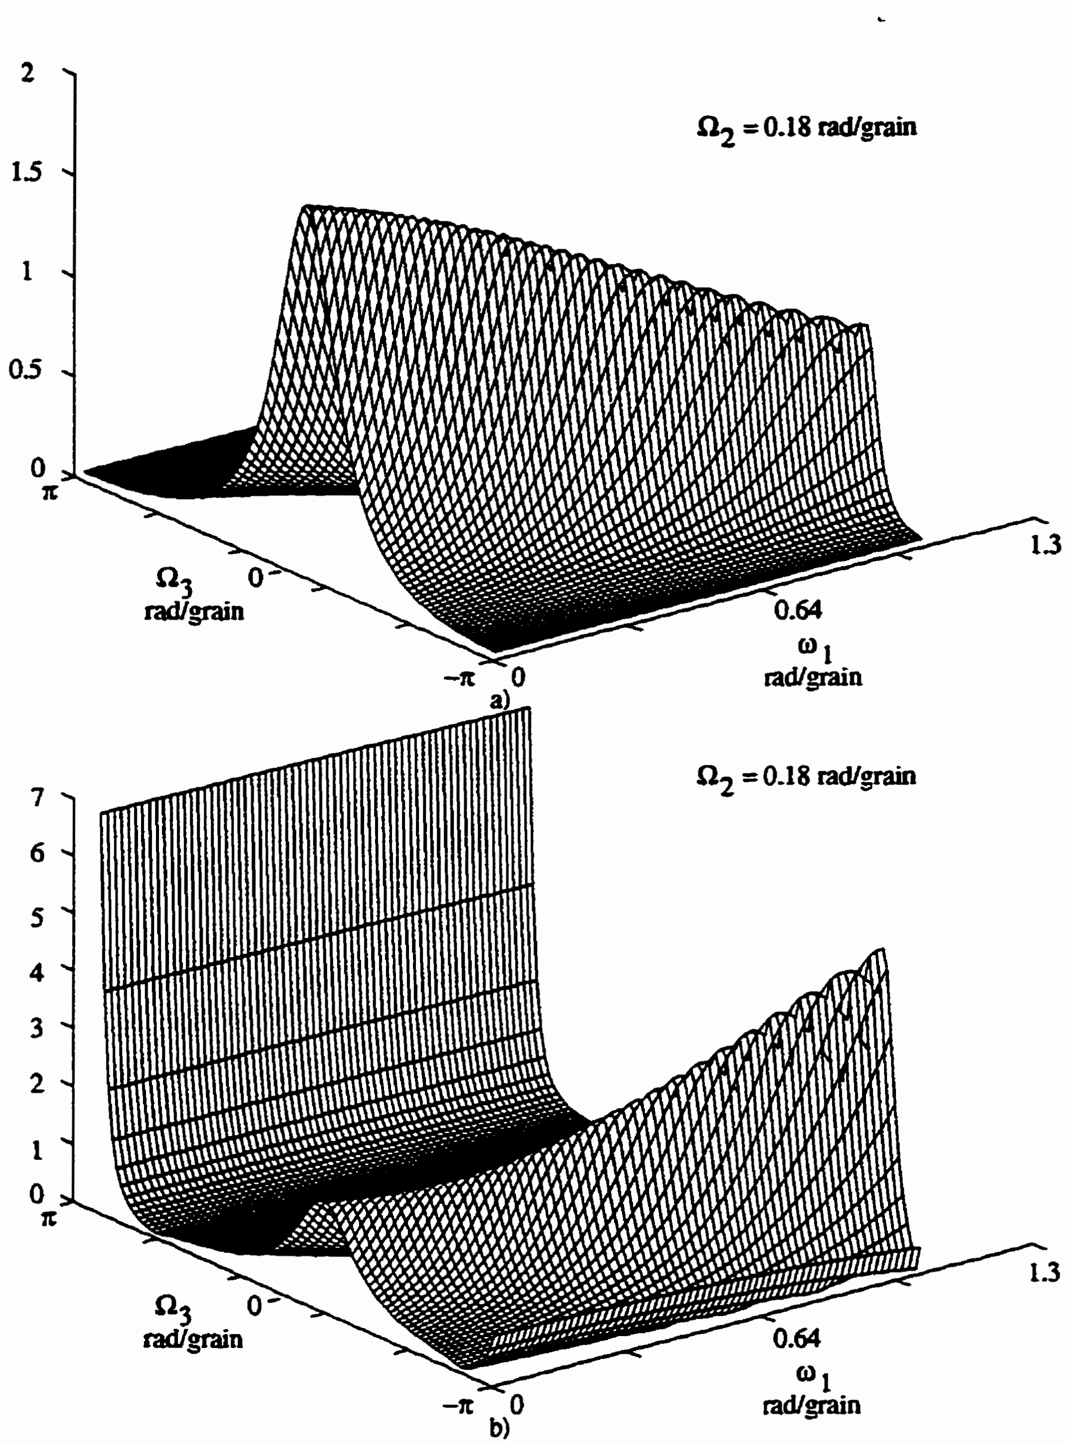

Figure 6.4: Lower Bound Worst Case Magnitude Sensitivity to Delay Element Gain Errors with Reduced Nominal Gains. a) Row delay gain errors. b) Frame delay gain errors.

6.3 Summary

Sensitivity measures are a well known way to quantify the effects of errors in component values on the response of an analogue filter. Specifically, the worst case and Schoeffler sensitivities provide useful measures of the possible and probable deviation from the ideal. In this chapter, lower bounds are found on the worst case and Schoeffler sensitivities to errors in the delay elements, for any mixD LSI transfer function. These lower bounds are useful for the selection of an appropriate filter structure.

As an example, the worst case sensitivity of the magnitude response of the IDD filter designed in chapter 3 to gain errors in the delay elements is compared to both the equivalent direct form filter and the lower bound. The IDD filter, which has a ladder form structure, has sensitivity equal to the lower bound, which is superior to the direct form version by up to an order of magnitude in the passband. This confirms that the ladder form technique is appropriate for the first order planar pass filter.

Further investigation reveals that the modification of the bilinear transform introduced in chapter 3 to increase the stability margin also reduces the sensitivity to delay gain errors in the passband.

It is also useful to point out that the worst case sensitivity of frequency planar digital ladder form filters to errors in the coefficients is lower than that of the corresponding direct form filters and wave digital filters [11]. It seems likely that this would hold true for the equivalent mixD filter structures.

Chapter 7

Practical BIBO Stability of M-D Systems

A filter must be stable for it to function effectively. Most definitions of stability imply that any transient effects eventually die away and that a desired steady state response will dominate the output. The most commonly used stability criterion in 1-D is bounded-input bounded-output (BIBO) stability, which states that for any bounded value input signal, the output will also be of bounded value. The BIBO stability condition has been extended to the M-D case and is extensively used in filter design; however, the direct extension to M-D is both difficult to test and unnecessarily restrictive.

In [31] Agathoklis and Bruton describe a practical-BIBO (PBIBO) stability condition for M-D discrete systems that is simpler to apply and less restrictive than the BIBO stability condition; yet is a sufficient condition for the effective function of the vast majority of practical systems. Both BIBO and PBIBO stability imply that for any bounded value input signal, the output will also be of bounded value; however, while BIBO stability bounds the value of the output over its entire domain, PBIBO stability bounds only those values of the output which are computable in finite time with a finite sized system [33]. Ignoring values of the output which can not be computed has no detrimental effect on the function of a system, but in many cases

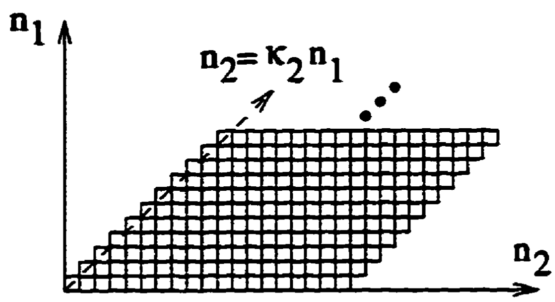

Figure 7.1: 2-D Example of a Computable Non-Rectangular Region of Support

simplifies and even improves the design.

