

ON THE RECOVERY OF IMAGES FROM PARTIAL INFORMATION USING  $\nabla^2 G$   
FILTERING

by  
JAMES ALLEN REIMER  
B.Sc.(EE), The University of Manitoba, 1980  
M.Sc., The University of Manitoba, 1982

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF  
THE REQUIREMENTS FOR THE DEGREE OF  
DOCTOR OF PHILOSOPHY

in  
THE FACULTY OF GRADUATE STUDIES  
ELECTRICAL ENGINEERING

We accept this thesis as conforming  
to the required standard

THE UNIVERSITY OF BRITISH COLUMBIA  
September, 1987

©James Allen Reimer, 1987

In presenting this thesis in partial fulfilment of the requirements for an advanced degree at the University of British Columbia, I agree that the Library shall make it freely available for reference and study. I further agree that permission for extensive copying of this thesis for scholarly purposes may be granted by the head of my department or by his or her representatives. It is understood that copying or publication of this thesis for financial gain shall not be allowed without my written permission.

Department of ELECTRICAL ENGINEERING

The University of British Columbia  
1956 Main Mall  
Vancouver, Canada  
V6T 1Y3

Date December 7, 1987

# Abstract

This thesis considers the recovery of a sampled image from partial information, based on the ‘edges’ or zero crossings found in  $\nabla^2 G$  filtered versions of the image. A scheme is presented for separating an image into a family of multiresolution images, using low pass filtering, subsampling, and  $\nabla^2 G$  filtering. A scheme is also presented for merging this family of  $\nabla^2 G$  filtered images to rebuild the original. The recovery of each of the  $\nabla^2 G$  filtered images from their ‘edges’ or zero crossings is then considered.

It has been suggested that  $\nabla^2 G$  filtered images might be characterized by their zero crossing locations. It is shown that  $\nabla^2 G$  filtered images, filtered in 1-D or 2-D are not, in general, uniquely given within a scalar by their zero crossing locations. Two theorems in support of such a suggestion are considered. The differences between the constraints of Logan’s theorem and  $\nabla^2 G$  filtering are considered, and it is shown that the zero crossings which result from these two situations differ significantly in number and location. Logan’s theorem is therefore not applicable to  $\nabla^2 G$  filtered images. A recent theorem by Curtis on the adequacy of zero crossings of 2-D functions is also considered. It is shown that the requirements of Curtis’ theorem are not satisfied by all  $\nabla^2 G$  filtered images. Further, it is shown that it is very difficult to establish if an image meets the requirements of Curtis’ theorem. Examples of different  $\nabla^2 G$  filtered images with the same zero crossings are also presented.

While not all  $\nabla^2 G$  filtered images are uniquely characterized by their zero crossing locations, the practical recovery of real camera images from this partial information is considered. An iterative scheme is developed for the reconstruction of a  $\nabla^2 G$  filtered image from its sampled zero crossings. The zero crossing samples are localized to the

original image sample grid. Experimental results are presented which show that the recovered images, while retaining many of the features of the original, suffer significant loss. It is shown that, in general, the full recovery of these images in a practical situation is not possible from this partial information.

From this experimental experience, it is proposed that  $\nabla^2 G$  filtered images might be practically recovered from their zero crossings, with some additional characterization of the image in the vicinity of each zero crossing point. A simple, non-iterative scheme is developed for extracting a characterization of the  $\nabla^2 G$  filtered image, through the use of an image edge model and a local estimation of a contrast figure in the vicinity of each zero crossing sample. A redrawing algorithm is then used to recover an approximation of the  $\nabla^2 G$  filtered image from its zero crossing locations and the extracted characterizations. This system is evaluated using natural scene and synthetic images. Resulting image quality is good, but is shown to vary depending on the nature of the image. The advantages and disadvantages of this technique are discussed.

The primary shortcoming of the implemented local estimation technique is an assumption of edge independence. A second approach is developed for characterizing the  $\nabla^2 G$  filtered image zero crossings, which eliminates this assumption. This method is based on 2-D filtering, and provides a new technique for the recovery of a  $\nabla^2 G$  filtered image from its sampled zero crossings. The method does not involve iteration or the solution of simultaneous equations. Good image reconstruction is shown for natural scene images, with the  $\nabla^2 G$  filtered image zero crossings localized only to the original image sample grid. The advantages and disadvantages of this technique are discussed.

The application of this recovery from partial information technique is then considered for image compression. A simple coding scheme is developed for representing the zero crossing segments with linear vector segments. A comparative study is then considered, examining the tradeoffs between compression tuning parameters and the resulting recovered image quality.



# Contents

<b>Abstract</b>	<b>ii</b>
<b>Acknowledgements</b>	<b>xi</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Motivation for Study . . . . .	2
1.2 The Underwater Transmission Problem . . . . .	3
1.3 Synopsis of Image Compression Techniques . . . . .	4
1.4 Motivation for $\nabla^2 G$ Approach . . . . .	6
1.5 Scope of Thesis . . . . .	10
<b>2 Implementation of <math>\nabla^2 G</math> Filters</b>	<b>12</b>
2.1 Pyramid Filtering Technique . . . . .	12
2.2 Low Pass Filter Design . . . . .	13
2.3 $\nabla^2 G$ Filter Design . . . . .	16
2.4 Recovering an Image From a $\nabla^2 G$ Pyramid . . . . .	17
2.5 Zero Crossing Detection . . . . .	22
<b>3 Characterizing <math>\nabla^2 G</math> Filtered Images by their Zero Crossings</b>	<b>26</b>
3.1 Logan's Theorem . . . . .	27
3.2 Curtis' Theorem . . . . .	29

3.3	Scale Space Considerations . . . . .	33
3.4	Different $\nabla^2 G$ Filtered Images with the Same ZC Locations . . . . .	34
3.5	Problems in a Sampled Environment . . . . .	36
3.6	Practical Recovery from ZC Locations . . . . .	37
3.7	The Importance of Contrast . . . . .	41
<b>4</b>	<b>Recovery From Zero Crossings Using Local Estimation</b>	<b>44</b>
4.1	Motivation . . . . .	44
4.2	Description of Method . . . . .	47
4.3	Experimental Results . . . . .	51
4.4	Applicability and Limitations . . . . .	54
<b>5</b>	<b>Recovery From Zero Crossings Using Filtering</b>	<b>55</b>
5.1	Description of Method . . . . .	55
5.2	Design of Reconstruction Filter - No Smoothing . . . . .	57
5.3	Design of Inverse Filter - No Smoothing . . . . .	63
5.4	Experimental Results - No Smoothing . . . . .	67
5.5	Design of Reconstruction Filter - Smoothed . . . . .	67
5.6	Design of Inverse Filter - Smoothed . . . . .	70
5.7	Experimental Results - With Smoothing . . . . .	72
5.8	Applicability and Limitations . . . . .	74
<b>6</b>	<b>Application: Image Compression</b>	<b>81</b>
6.1	Representation of Edges With Vector Segments . . . . .	81
6.2	Approximation Parameters . . . . .	88
6.3	Experimental Results . . . . .	89
<b>7</b>	<b>Conclusions and Recommendations</b>	<b>93</b>

7.1 Recommendations For Further Study . . . . .	95
<b>References</b>	<b>99</b>
<b>Appendix</b>	<b>106</b>

# List of Tables

6.1	Image Compression Rates, Corresponding to the ‘Girl’ Image Matrix in Figure 6.1 (in bits/pixel) . . . . .	91
-----	---	----