The authors of [31] present conditions on both the impulse response and the z-transform transfer function of a linear shift-invariant discrete system for PBIBO stability that assume a rectangular region of support (ROS), or region of calculation. That is, they assume that all indices of the signals except one are bounded by a finite number and that the calculation proceeds along one of the dimensional axes; which is the case for the vast majority of systems. However, for the filters described in chapters 3 and 4, as well as some discrete systems, that is not always the case. Refering to Figure 3.18c it is apparent that the ROS for these filters is not rectangular when delay variations are applied: and, since these variations are to be used to steer the filter trajectory, it is important to determine if the filters will remain stable under these conditions.

The conditions for PBIBO stability given in [31] do not apply to systems with a non-rectangular ROS, such as that shown in Figure 7.1. For example, if the discrete difference equation

$$y(n_1, n_2) = x(n_1, n_2) + y(n_1 - 1, n_2 - 1) \tag{7.1}$$

which has the z-transform transfer function

$$H(z_1^{-1}, z_2^{-1}) = \frac{A(z_1^{-1}, z_2^{-1})}{B(z_1^{-1}, z_2^{-1})} = \frac{1}{1 - z_1^{-1} z_2^{-1}} \tag{7.2}$$

is calculated along the ROS given by $n_2 = n_1, n_1 \geq 0$ (or any ROS including this line) with the bounded input $x(n_1, n_2) = 1$; the output is a ramp function which increases without bound [26]. Thus, this system is unstable for the given ROS. However, both conditions for PBIBO stability given in [31] hold; and therefore different conditions are required for systems with a non-rectangular ROS.

In [33] the authors show that a computable ROS may only be of unbounded extent in one direction and this chapter extends the conditions for PBIBO stability to systems in which this direction is not along one of the dimensional axis. Thus, while the ROS is of bounded extent in all but one direction, the indices of the signal are not necessarily bounded. Also, the conditions are expanded to apply to mixD systems.

7.1 Non-Rectangular Regions of Support for Mixed Domain Systems

The conditions for PBIBO stability depend on the ROS of the system. Of the many computable non-rectangular regions of support, two of interest are considered in the remainder of this chapter. One can be described in terms of an unbounded continuous dimension and the other in terms of an unbounded discrete dimension.

To describe a non-rectangular ROS in terms of an unbounded continuous dimension, it is simplest to label that dimension as t_1 and the others in order such that the region can be written as

$$V_{ROS(t)} = \left\{ \mathbf{t},\mathbf{n} \left| \begin{array}{ccccc} 0 & \leq & t_1 & \leq & \infty, (T_1 = \infty) \\ \kappa_2 t_1 & \leq & t_2 & \leq & \kappa_2 t_1 + T_2 \\ & & \vdots & & \\ \kappa_{p-1} t_1 & \leq & t_{p-1} & \leq & \kappa_{p-1} t_1 + T_{p-1} \\ \kappa_p t_1 & \leq & n_p & \leq & \kappa_p t_1 + N_p \\ & & \vdots & & \\ \kappa_M t_1 & \leq & n_M & \leq & \kappa_M t_1 + N_M \end{array} \right. , \begin{array}{c} \mathbf{t} \in \mathcal{R}^{p-1} \\ \mathbf{n} \in \mathcal{I}^{M-p} \end{array} \right\} \tag{7.3}$$

where several coefficients, labelled $\kappa_j, \ldots, \kappa_{p-1}$ and $\kappa_q, \ldots, \kappa_M$, may be zero. This results in five possible classes of dimensions. The first continuous index, t_1, is unbounded; while the other indices are either continuous or discrete and bounded either by 0 and a positive constant, T_i; or by a constant times the first index, $\kappa_i t_1$, and a that plus a positive constant, T_i. This equation can thus describe a wide range of shapes over which it is possible to calculate the output signal.

Similarly, to describe a non-rectangular ROS is terms of an unbounded discrete dimension, the dimensions can be labelled so that

$$V_{ROS(n)} = \left\{ \mathbf{t}, \mathbf{n} \left| \begin{array}{ccccc} \kappa_1 n_p & \le & t_1 & \le & \kappa_1 n_p + T_1 \\ & & \vdots & & \\ \kappa_{p-1} n_p & \le & t_{p-1} & \le & \kappa_{p-1} n_p + T_{p-1} \\ 0 & \le & n_p & \le & \infty, (N_p = \infty) \\ \kappa_{p+1} n_p & \le & n_{p+1} & \le & \kappa_{p+1} n_p + N_{p+1} \\ & & \vdots & & \\ \kappa_M n_p & \le & n_M & \le & \kappa_M n_p + N_M \end{array} \right. , \begin{array}{c} \mathbf{t} \in \mathcal{R}^{p-1} \\ \mathbf{n} \in \mathcal{I}^{M-p} \end{array} \right\} \tag{7.4}$$

where again $\kappa_j, \ldots, \kappa_{p-1} = 0$ and $\kappa_q, \ldots, \kappa_M = 0$. This also results in five possible classes of dimensions which are the same as above except that the unbounded index, n_p, is discrete. Many other, more complexly shaped regions of support are possible, with limits dependent on more than one other dimension.

The rectangular ROS can be defined by

$$\begin{gathered} V_{\widehat{ROS}(k)} = \left\{ \mathbf{t}, \mathbf{n} \left| \begin{array}{ll} 0 \le t_i \le T_i, & \mathbf{t} \in \mathcal{R}^{p-1} \\ 0 \le n_i \le N_i, & \mathbf{n} \in \mathcal{I}^{M-p} \end{array} \right. \right\} \\ \mathbf{T}_{(k)} = \{T_1, \ldots, T_{p-1}\} \\ \mathbf{N}_{(k)} = \{N_p, \ldots, N_M\} \end{gathered} \tag{7.5}$$

where $N_{i \neq k}$ are finite integers and $T_{i \neq k}$ are finite real numbers and $T_k = \infty$ if $1 \le k < p$ or $N_k = \infty$ if $p \le k \le M$.

The impulse response, input and output of a system with the ROS $V_{ROS(t)}$ can be transformed into those with rectangular ROS $V_{\widehat{ROS}(1)}$ as

$$\begin{aligned}
\hat{h}(\mathbf{t},\mathbf{n}) &\equiv h(t_1, t_2+\kappa_2 t_1, \ldots, n_p+\kappa_p t_1, \ldots) = h(\mathbf{t}+\mathbf{K}_t(t_1), \mathbf{n}+\mathbf{K}_n(t_1)) \\
\hat{x}(\mathbf{t},\mathbf{n}) &\equiv x(t_1, t_2+\kappa_2 t_1, \ldots, n_p+\kappa_p t_1, \ldots) = x(\mathbf{t}+\mathbf{K}_t(t_1), \mathbf{n}+\mathbf{K}_n(t_1)) \\
\hat{y}(\mathbf{t},\mathbf{n}) &\equiv y(t_1, t_2+\kappa_2 t_1, \ldots, n_p+\kappa_p t_1, \ldots) = y(\mathbf{t}+\mathbf{K}_t(t_1), \mathbf{n}+\mathbf{K}_n(t_1))
\end{aligned} \tag{7.6}$$

and a system with the ROS $V_{ROS(n)}$ can be transformed into one with a rectangular ROS $V_{\widehat{ROS}(p)}$ as

$$\begin{aligned}
\hat{h}(\mathbf{t},\mathbf{n}) &\equiv h(t_1+\kappa_1 n_p, \ldots, n_p, n_{p+1}+\kappa_{p+1} n_p, \ldots) = h(\mathbf{t}+\mathbf{K}_t(n_p), \mathbf{n}+\mathbf{K}_n(n_p)) \\
\hat{x}(\mathbf{t},\mathbf{n}) &\equiv x(t_1+\kappa_1 n_p, \ldots, n_p, n_{p+1}+\kappa_{p+1} n_p, \ldots) = x(\mathbf{t}+\mathbf{K}_t(n_p), \mathbf{n}+\mathbf{K}_n(n_p)) \\
\hat{y}(\mathbf{t},\mathbf{n}) &\equiv y(t_1+\kappa_1 n_p, \ldots, n_p, n_{p+1}+\kappa_{p+1} n_p, \ldots) = y(\mathbf{t}+\mathbf{K}_t(n_p), \mathbf{n}+\mathbf{K}_n(n_p))
\end{aligned} \tag{7.7}$$

7.2 PBIBO Stability in Rectangular Regions of Support

7.2.1 Discrete Domain Systems

The necessary and sufficient conditions given in [31] for PBIBO stability under any rectangular ROS $V_{\widehat{ROS}(k)}$ that are stated here were derived for a discrete domain system. First, a linear shift invariant M-D discrete domain system is PBIBO stable iff the following M inequalities are satisfied:

$$\sum_{n_1=0}^{N_1} \sum_{n_2=0}^{N_2} \cdots \sum_{n_k=0}^{N_k=\infty} \cdots \sum_{n_M=0}^{N_M} |h(n_1, n_2, \ldots, n_M)| < \infty \tag{7.8}$$

for $k = 1, 2, \ldots, M$ where $N_1, N_2, \ldots, N_{k-1}, N_{k+1}, \ldots, N_M$ are finite positive integers and $h(n_1, n_2, \ldots, n_M)$ is the impulse response. And equivalently in the z-domain, an M-D discrete domain system described by the steady state transfer function

$$H(\mathbf{z}^{-1}) = \frac{A(\mathbf{z}^{-1})}{B(\mathbf{z}^{-1})} \tag{7.9}$$

is PBIBO stable iff

$$B(0,\ldots,0,z_k^{-1},0,\ldots,0) \neq 0 \tag{7.10}$$

for $z_k^{-1} \in \bar{U}, (k = 1,2,\ldots,n)$ where $\bar{U} = \{z||z| \leq 1\}$.

7.2.2 Mixed Domain Systems

Before extending the conditions for PBIBO stability given in equations 7.8 and 7.10 to the mixD case, it is appropriate to prove formally the time domain condition for BIBO stability given in equation 2.34 in chapter 2.

Theorem 7.1 *A linear shift-invariant mixed domain system is BIBO stable iff*

$$\sum_{\mathbf{n}=-\infty}^{\infty} \int_{-\infty}^{\infty} |h(\mathbf{t},\mathbf{n})| \, d\mathbf{t} \leq S_1 < \infty \tag{7.11}$$

where $h(\mathbf{t},\mathbf{n})$ is the unit impulse response of the system and S_1 is a non-infinite real number.

Proof: Using the convolution sum to prove sufficiency yields

$$\begin{aligned}
|y(\mathbf{t},\mathbf{n})| &= \left| \sum_{\mathbf{k}=-\infty}^{\infty} \int_{-\infty}^{\infty} h(\mathbf{r},\mathbf{k})x(\mathbf{t}-\mathbf{r},\mathbf{n}-\mathbf{k}) \, d\mathbf{r} \right| \\
&\leq \sum_{\mathbf{k}=-\infty}^{\infty} \int_{-\infty}^{\infty} |h(\mathbf{r},\mathbf{k})| \, |x(\mathbf{t}-\mathbf{r},\mathbf{n}-\mathbf{k})| \, d\mathbf{r} \\
&\leq \max_{\mathbf{t},\mathbf{n}}(|x(\mathbf{t},\mathbf{n})|) \sum_{\mathbf{k}=-\infty}^{\infty} \int_{-\infty}^{\infty} |h(\mathbf{r},\mathbf{k})| \, d\mathbf{r} \\
&< \infty.
\end{aligned} \tag{7.12}$$

To prove necessity, assume that

$$\sum_{\mathbf{n}=-\infty}^{\infty} \int_{-\infty}^{\infty} |h(\mathbf{t},\mathbf{n})| \, d\mathbf{t} = \infty \tag{7.13}$$

and that for some $\mathbf{t}, \mathbf{n}$, $x(\mathbf{t}-\mathbf{r}, \mathbf{n}-\mathbf{k}) = \mathrm{sign}(h(\mathbf{r}, \mathbf{k}))$. Then

$$\begin{aligned} y(\mathbf{t}, \mathbf{n}) &= \sum_{\mathbf{k}=-\infty}^{\infty} \int_{-\infty}^{\infty} h(\mathbf{t}, \mathbf{n}) \mathrm{sign}(h(\mathbf{r}, \mathbf{k})) \\ &= \sum_{\mathbf{k}=-\infty}^{\infty} \int_{-\infty}^{\infty} |h(\mathbf{t}, \mathbf{n})| \\ &= \infty. \end{aligned} \tag{7.14}$$

If, on the other hand, $x(\mathbf{t}, \mathbf{n})$ and $h(\mathbf{t}, \mathbf{n})$ are constrained to be causal, then with the same definitions $\lim_{\mathbf{t}, \mathbf{n} \to \infty} y(\mathbf{t}, \mathbf{n}) = \infty$. QED.

The following condition for PBIBO stability is much less restrictive and thus of more use.

Theorem 7.2 *A linear shift-invariant mixed domain system is PBIBO stable over the rectangular region of support $V_{\widehat{ROS}(k)}$ iff*

$$\sum_{\mathbf{n}=0}^{\mathbf{N}_{(k)}} \int_{0}^{\mathbf{T}_{(k)}} |h(\mathbf{t}, \mathbf{n})| \, d\mathbf{t} \leq S_1 < \infty \tag{7.15}$$

where $h(\mathbf{t}, \mathbf{n})$ is the impulse response of the system, S_1 is a non-infinite real number and $V_{\widehat{ROS}(k)}$, $\mathbf{N}_{(k)}$ and $\mathbf{T}_{(k)}$ are defined by equation 7.5.

Proof: Using the convolution sum to prove sufficiency yields

$$\begin{aligned} |y(\mathbf{t}, \mathbf{n})| &= \left| \sum_{\mathbf{k}=0}^{\mathbf{n}} \int_{0}^{\mathbf{t}} h(\mathbf{r}, \mathbf{k}) x(\mathbf{t}-\mathbf{r}, \mathbf{n}-\mathbf{k}) \, d\mathbf{r} \right| \\ &\leq \sum_{\mathbf{k}=0}^{\mathbf{N}_{(k)}} \int_{0}^{\mathbf{T}_{(k)}} |h(\mathbf{r}, \mathbf{k})| \, |x(\mathbf{t}-\mathbf{r}, \mathbf{n}-\mathbf{k})| \, d\mathbf{r} \\ &\leq \max_{\mathbf{t}, \mathbf{n}}(|x(\mathbf{t}, \mathbf{n})|) \sum_{\mathbf{k}=0}^{\mathbf{N}_{(k)}} \int_{0}^{\mathbf{T}_{(k)}} |h(\mathbf{r}, \mathbf{k})| \, d\mathbf{r} \\ &< \infty. \end{aligned} \tag{7.16}$$

Necessity can be shown exactly as in equation 7.14 with $\mathbf{t}, \mathbf{n}$ approaching $\mathbf{T}_{(k)}, \mathbf{N}_{(k)}$.

QED.

MixD filters are usually designed in the SZ-domain. so the following condition is usually of more use.

Theorem 7.3 *A linear shift-invariant M-D mixed domain system described by the transfer function*

$$H(\mathbf{s},\mathbf{z}^{-1}) = \frac{A(\mathbf{s},\mathbf{z}^{-1})}{B(\mathbf{s},\mathbf{z}^{-1})} \tag{7.17}$$

is PBIBO stable over the rectangular region of support $V_{\widehat{ROS}(k)}$ iff

$$\begin{gathered} B(\underbrace{1,\ldots,s_k,\ldots,1}_{1\ to\ p-1},\underbrace{0,\ldots,0}_{p\ to\ M}) \neq 0, Re[s_k] \geq 0 \quad \text{if } 1 \leq k < p \\ \text{or} \\ B(\underbrace{1,\ldots,1}_{1\ to\ p-1},\underbrace{0,\ldots,z_k^{-l},\ldots,0}_{p\ to\ M}) \neq 0, |z_k^{-1}| \leq 1 \quad \text{if } p \leq k \leq M. \end{gathered} \tag{7.18}$$

Proof: The proof closely parallels the proof of theorem 2 in [31] and theorem 4 in [26].