# List of Figures

2.1	1/2 Bandwidth 1-D Low Pass Filter . . . . .	15
2.2	1/2 Bandwidth 2-D Low Pass Filter . . . . .	15
2.3	$\nabla^2 G$ 1-D Filter . . . . .	18
2.4	$\nabla^2 G$ 2-D Filter . . . . .	18
2.5	$\nabla^2 G$ Family of Filters . . . . .	19
2.6	1-D Sum of $\nabla^2 G$ and Near-dc Channels . . . . .	20
2.7	$\nabla^2 G$ Image Generation and Image Recovery . . . . .	21
2.8	Low Pass Filtered Image Pyramid: ‘Girl’ . . . . .	23
2.9	Low Pass Filtered Image Pyramid: ‘Cake’ . . . . .	23
2.10	$\nabla^2 G$ Filtered Image Pyramid: ‘Girl’ . . . . .	24
2.11	$\nabla^2 G$ Filtered Image Pyramid: ‘Cake’ . . . . .	24
2.12	Comparison: Original Image, and Best Possible Remerged Image: ‘Girl’	25
2.13	Zero Crossing Image Pyramid: ‘Girl’ . . . . .	25
3.1	Ideal Bandpass Filter Step Response . . . . .	30
3.2	$\nabla^2 G$ Filter Step Response . . . . .	30
3.3	Two Edge Image Intensity Profile, and its $\nabla^2 G$ Filtered Response . . . .	35
3.4	Staircase Image Intensity Profile, and its $\nabla^2 G$ Filtered Response . . . .	35
3.5	Image Intensity Profile, and its $\nabla^2 G$ Filtered Response . . . . .	36
3.6	‘Girl’ $\nabla^2 G$ Pyramid Recovered From ZC Locations . . . . .	42

3.7	Comparison: Best Possible Remerged Image, and Image Reconstructed From ZC Locations: 'Girl' . . . . .	42
4.1	Block Diagram: Image Characterization and Reconstruction with Local Estimation . . . . .	47
4.2	Step Edge Model Configuration . . . . .	49
4.3	Quantized Edge Orientations . . . . .	50
4.4	Reconstructed $\nabla^2 G$ Pyramid, Using Local ZC and Local Estimation: 'Cake' . . . . .	52
4.5	Comparison: Best Possible Remerged Image, and Image Reconstructed Using ZC and Local Estimation: 'Cake' . . . . .	52
4.6	Reconstructed $\nabla^2 G$ Pyramid, Using Local ZC and Local Estimation: 'Girl' . . . . .	53
4.7	Comparison: Best Possible Remerged Image, and Image Reconstructed Using ZC and Local Estimation: 'Girl' . . . . .	53
5.1	Illustration: Inverse and Reconstruction Filtering Steps . . . . .	58
5.2	Block Diagram: Inverse and Reconstruction Filtering Steps . . . . .	58
5.3	1-D Space Domain Reconstruction Filter . . . . .	59
5.4	Imaginary Portion of the Frequency Response of a 1-D $\nabla^2 G$ Filter Step Response . . . . .	60
5.5	Frequency Plane Organization And Labelling . . . . .	61
5.6	2-D Reconstruction Filter (Frequency Domain), Specified in quadrants 1 and 4, and along the $\omega_1 = 0$ and $\omega_2 = 0$ axis . . . . .	62
5.7	2-D Reconstruction Filter Frequency Response, Imaginary Part, No Smoothing . . . . .	64
5.8	2-D Reconstruction Filter Frequency Response Magnitude, No Smoothing . . . . .	64
5.9	2-D Inverse Filter Frequency Response, Imaginary Part, No Smoothing . . . . .	66
5.10	2-D Inverse Filter Frequency Response Magnitude, No Smoothing . . . . .	66
5.11	Inverse Filtered Pyramid, No Smoothing: 'Girl' . . . . .	68
5.12	Reconstruction Filtered Pyramid, No Smoothing: 'Girl' . . . . .	68