From the fact that $B(1,\ldots,1,0,\ldots,0) \neq 0$ there exists a region $A_\varepsilon^M = \{\mathbf{s},\mathbf{z}^{-1} \mid |s_i - 1| < \varepsilon, |z_i^{-1}| < \varepsilon\}$ where $H(\mathbf{s},\mathbf{z}^{-1})$ is analytic and there exists an SZ-transform

$$H(\mathbf{s},\mathbf{z}^{-1}) = \sum_{\mathbf{n}=0}^{\infty} \int_0^{\infty} h(\mathbf{t},\mathbf{n}) e^{-\mathbf{s}^T\mathbf{t}} z^{-\mathbf{n}}\, d\mathbf{t} \tag{7.19}$$

which can be differentiated term-wise in each z_i^{-1}, $i \neq k$, n_i times

$$\left.\frac{\partial^{n_p+\cdots+n_{i\neq k}+\cdots+n_M}}{\partial z_p^{-n_p}\cdots\partial z_{i\neq k}^{-n_i}\cdots\partial z_M^{-n_M}} H(\mathbf{s},\mathbf{z}^{-1})\right|_{\substack{z_p^{-n_p}=\cdots=z_{i\neq k}^{-n_i}=\cdots=z_M^{-n_M}=0 \\ s_1=\cdots=s_{j\neq k}=\cdots=s_{p-1}=1}} \tag{7.20}$$

$$= \begin{cases} (n_p!)\cdots(n_M!)\displaystyle\int_0^{\infty} e^{-\mathbf{t}_{(-k)}} \int_0^{\infty} h(\mathbf{t},\mathbf{n}) e^{-s_k t_k}\, dt_k\, d\mathbf{t}_{(-\mathbf{k})} & \text{if } 1 \leq k < p \\ (n_p!)\cdots(n_{i\neq k}!)\cdots(n_M!)\displaystyle\int_0^{\infty} e^{-\mathbf{t}} \sum_{n_k=0}^{\infty} h(\mathbf{t},\mathbf{n}) z_k^{-n_k}\, d\mathbf{t} & \text{if } p \leq k \leq M. \end{cases}$$

where $\mathbf{t}_{(-k)}$ is the vector $\mathbf{t}$ without the entry t_k. In the case that $1 \leq k < p$ the left hand side is a rational function of the form

$$\frac{A_k(s_k)}{[B(1,\ldots,s_k,\ldots,1,0,\ldots,0)]^{n_p+\cdots+n_M}} \tag{7.21}$$

which, if equation 7.18 holds, is a bounded 1-D transfer function so

$$\int_0^\infty |h(\mathbf{t}, \mathbf{n})| dt_k < \infty \tag{7.22}$$

for all $t_{i \neq k}, n_i$ finite. In the case that $p \leq k \leq M$, the left hand side of equation 7.20 is a rational function of the form

$$\frac{A_k(z_k^{-1})}{[B(1, \ldots, 1, 0, \ldots, z_k^{-1}, \ldots, 0)]^{n_p + \cdots + n_{i \neq k} + \cdots + n_M}} \tag{7.23}$$

which, if equation 7.18 holds, is a bounded 1-D transfer function so

$$\sum_{n_k=0}^{\infty} |h(\mathbf{t}, \mathbf{n})| < \infty \tag{7.24}$$

for all $t_i, n_{i \neq k}$ finite. Equation 7.15 follows as the sum of a finite number of integrals over finite intervals of elements of the form of equation 7.22 or 7.24. QED.

PBIBO stability for discrete domain and continuous domain systems are a special case of the mixD results presented above. They are, however, restricted to rectangular regions of support.

7.3 PBIBO Stability in Non-Rectangular Regions of Support for Mixed Domain Systems

The conditons for PBIBO stability with non-rectangular regions of support for mixD systems are similar to (7.8) and (7.10) in both form and proof, and are relatively simple extensions of the concepts presented in [31]. The conditions in the spatio-temporal domain are given by theorem 7.4.

Theorem 7.4 *A linear shift-invariant M-D mixed domain system is practical-BIBO stable over the ROS $V_{ROS(t)}$ or $V_{ROS(n)}$ iff the following inequality is satisfied:*

$$\int_0^{\mathbf{T}_{(k)}} \sum_0^{\mathbf{N}_{(k)}} |\hat{h}(\mathbf{t}, \mathbf{n})| \, d\mathbf{t} < \infty \tag{7.25}$$

where $k = 1$ for $V_{ROS(t)}$ and $k = p$ for $V_{ROS(n)}$. $\hat{h}(\mathbf{t}, \mathbf{n})$ is the transformed impulse response defined by equation 7.6 or 7.7 and $\mathbf{T}_{(k)}$ and $\mathbf{N}_{(k)}$ are defined by equation 7.5.

Proof: As in [31], we use the M-D convolution integral-sum to show sufficiency. For the $V_{ROS(n)}$ case:

$$\begin{aligned}
y(\mathbf{t},\mathbf{n}) &= \sum_{\mathbf{k}=-\infty}^{\infty}\int_{-\infty}^{\infty} h(\mathbf{t}-\mathbf{r},\mathbf{n}-\mathbf{k})x(\mathbf{r},\mathbf{k}), \quad \mathbf{t},\mathbf{n}\in V_{ROS(n)} \\
&= \sum_{\mathbf{k}=\mathbf{K}_n(k_p)}^{\mathbf{n}+\mathbf{K}_n(k_p-n_p)}\int_{\mathbf{K}_t(k_p)}^{\mathbf{t}+\mathbf{K}_t(k_p-n_p)} h(\mathbf{t}-\mathbf{r},\mathbf{n}-\mathbf{k})x(\mathbf{r},\mathbf{k}), \quad \mathbf{t},\mathbf{n}\in V_{ROS(n)} \\
\hat{y}(\mathbf{t},\mathbf{n}) &= \sum_{\mathbf{k}=\mathbf{K}_n(k_p)}^{\mathbf{n}+\mathbf{K}_n(k_p)}\int_{\mathbf{K}_t(k_p)}^{\mathbf{t}+\mathbf{K}_t(k_p)} h(\mathbf{t}+\mathbf{K}_t(n_p)-\mathbf{r},\mathbf{n}+\mathbf{K}_n(n_p)-\mathbf{k})x(\mathbf{r},\mathbf{k}), \quad \mathbf{t},\mathbf{n}\in V_{\widehat{ROS}(p)} \\
&= \sum_{\mathbf{k}=0}^{\mathbf{n}}\int_{0}^{\mathbf{t}} h(\mathbf{t}+\mathbf{K}_t(n_p-k_p)-\mathbf{r},\mathbf{n}+\mathbf{K}(n_p-k_p)-\mathbf{k})x(\mathbf{r}+\mathbf{K}_t(k_p),\mathbf{k}+\mathbf{K}_n(k_p)) \\
&= \sum_{\mathbf{k}=0}^{\mathbf{n}}\int_{0}^{\mathbf{t}} \hat{h}(\mathbf{t}-\mathbf{r},\mathbf{n}-\mathbf{k})\hat{x}(\mathbf{r},\mathbf{k}), \quad \mathbf{t},\mathbf{n}\in V_{\widehat{ROS}(p)} \qquad (7.26)
\end{aligned}$$

The $V_{ROS(t)}$ case is similar, with manipulations via $\mathbf{K}_t(t_1)$ and $\mathbf{K}_n(t_1)$ and resulting in $\mathbf{n},\mathbf{t}\in V_{\widehat{ROS}(t)}$. If $\hat{y}$ is bounded for bounded $\hat{x}$, then y is bounded for bounded x and the system is PBIBO stable. If x is bounded by S then

$$\begin{aligned}
|\hat{y}(\mathbf{t},\mathbf{n})| &= \left|\sum_{\mathbf{k}=0}^{\mathbf{n}}\int_{0}^{\mathbf{t}} \hat{h}(\mathbf{t}-\mathbf{r},\mathbf{n}-\mathbf{k})\hat{x}(\mathbf{r},\mathbf{k})\right| d\mathbf{r}, \quad \mathbf{t},\mathbf{n}\in V_{\widehat{ROS}(n)} \\
&\le \max_{\mathbf{t},\mathbf{n}}\{|\hat{x}(\mathbf{t},\mathbf{n})|\}\sum_{\mathbf{k}=0}^{\mathbf{n}}\int_{0}^{\mathbf{t}} |\hat{h}(\mathbf{r},\mathbf{k})|d\mathbf{r}, \quad \mathbf{t},\mathbf{n}\in V_{\widehat{ROS}(n)} \\
&\le S\sum_{\mathbf{k}=0}^{\mathbf{N}_{(p)}}\int_{0}^{\mathbf{T}_{(p)}} |\hat{h}(\mathbf{r},\mathbf{k})|d\mathbf{r}, \quad \mathbf{t},\mathbf{n}\in V_{\widehat{ROS}(n)} \\
&< \infty \qquad (7.27)
\end{aligned}$$

from (7.25). The $V_{ROS(t)}$ case is again similar.

To prove necessity, assume

$$\sum_{\mathbf{n}=0}^{\mathbf{N}_{(p)}}\int_{0}^{\mathbf{T}_{(p)}} |\hat{h}(\mathbf{t},\mathbf{n})| = \infty \qquad (7.28)$$

and consider the signal

$$x(\mathbf{T}_{(p)}-\mathbf{t},\mathbf{N}_{(p)}-\mathbf{n}) = \mathrm{sign}(h(\mathbf{t},\mathbf{n})), \quad \mathbf{t},\mathbf{n}\in V_{ROS(p)} \qquad (7.29)$$

as input. This gives

$$\lim_{N_p\to\infty} |\hat{y}(\mathbf{T}_{(k)}, \mathbf{N}_{(k)})| = \lim_{N_p\to\infty} \sum_{\mathbf{n}=0}^{\mathbf{N}_{(k)}} \int_{0}^{\mathbf{T}_{(k)}} |\hat{h}(\mathbf{t},\mathbf{n})| = \infty. \tag{7.30}$$

QED.

The conditions corresponding to theorem 7.4 in the SZ-domain are given by theorem 7.5.

Theorem 7.5 *A linear shift-invariant M-D mixed domain system described by the transfer function given in equation 7.17 is PBIBO stable over the ROS $V_{ROS(t)}$ iff*

$$B(s_1 - K_t(1)^T \mathbf{s}_{(-1)} - K_n(1)^T \ln \mathbf{z}, \mathbf{s}_{(-1)}, \mathbf{z}^{-1})\Big|_{s_{i\neq 1}=1, z_i^{-1}=0} \neq 0, \forall \mathrm{Re}[s_1] \geq 0 \tag{7.31}$$

or over the ROS $V_{ROS(n)}$ iff

$$B(\mathbf{s}, z_p^{-1} e^{\mathbf{s}^T K_t(1)} \mathbf{z}_{(-p)}^{K_n(1)}, \mathbf{z}_{(-p)})\Big|_{s_i=1, z_{i\neq p}^{-1}=0} \neq 0, \forall |z_p^{-1}| \leq 1 \tag{7.32}$$

where $\mathbf{s}_{(-1)}$ is $\mathbf{s}$ without the entry s_1 and $\mathbf{z}_{(-p)}$ is $\mathbf{z}$ without the entry z_p.

These conditions correspond to transforming the SZ-domain system from the non-rectangular region of support to a rectangular region of support and then applying theorem 7.3.

Proof: The transfer function, $H(\mathbf{s}, \mathbf{z}^{-1})$ is equal to the SZ transform of the impulse response, $h(\mathbf{t}, \mathbf{n})$. Defining the transformed transfer function, $\hat{H}(\mathbf{s}, \mathbf{z}^{-1})$ as the SZ transform of the transformed impulse response, $\hat{h}(\mathbf{t}, \mathbf{n})$, we have for the $V_{ROS(t)}$ case

$$\begin{aligned}
\hat{H}(\mathbf{s}, \mathbf{z}^{-1}) &= \int_{-\infty}^{\infty} \sum_{\mathbf{n}=-\infty}^{\infty} \hat{h}(\mathbf{r}, \mathbf{k}) e^{-\mathbf{s}^T \mathbf{r}} \mathbf{z}^{-\mathbf{k}} \, d\mathbf{r} \\
&= \int_{\infty}^{\infty} \sum_{\mathbf{n}=-\infty}^{\infty} h(\mathbf{r} + \mathbf{K}_t(t_1), \mathbf{k} + \mathbf{K}_n(t_1)) e^{-\mathbf{s}^T \mathbf{r}} \mathbf{z}^{-\mathbf{k}} \, d\mathbf{r} \\
&= \int_{-\infty}^{\infty} \sum_{\mathbf{n}=-\infty}^{\infty} h(\mathbf{t}, \mathbf{n}) e^{-\mathbf{s}^T (\mathbf{t} - \mathbf{K}_t(t_1))} \mathbf{z}^{-(\mathbf{n} - \mathbf{K}_n(t_1))} \, d\mathbf{t} \\
&= H(s_1 - \mathbf{K}_t(1)^T \mathbf{s}_{(-1)} - \mathbf{K}_n(1)^T \ln \mathbf{z}, \mathbf{s}_{(-1)}, \mathbf{z}^{-1})
\end{aligned} \tag{7.33}$$

and for the $V_{ROS(n)}$ case

$$\begin{aligned}
\hat{H}(\mathbf{s},\mathbf{z}^{-1}) &= \int_{-\infty}^{\infty} \sum_{\mathbf{n}=-\infty}^{\infty} \hat{h}(\mathbf{r},\mathbf{k}) e^{-\mathbf{s}^T\mathbf{r}} \mathbf{z}^{-\mathbf{k}}\, d\mathbf{r} \\
&= \int_{-\infty}^{\infty} \sum_{\mathbf{n}=-\infty}^{\infty} h(\mathbf{r}+\mathbf{K}_t(n_p),\mathbf{k}+\mathbf{K}_n(n_p)) e^{-\mathbf{s}^T\mathbf{r}} \mathbf{z}^{-\mathbf{k}}\, d\mathbf{r} \\
&= \int_{-\infty}^{\infty} \sum_{\mathbf{n}=-\infty}^{\infty} h(\mathbf{t},\mathbf{n}) e^{-\mathbf{s}^T(\mathbf{t}-\mathbf{K}_t(n_p))} \mathbf{z}^{-(\mathbf{n}-\mathbf{K}_n(n_p))}\, d\mathbf{t} \\
&= H(\mathbf{s}, z_p^{-1} e^{\mathbf{s}^T\mathbf{K}_t(1)} \mathbf{z}_{(-p)}^{\mathbf{K}_n(1)}, \mathbf{z}_{(-p)})
\end{aligned} \tag{7.34}$$

If $\hat{H}(\mathbf{s},\mathbf{z}^{-1})$ is PBIBO stable over the ROS $V_{\widehat{ROS}(1)}$ then $H(\mathbf{s},\mathbf{z}^{-1})$ is PBIBO stable over the ROS $V_{ROS(t)}$, and similarly for the second case. The functions B in equations 7.31 and 7.32 are the denominators of $\hat{H}$ for the two cases respectively and the substitutions and conditions correspond to theorem 7.3. QED.

7.4 Design of PBIBO Stable Systems from Continuous Positive M-D Networks

In [31] the authors show that M-D discrete domain systems derived from continuous domain positive M-D networks via the bilinear transform are always PBIBO stable, given a rectangular ROS, though they may not always be BIBO stable [52]. The same arguments can be applied to show that mixD systems derived from continuous positive M-D networks are PBIBO stable in both rectangular regions of support and some non-rectangular regions of support.