5.13	Comparison: Best Possible Remerged Image, and the Image Recovered Using Reconstruction Filtering - No Smoothing: 'Girl' . . . . .	69
5.14	Comparison: Best Possible Remerged Image, and the Image Recovered Using Reconstruction Filtering, Using Only the Zero Crossing Samples - No Smoothing: 'Girl' . . . . .	69
5.15	2-D Reconstruction Filter Frequency Response, Imaginary Part, With Smoothing . . . . .	71
5.16	2-D Reconstruction Filter Frequency Response Magnitude, With Smoothing . . . . .	71
5.17	2-D Inverse Filter Frequency Response, Imaginary Part, With Smoothing	73
5.18	2-D Inverse Filter Frequency Response Magnitude, With Smoothing . . .	73
5.19	Inverse Filtered Pyramid, With Smoothing: 'Girl' . . . . .	75
5.20	Reconstruction Filtered Pyramid, With Smoothing: 'Girl' . . . . .	75
5.21	Reconstruction Filtered Pyramid, With Smoothing: 'Cake' . . . . .	76
5.22	Comparison: Best Possible Remerged Image, and the Image Recovered Using Reconstruction Filtering - With Smoothing: 'Girl' . . . . .	76
5.23	Comparison: Best Possible Remerged Image, and the Image Recovered Using Reconstruction Filtering, Using Only the Zero Crossing Samples - With Smoothing: 'Girl' . . . . .	78
5.24	Comparison of Images, Showing Progressive Loss By Processing Step: 'Girl' . . . . .	79
5.25	Comparison of Images, Showing Progressive Loss By Processing Step: 'Cake' . . . . .	80
6.1	Illustration of Least Squares Problem Situation . . . . .	86
6.2	Comparison of Reconstructed Image Quality, With Varying Edge Threshold and Edge Localization: 'Girl' . . . . .	90

# Acknowledgements

I would like to thank my supervisor, Dr. Peter Lawrence, for the generous enthusiasm and encouragement he provided throughout this research. His dedication is sincerely appreciated.

I would also like to thank the other members of my supervisory committee: Dr. Michael Beddoes, Dr. Rabab Ward, and Dr. Robert Woodham. The committee provided a valuable resource.

I am indebted to many colleagues for their help, particularly that of Dr. James Clark and Mr. Robert Ross.

The cooperation of Mr. Jack Wilson of Robotic Systems International Ltd., and Mr. Eric Jackson of International Submarine Engineering is gratefully acknowledged.

This research was made possible through an education leave of absence from IBM Canada Ltd., with financial support from a British Columbia Science Council G.R.E.A.T. Scholarship and the Natural Sciences and Engineering Research Council.

Lastly, I would like to thank my wife Elaine, for her patience and support.



# Chapter 1

## Introduction

Vision is probably man's most valuable sense. It is not surprising, therefore, that a better understanding of the vision process is a goal being sought by many different disciplines. Fields as diverse as psychology, physiology, computer science, robotics, and communication theory all consider the human vision and perception process. While each field has its own set of goals or objectives, there is much to be shared between these fields of research.

An important issue in understanding the vision process is the representation of images with partial information. By understanding how humans perceive images, insight might be gained into what the essential components of an image really are (for the human viewer). This issue is important for a basic understanding of the human vision system, but also can influence our approach to image processing tasks. Two such image processing tasks are computational vision for image interpretation [3,29] and image compression [71].

One objective of lower-level computational vision is to extract the essential components of the image. Sampled images contain large quantities of data, and in their raw form are difficult to interpret. Identifying the essential components of an image not only helps in the process of image interpretation, but also generally reduces the volume of data. A basic issue in such a low-level processing step is whether these essential components contain *all* the information in the original image. If these components do contain all the information

from the original image, further processing can continue without inherent loss. It should also be possible to reconstruct the original image from this partial information. If these extracted components do not contain all the original image information, these losses must be well understood. In this sense, the fields of computational vision and image compression share a common goal. In both fields, the identification of a compact set of data which will characterize an image is a basic objective.

## 1.1 Motivation for Study

This study will consider the representation and recovery of an image from a partial set of information using  $\nabla^2 G$  filtering. This study was motivated by an interest in image processing, and an image transmission problem in the subsea industry. In recent years, the use of unmanned submarines has become common for applications as diverse as oil drilling rig operations, recovering the space shuttle wreckage, and locating the Titanic. Often, these unmanned submarines are equipped with manipulator arms. An operator on the surface pilots the submarine and operates the manipulator arms. Without remote vision, it is very difficult to guide the submarine and essentially impossible to operate the manipulator arms.