The same continuous M-D structure as in [31], which is a general M-D reactance 2-port as shown in Figure 7.2, containing only unit M-D capacitors, resistors, gyrators and transformers, is considered here. All M-D inductors and capacitors can be realized as combinations of gyrators, transformers and unit capacitors. The input is a voltage across port 1 and the output is the voltage across a non-zero resistor terminating port

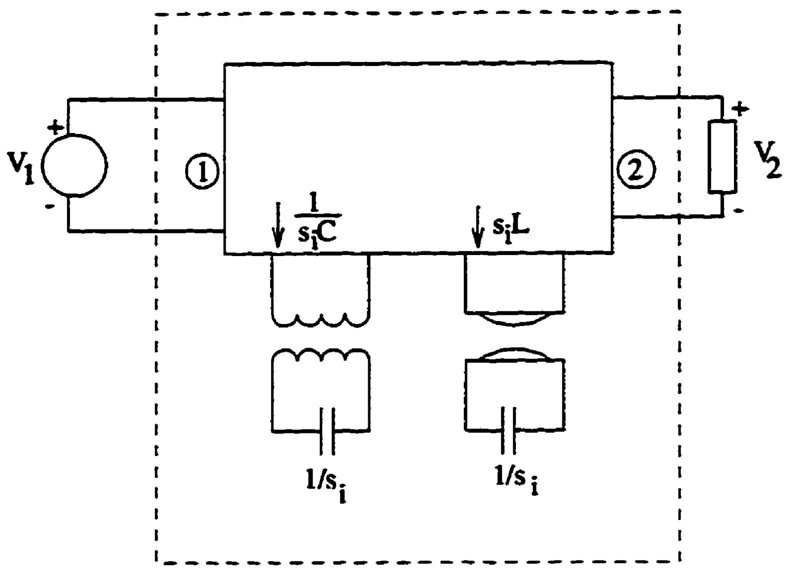

Figure 7.2: A M-D Reactance 2-Port Terminated in a Resistance

2 so the general M-D transfer function can be written

$$T(\mathbf{s}) = \frac{V_2(\mathbf{s})}{V_1(\mathbf{s})}. \tag{7.35}$$

This continuous domain system can then be transformed into a mixD system by applying the bilinear transform (equation 2.22) or the modified bilinear transform (equation 3.10) to each dimension which is to be discrete in the final design.

7.4.1 Mixed Domain Systems under Rectangular Regions of Support

Theorem 7.6 is the extension of theorem 3 in [31] to mixD systems.

Theorem 7.6 *The voltage-transfer function of an M-D reactance 2-port, where the output voltage is across a resistor, leads to a PBIBO stable mixed domain system over any rectangular region of support,* $V_{\widetilde{ROS}(k)}$, *after the application of the bilinear or modified bilinear transformation to* $M-(p-1)$ *dimensions.*

Proof: If the continuous domain transfer function of the M-D reactance 2-port, as shown in Figure 7.2, is given by equation 7.35 then the 1-D transfer function $T_k(s_k)$

From equation 7.40 we have that

$$B(s_1, z_p^{-1}e^{s_1 K_{t1}(1)}\mathbf{z}_{(-p)}^{\mathbf{K}_n(1)}, \mathbf{z}_{(-p)}^{-1}) \neq 0, \forall \left| z_p^{-1}e^{s_1 K_{t1}(1)}\mathbf{z}_{(-p)}^{\mathbf{K}_n(1)} \right| \leq 1, s_1 = 1, \mathbf{z}_{(-p)}^{-1} = \mathbf{0} \tag{7.42}$$

Again, since z_p appears only in the form of the bilinear transformation no cancellation can occur and the condition holds if $|e^{K_{t1}(1)}| \leq 1$, all the elements of $K_n(1)$ are negative and $|z_p^{-1}| \leq 1$. Theorem 7.5 then applies. QED.

7.5 Conclusions and Further Work

In this chapter a number of conditions for the stability of mixD systems are derived. These include conditions in the spatio-temporal and SZ-domains for PBIBO stability under both rectangular and non-rectangular regions of support. Also, filters designed from continuous domain M-D reactance 2-ports are shown to be PBIBO stable after bilinear transformation into mixD filters under rectangular and some non-rectangular regions of support.

A major concern and the primary motivation for this work concerns the stability of the IDD and Bowl filters discussed in chapters 3 and 4 with delay variations. In section 3.5 it was shown that changes in the delay element lengths corresponded to applying the original filter to skewed inputs over a non-rectangular region of support. In terms of equation 3.11, the region of support $V_{ROS(n)}$ for these filters is given by

$$V_{ROS(n)} = \left\{ \mathbf{t}, \mathbf{n} \left| \begin{array}{ccccc} -m_1 n_3 & \leq & t_1 & \leq & -m_1 n_3 + T_1 \\ -m_2 n_3 & \leq & n_2 & \leq & -m_2 n_3 + N_2 \\ 0 & \leq & n_3 & \leq & \infty, (N_3 = \infty) \end{array} \right. , \begin{array}{c} \mathbf{t} \in \mathcal{R} \\ \mathbf{n} \in \mathcal{I}^2 \end{array} \right\} \tag{7.43}$$

so that, according to theorem 7.7 the filters will be stable if the delay lengths are increased, but not if they are decreased.

The Challenger Real-Time Video Processor [1, 55] was used to implement the fully digital version of the frequency planar pass filter (using the transfer function coefficients given in table 6.1) with both rectangular and non-rectangular regions of support by varying the length of the frame delay. It was confirmed that increasing the length of the delay led to stable responses and decreasing the length of the delay lead to unstable responses, as expected by theorem 7.7. The trials were not exhaustive.

However, using the Endeavour Analogue Video Filter, the opposite was found: increasing the length of the frame delay resulted in unstable responses and decreasing it resulted in stable responses. The results shown in chapter 5 confirm this. The reasons for this are not understood, but the effect is advantageous, since lengthening the frame delay does not skew the passband into octants other than the first while shortening it does.

Chapter 8

Conclusions and Recommendations for Further Research

8.1 Conclusions

In this thesis the design, implemenation and testing of real-time, controllable, analogue 3-D LT filters for video signals is described. These LT filters enhance or reject signals on the basis of their velocity and are intended as a building block for various applications in which velocity information is important. Because of the large amount of data in video signals, the digital, discrete domain approach to processing them requires a great deal of expensive hardware to achieve real-time operation. The analogue, mixD approach is intended to overcome this obstacle, using the inherent speed and lower cost of analogue computing elements. The major improvement of the filters described in this thesis over previous implementations is that they are controllable remotely in real-time, rather than by the manipulation of potentiometers and switches. Also, this is the first time that the 3-D frequency response of a real-time filter for video signals has been experimentally measured in an automated manner and over a large portion of the frequency domain. In the interest of improving the performance of the filters, a number of sensitivity issues unique to mixD filters are investigated. Finally, the stability of these filters under manipulation of the delay elements, which enable them to be much more flexible, is addressed.

Chapters one and two present an introduction to the basic mathematical tools needed in the design of mixD systems; to the class of linear trajectory signals, on which the velocity selectivity of the filters to be designed is based; and to raster scanning and in particular the NTSC raster scan video format, which is used for input and output to the system. Objects in video sequences that travel at a constant speed in a straight line belong to the class of linear trajectory signals. In the frequency domain, the energy in these signals lies entirely in a plane related to the object's velocity, which is the connection between frequency planar filters, LT filters and velocity selective filters. The mixD approach to processing video signals is shown to be advantageous in that the most common format for broadcast and storage of video sequences, NTSC raster scan format, is a mixD signal. MixD filters can process this raster scanned signal directly, without the sampling and conversion needed by digital systems. The raster scan operation is treated mathematically as a transformation of variables, which is used in later chapters in the design process, to explain the effects of manipulating the delay elements, and to derive the frequency response measurement technique.

Chapter three describes the design of two linear trajectory mixD filters using the ladder form structure. The first order frequency planar pass filter is the simplest linear trajectory filter and is derived from a continuous domain prototype containing one 3-D inductor and one resistor. The second order frequency bowl pass filter has better velocity selectivity in the low frequency region, but uses more operating elements. It is derived from a continuous domain prototype containing two 3-D inductors and two resistors. The only difference between these designs and those developed in [4, 5] is the modification of the bilinear transform, which increases the stability margin and reduces sensitivity to delay element gain errors. In their original form these filters are limited to passband trajectories that lie in the first octant. A technique to overcome this limitation by changing the length of the delay elements is discussed and its effects described with a simple model.

A major goal of this research work was to develop a method to provide parameter control of the real-time filter, while in operation, without having to manipulate the circuit elements mechanically. In chapter four, the implementation of the filters designed in chapter three, using wide bandwidth analogue multipliers and high gain-bandwidth product operational amplifiers as computing elements, is described. The filter parameters are supplied as currents to the analogue multipliers either directly from an exterior source or from onboard D/A converters. The D/A converters, along with some other registers which determine filter structure and delay element operation, can be controlled via a GUI running on a workstation to which the filters are connected. The hardware and software that makes up this implementation is called the Endeavour Analogue Video Filter.

Chapter five describes the characterization of the response of the filters in both the spatio-temporal domain and the frequency domain. They are shown to be effective, visually, in enhancing objects with a given velocity and attenuating others. The filters are prone to instability due to overflow non-linearities, so some suggestions for improvement are made. Also, the issue of calibration is touched on. The technique for measuring the 3-D frequency response of a raster scan based filter suggested in [5] is applied, for the first time, to the filters; and magnitude responses for a number of different settings are given. While the measured responses are good, they differ from the ideal, mostly by a 1-D lowpass response. A simple 1-D highpass postfilter and some changes to the circuit are suggested to compensate.

In chapter six, the sensitivity of the filters to errors in circuit component values is investigated. Theoretical lower bounds on the worst case and Schoeffler sensitivities to delay element errors for the implementation of any given transfer function are derived. These can be used to compare different filter structures, before they are constructed. It is shown that the ladder form structure used in the Endeavour Analogue Video Filter has excellent sensitivity properties and that the modification to the bilinear transform introduced in chapter three to increase the stability margin also reduces the sensitivity of the filter to delay element gain errors.

In chapter seven, the conditions for PBIBO stability developed in [31] for discrete domain filters are extended to mixD filters and to non-rectangular regions of support. It is shown that mixD filters derived from passive continuous domain prototypes via the bilinear or modified bilinear transforms are PBIBO stable under all rectangular and some non-rectangular regions of support. Non-rectangular regions of support arise in the case of mixD filters when the length of the delay elements are manipulated. The stability conditions were tested both with a discrete domain implementation, for which they held; and with the Endeavour Analogue Video Filter, for which they were *reversed*. The reason for this is not understood and should be persued in future research in the area.

8.2 Recommendations for Future Research

The field of mixD filter design and implementation is quite new, so many areas of interest are still unexplored. Even in the areas that have been examined the results are preliminary and invite further exploration. Some specific areas that may be of interest are suggested below.

To further reduce the cost and power requirements of implementing these filters, it would be of interest to explore alternate implementations of the delay elements. Several options have recently become available or may be possible in the near future, including analogue memories and large digital FIFO memories. Currently the size and persistance of the analogue memories are too small for a frame delay, but progress is being made. With recently released digital FIFO products it is possible to reduce the number of ICs associated with the delays from 178 to about 34. Another major reduction in size and power would come from a VLSI implementation of the analogue functions.

Improvements to the implementation might include a fixed lossy integrator, in place of the current variable one. This would eliminate the delay through the feedback loop that is probably causing the extra lowpass characteristic measured in chapter five. The passband would still be completely controllable, though the bandwidth would only be adjustable in discrete steps. It would also reduce the analogue hardware requirements significantly. Another solution to this problem would be a 1-D highpass postfilter.

The results of the theoretic investigation into stability under delay variations in chapter seven are inconclusive. Further study would be appropriate.

Finally, one of the purposes of developing these filters is the implementation of affordable adaptive tracking algorithms, classification systems and other applications. An adaptor module added to the present platform would be instrumental in achieving this end.

References

[1] C. J. Kulach, L. T. Bruton, and N. R. Bartley. Real time implementation of a general 3-dimensional first-order recursive discrete-time transfer function for video. *IEEE Intl. Symp. on Circuits and Systems*, 1996.

[2] M. A. Sid-Ahmed. Two-dimensional analog filters: A new form of realization. *IEEE Trans. on Circuits and Systems*, 36(1), January 1989.

[3] Roger K. Bertschmann. Design and implementation of analog 2-D linear trajectory filters. Master's thesis, The University of Calgary, 1993.

[4] R. K. Bertschmann, N. R. Bartley, and L. T. Bruton. A 3-D integrator-differentiator double-loop (IDD) filter for raster-scan video processing. *IEEE Trans. on Circuits and Systems for Video Technology*. Under Review.

[5] M.-K. Najafi-Koopai. On the frequency response measurement of raster-scanned video systems and motion estimation using three-dimensional LT filters. Master's thesis, The University of Calgary, 1994.

[6] C. Koch. Implementing early vision algorithm in analog hardware: An overview. *Visual Information Processing: From Neurons to Chips, SPIE*, 1473:2–16, 1991.

[7] K. Strohbehn, R. C. Meitzler, A. G. Andreou, and R. E. Jenkins. Analog image processing with silicon retinas. *Johns Hopkins APL Technical Digest*, 15(3):178–186, 1994.

[8] C. P. Chong, C. A. T. Salama, and K. C. Smith. Image-motion detection using analog VLSI. *IEEE Journal of Solid-State Circuits*, 27(1):93–96, January 1992.

[9] L. T. Bruton and N. R. Bartley. Three dimensional image processing using the concept of network resonance. *IEEE Trans. Circuits and Systems*, CAS-32(7):664–672, 1985.

[10] L. T. Bruton and N. R. Bartley. The enhancement and tracking of moving objects in digital images using adaptive three-dimensional recursive filters. *IEEE Trans. on Circuits and Systems*, CAS-33(6), June 1986.

[11] Y. Zhang and L. T. Bruton. Differentiator-type three-dimensional recursive ladder filters having frequency-planar or frequency-beam-shaped passbands. *IEEE Trans. on Circuits and Systems for Video Technology*, 2(3), September 1992.

[12] Y. Zhang. *Three-Dimensional Recursive Digital Filters for the Processing of Image Sequences.* PhD thesis, The University of Calgary, 1993.

[13] T. J. Fowlow and L. T. Bruton. The design and application of a high quality three dimensional linear trajectory filter. *Proc. of the Int. Symp. on Circuits and Systems (ISCAS '88)*, 1988.

[14] L. T. Bruton, N. R. Bartley, and Z. Q. Liu. On the classification of moving objects in image sequences using 3-D adaptive recursive tracking filters and neural networks. *Proc. of Asilomar Conf. on Signals, Systems, and Computers*, 2:1006–1010, October 1995.

[15] National Semiconductor Corporation, 2900 Semiconductor Drive, P.O. Box 58090, Santa Clara CA. *1993 Video Products Databook.*

[16] N. K. Bose. *Applied Multidimensional Systems Theory.* Van Nostrand Reinhold, Scarborough ON, 1982.

[17] D. E. Dudgeon and R. M. Mersereau. *Multidimensional Digital Signal Processing.* Prentice-Hall, Inc., Englewood Cliffs, NJ, USA, 1984.

[18] N. K. Bose. Problems and progress in multidimensional systems theory. *Proc. of the IEEE*, 65(6):824–840, June 1977.

[19] N. K. Bose. Multidimensional digital signal processing: Problems, progress, and future scope. *Proc. of the IEEE*, 78(4), April 1990.

[20] M. Ohki, M. E. Zervakis, and A. N. Venetsanopoulos. 3-D digital filters. In C. T. Leondes, editor, *Control and Dynamic Systems, Vol. 69, Multidimensional Systems: Signal Processing and Modeling Techniques*, pages 49–88. Academic Press, 1995.