Communicating through water, however, is difficult. Electromagnetic transmission can occur only at very low frequencies without high loss. Acoustic transmission provides the most effective line of sight bandwidth, but this bandwidth is still much less than that required for uncompressed image transmission. Even with current state of the art image compression techniques, good quality images cannot be transmitted through the available bandwidth. As a result, the submarines are commonly connected to the surface by a tether cable for communications and power (although power could be supplied with on board batteries for tetherless operation) [2].

The submarine and tether are heavily influenced by water currents, reactions to forces exerted by the manipulator, and the dynamics of a tether moving with the surface ship. Also, the tether presents many operational problems during launching and recovery, and

in tangling about underwater structures. The elimination of the tether would significantly reduce these operational difficulties. Remote underwater vision for tele-operation is not, however, possible today without a cable connection. Therefore, new image compression techniques are required to compress the images for the available acoustic channel bandwidth. This image transmission problem has been identified as one of the most serious facing the subsea industry [1].

## 1.2 The Underwater Transmission Problem

The bandwidth available through an underwater acoustical channel depends on the transmission range, and also the directional selectivity of the antennas. For transmission from a moving submarine, a channel capacity of 30 - 50 Kbits/second is technologically possible for ranges of several hundred meters [9,68]. Given this severe limit on the available bandwidth, it is necessary to consider where compromise can be taken in the imaging for tele-manipulation. Studies of imaging for tele-operator control of a remote manipulator have shown that a minimum image update rate of approximately 6 frames per second is required for effective tele-operation [68]. This study was based on a system that produced an image refreshed at 28 frames/second, implying that a given frame was presented several times before being updated. For frame update rates below 6 frames/second, operators began to use a “move-and-wait” strategy. This style of operation resulted in a considerable loss in performance and increase in variability. This study also found that an image resolution of approximately 256 X 256 pixels is required for ‘good’ underwater imaging. Recently, a ‘flight-simulator’ has been developed for studying tele-operation [39]. This study has found that a reduced frame rate is quite noticeable below 10 frames/second. Further, they have found that the impact of the reduced frame rate depends on the speed at which objects are moving within the frame, and the imaging precision required for the task (precise approach of the manipulator to an object, for example).

The propagation delay of sound in water must also be considered. The speed of sound in water is approximately 1600 meters/second. Since frame update rates below

6 frames/second lead to a “move-and-wait” control strategy [68], the same effect may exist for image transmission delays longer than 1/6 second. Assuming zero time for processing and presenting the image, this result implies that approximately 166 ms is available for acoustic transmission - a range of approximately 267 meters. This limitation may be partially addressed through the use of cable attached antennas, with acoustical transmission over a shorter range in the submarine’s localized area.

Uncompressed, a 256 X 256 monochrome image with 8 bits/pixel resolution and 30 frames/second results in a data rate of 15.7 Mbits/second. Assuming a frame rate of 6 frames/second, this data rate is reduced to 3.1 Mbits/second (with 8 bits/pixel resolution). The available bandwidth is of the order of 50 Kbits/second, so a compression is required from 8 bits/pixel to approximately 0.13 bits/pixel (approximately 63:1 compression).

Commercial acoustic image transmission systems are available today. One such system has been developed by Thompson-CSF in France [64]. This system communicates over a 20 Kbits/second acoustic link, and compresses the images using a simple block replacement scheme. Frames are updated at approximately one second intervals. Successive frames are examined for spatially correspondent blocks which have changed. If little has changed between successive frames, small changed blocks are transmitted. If the number of changes exceeds the channel capacity, larger image blocks are transmitted. For a sequence of images sweeping across a scene, this scheme results in a sequence of very low passed images being presented. Only when the remote scene is still does the full resolution scene reach the viewer. This scheme is therefore not desirable for imaging from a moving submarine.