[21] A. Fettweis. Multidimensional circuit and systems theory. *Proc. of the Int. Symp. on Circuits and Systems*, pages 951–957, May 1984.

[22] P. A. Ramamoorthy and L. T. Bruton. Design of stable two-dimensional analogue and digital filters with applications in image processing. *Circuit Theory and Applications*, 7:229–245, 1979.

[23] S. Erfani, M. Ahmadi, and V. Ramachandran. Alternative approach to digital spectral transformations. *IEEE Trans. on Circuits and Systems*, 35(11):1461–1463, November 1988.

[24] L. T. Bruton. On the relationships between analog circuits and multidimensional digital algorithms. *Proc. of the IEEE Int. Symp. on Circuits and Systems*, pages 456–459, May 1986.

[25] P. Agathoklis, L. T. Bruton, and N. R. Bartley. The elimination of spikes in the magnitude frequency response of 2-D discrete filters by increasing the stability margin. *IEEE Trans. on Circuits and Systems*, CAS-32(5):451–458, May 1985.

[26] D. Goodman. Some stability properties of two-dimensional linear shift-invariant digital filters. *IEEE Trans. on Circuits and Systems*, CAS-24(4):201–208, April 1977.

[27] L. T. Bruton. Enel 699: Multidimensional signal processing course notes. Dept. Electrical and Computer Engineering, University of Calgary, 1995.

[28] V. A. Ditkin and A. P. Prudnikov. *Operational Calculus in Two Variables and its Applications*. Pergamon Press, Headington Hill Hall, Oxford, 1962. Translated by D. M. G. Wishart.

[29] L. T. Bruton, N. R. Bartley, and R. A. Stein. Stable 2D recursive filter design using equi-terminated lossless N-port structures. *Proc. European Conf. on Circuit Theory and Design, The Hague, Netherlands*, pages 300–307, August 1981.

[30] Y. Zhang and L. T. Bruton. Aplications of 3-D LCR networks in the design of 3-D recursive filters for processing image sequences. *IEEE Trans. on Circuits and Systems for Video Technology*, 4(4):369–382, August 1994.

[31] P. Agathoklis and Bruton L. T. Practical-BIBO stability of n-dimensional discrete systems. *IEE Proc.*, 130(6), December 1983.

[32] L. Zhu, W. Li, and A. Fettweis. A novel approach to the design of recursive fan filters. *IEEE Trans. Circuits and Systems—II:Analog and Digital Signal Processing*, 42(7):492–497, July 1995.

[33] M. S. Lazar and Bruton L. T. On the practical BIBO stability of multidimensional filters. *Proc. of the IEEE Int. Symp. on Circuits and Systems (ISCAS '93), Chicago, IL*, pages 571–574, May 1993.

[34] A. A. Choudhury and L. T. Bruton. Multidimensional filtering using combined discrete fourier transform and linear difference equation methods. *IEEE Trans on Circuits and Systems*, 37(2), February 1990.

[35] K.S. Knudsen and L.T. Bruton. Mixed domain filtering of multidimensional signals. *IEEE Trans. on Circuits and Systems for Video Technology*, CSVT-1(3):260–268, September 1992.

[36] M. S. Lazar and Bruton L. T. Three-dimensional linear trajectory filtering using the DWT and the mixed-domain approach. *Proc. of the IEEE Int. Symp. on Circuits and Systems*, pages 2.45–2.48, June 1994.

[37] D. C. Youla, J. D. Rhodes, and P. C. Marston. Driving-point synthesis of resistor-terminated cascades composed of lumped lossless passive 2-ports and commensurate TEM lines. *IEEE Trans. Circuit Theory*, CT-19(6):648–663, November 1972.

[38] D. C. Youla, J. D. Rhodes, and P. C. Marston. Recent developments in the synthesis of a class of lumped-distributed filters. *IEEE Trans. Circuit Theory and Applications*, 1:59–70, 1973.

[39] A. Michael Noll. *Television Technology: Fundamentals and Future Prospects*. Artech House, Inc., Norwood, MA, USA, 1988.

[40] T. L. Marzetta. Fan filters, the 3-D radon transform, and image sequence analysis. *IEEE Trans. Image Processing*, 3(3):253–264, May 1994.

[41] W. T. Freeman and E. H. Adelson. The design and use of steerable filters. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 13(9):891–906, September 1991.

[42] C. S. Gargour and V. Ramachandran. Generation of stable 2-D transfer functions having variable magnitude characteristics. In C. T. Leondes, editor, *Control and Dynamic Systems, Vol. 69, Multidimensional Systems: Signal Processing and Modeling Techniques*, pages 255–297. Academic Press, 1995.

[43] L. T. Bruton and N. R. Bartley. Highly selective three-dimensional recursive beam filters using intersecting resonant planes. *IEEE Trans. Circuits and Systems*, CAS-30(3):190–193, March 1983.

[44] Q. Liu and L.T. Bruton. Sensitivity analysis of 3-D recursive digital beam filter structures. *Proc. of 1988 Asilomar Conf. on Circuits, Systems, and Computers, La Jolla, CA*, November 1988.

[45] Q. Liu and L.T. Bruton. Design of 3-D planar and beam recursive digital filters using spectral transformations. *IEEE Trans. Circuits and Systems*, CAS-36(3):365–374, March 1989.

[46] D. J. Fleet and A. D. Jepson. Hierarchical construction of orientation and velocity selective filters. *IEEE Trans. Pattern Anal. Machine Intell.*, 11(3):315–325, March 1989.

[47] T. L. Marzetta. Uniformly optimal 3-D fan filters for optical moving target detection. *Proc. ICASSP '93*, pages V543–V545, 1993.

[48] L. T. Bruton. *RC-Active Circuits Theory and Design*. Prentice-Hall, Inc., Englewood Cliffs, NJ, USA, 1980.

[49] T. J. Fowlow and L. T. Bruton. Attenuation characteristics of three-dimensional planar-resonant recursive digital filters. *IEEE Trans. on Circuits and Systems*, CAS-35(5):595–599, May 1988.

[50] L. T. Bruton and N. R. Bartley. Using nonessential singularities of the second kind in two-dimensional filter design. *IEEE Trans. Circuits and Systems*, CAS-36(1):113–116, January 1989.

[51] Z. Lin and L.T. Bruton. BIBO stability of inverse 2-D digital filters in the presence of nonessential singularities of the second kind. *IEEE Trans. on Circuits and Systems*, CAS-36(2):244–254, February 1989.

[52] D. Goodman. Some difficulties with the double bilinear transformation in 2-D recursive filter design. *Proc. of the IEEE*, 66(7), July 1978.

[53] M. N. S. Swamy, L. M. Roytman, and E. I. Plotkin. On stability properties of three- and higher dimensional linear shift-invariant digital filters. *IEEE Trans. Circuits and Systems*, CAS-32(9):888–892, September 1985.

[54] L. Wang and D. Xiyu. Nonessential singularities of the second kind and stability for multidimensional digital filters. *Multidimensional Systems and Signal Processing*, 3:363–380, 1992.

[55] C. J. Kulach. Real-time first-order recursive 3-dimensional digital filter for video signals. Master's thesis, The University of Calgary, 1996.

[56] N. Bartley. Technical description of the digitally-sampled analog delay line. Internal technical report, Dept. Electrical and Computer Engineering, University of Calgary, August 1993.

[57] E. Franchi, M. Tartagni, R. Guerrieri, and G. Baccarani. Random access analog memory for early vision. *IEEE Journal of Solid-State Circuits*, 27(7):1105–1109, July 1992.

[58] R. K. Bertschmann, N. R. Bartley, and L. T. Bruton. A 3-D integrator-differentiator double-loop (IDD) filter for raster-scan video processing. *IEEE Trans on Circuits and Systems for Video Technology*, Under Review.

[59] L. T. Bruton, J. W. Haslett, and R. C. Atkins. Nonideal performance of analogue discrete networks. *IEEE Trans. on Circuit Theory and Appl.*, 5:107–118, 1977.