### **1.3 Synopsis of Image Compression Techniques**

The field of sampled-image compression has been studied at length, and remains a very active area of research. I will not attempt to provide a comprehensive overview of past research in image compression; there are several fine review papers that provide such an overview [24,35,38,56,57]. Very briefly, image compression studies have generally

considered either the compression of a single image frame, or the compression of a sequence of frames (with no compression in the frame). The majority of this effort has been directed at the compression of a single image frame. A large number of approaches have been considered for this task, but the main effort has been on capitalizing on the geometrical and statistical characteristics of the image samples. Somewhat independent of technique, the best image compression for ‘good’ quality imaging has leveled off at approximately 10:1 for techniques based solely on information theory [36,38,45,53,97]. In recent years, the use of vector quantization has improved this figure to approximately 20:1 [19,89]. It should be noted, however, that this method is based on the use of a code book. As a result, the quality of image reconstruction is dependent on how well the current code book covers the current image. The stated compression rate does not cover the cost of retransmitting a new code book.

These information theory based techniques concentrate on the capture and recovery of the actual image samples. Given an image with a set of samples, there exists a measure of the information content or entropy of those samples, based on their statistical characteristics. This entropy measure provides a lower bound for the compression which can be obtained without loss in the image data. An information theory approach, however, does not consider the characteristics of the human receiver. The human receiver may not require the exact replication of the original image grey scale values to perceive the same image. Capitalizing on such a possibility, however, is difficult without a model for the human vision system (image acquisition, low-level processing, and perception). Unfortunately, while the understanding of the human vision system has progressed significantly, no comprehensive models of the human vision and perception process exist today. Several recent studies have considered the use of a human vision system model for image compression [16,38,59,66]. These studies have not, however, yet obtained the level of image compression required for the underwater image transmission problem.

A number of studies have considered image transmission at rates similar to those required for the underwater transmission problem [17,46,55,74]. These studies have considered the use of motion compensation, and in-the-frame statistical coding. For limited interframe movement and/or low resolution imaging, successful transmission has

been achieved at these rates. Transmission has not been accomplished, however, for grey scale imaging at the cited resolution and frame rate requirements. The use of cartoon-like region segmentation has also been used for low data rate imaging [63,94]. This approach works well for the transmission of distinct objects (sign language for the deaf, for example [41]), but has limited applicability for the transmission of images with complex shading, texture, and structures. For these types of images, the use of region segmentation with texture characterization has been examined [38]. While excellent compression rates have been reported with this approach (up to 70:1), the resulting images appear somewhat “paint-by-numbers”; the images are neatly divided into regions of uniform texture/intensity. As a result, these images reflect loss in shading, tone, and smoothness in the transition between segmented regions. It should also be noted that the image segmentation process is quite complex, and remains an active area of research.

## 1.4 Motivation for $\nabla^2 G$ Approach

This study will consider the representation and recovery of an image from a compact set of data. The application motivating this representation and recovery is the described image compression problem. Another motivating factor came from an interesting suggestion in the computational vision literature. A basic approach considered in computational vision studies is the evaluation and interpretation of images based on image primitives [21,49,50]. The image primitives most commonly considered are the ‘edges’ observed in an image. These edges are typically found by applying a second order differential operator, and examining the zero crossings.

Image intensity changes occur over a range of spatial frequencies. Abrupt changes in intensity correspond to high frequency edges; slowly varying intensity changes correspond to low frequency edges. Studies of the human vision system have found evidence of spatial filtering of retinal images into bandpass channels, which are processed in parallel [22,28,32,96]. These bandpass channels have been found to closely resemble the

characteristics of  $\nabla^2 G$  filters [49], which are given in the space domain by:

$$\nabla^2 G(\mathbf{r}) = \frac{-1}{\pi\sigma^4} \left( 1 - \frac{\mathbf{r}^2}{2\sigma^2} \right) e^{-\mathbf{r}^2/2\sigma^2} \quad (1.1)$$

where  $\mathbf{r} = \sqrt{x^2 + y^2}$ .  $\nabla^2 G$  filters are two-dimensional, Gaussian smoothed, second derivative filters. These filters have a bandpass characteristic, with a bandwidth of approximately one octave. A one-dimensional plot of this filter's frequency response can be seen in Figure 2.3. A two-dimensional plot of this frequency response can be seen in Figure 2.4.

Since the  $\nabla^2 G$  filter is a bandpass filter, it can be tuned (with  $\sigma$ ) to select a range of spatial frequencies. A family of  $\nabla^2 G$  filters can thereby be constructed to match the spatial-frequency channels observed in the human vision system. As mentioned earlier, changes in intensity can occur over a wide range of spatial frequencies. A change in intensity in an image (at a given spatial frequency) results in a spatial correspondent zero crossing in a  $\nabla^2 G$  filtered version of the image. This relationship can be used to localize the edges in an image within spatial-frequency channels, by locating the zero crossings in  $\nabla^2 G$  filtered versions of the image. Many computational vision studies have used the zero crossings found in a multiresolution family of  $\nabla^2 G$  filtered images as the basis for further image processing [8,21].

The  $\nabla^2 G$  operator has several important properties. Gaussian operators have been shown to provide minimum uncertainty for localization in space and frequency, for the set of all real operators [87]. Further, the Laplacian operator is the only linear, rotationally invariant second derivative operator. Also,  $\nabla^2 G$  filtered images do not ring on edge transitions. Therefore, while some debate exists about the relative performance of the  $\nabla^2 G$  filter as an edge detector (as compared to other edge detection techniques) [20,23], it was chosen because of these important properties and its relationship to the human vision system.

Marr and Hildreth have presented a comprehensive theory of edge detection based on  $\nabla^2 G$  filtering, and have argued that this method of edge detection is 'optimal' [49]. They have also considered whether a  $\nabla^2 G$  filtered image is uniquely characterized by its zero crossing locations. Marr et. al. [48] considered the applicability of Logan's theorem

[42] to early human vision processing. Logan has shown that a 1-D function strictly bandpassed to one octave can, under certain conditions, be completely characterized (to within a constant scalar) by its real zero crossings. Marr and Hildreth extended the application of Logan's theorem to  $\nabla^2 G$  filtered images, and suggested that there were "grounds for believing that this representation of the ( $\nabla^2 G$ ) image is complete". They pointed out that Logan's theorem was not strictly satisfied, but stated that Logan's theorem showed that the set of zero crossings extracted from  $\nabla^2 G$  filtered images was "rich in information". This discussion led to the consideration of the recovery of a  $\nabla^2 G$  filtered image from its zero crossings.

It was therefore decided to consider the representation of an image with partial information by separating it into a family of  $\nabla^2 G$  filtered images, and recovering each  $\nabla^2 G$  filtering plane from its zero crossings. Given the suggestion that the human viewer separates an image into a family of  $\nabla^2 G$  filtered images, and that the interpretation of the image is derived from these separated images, a motivation existed to consider a compression scheme based on this model. It was felt that 'good' perceived image quality might be possible even if the image were not reproduced exactly, so long as the image were reasonably reconstructed in the areas which were actually extracted through the early retinal processing. By separating the image into  $\nabla^2 G$  filtered channels and recovering each channel independently, it was felt that those aspects of the image most important to the human viewer might be preserved. Further, it was felt that the identification of a sparse set of image primitives might provide a new avenue for image compression.

This raises the issue of how to measure the goodness of a reconstructed image. Given an objective of achieving 'good' image reconstruction, one must consider how to measure the 'likeness' of two images. Unfortunately, no good objective measure of perceived image likeness exists today [66]. Traditionally, a simple mathematical measure (such as mean square error) has been used to measure the 'distance' in quality between two images. While it has long been recognized that such a measure is inappropriate, its use continues because of its simplicity. The exact replication of an image's grey scale values will obviously assure likeness. Such a replication may not be required, however, for perceived likeness. Further, studies on subjective image evaluation have found that



an inverse relationship can exist between a simple error measure and subjective quality [72]. In considering the possibility of an image approximation which capitalizes on the human vision system, it was felt that the use of a simple mathematical metric would be inappropriate (and possibly misleading). Instead, standard test images were used throughout this study. It was felt that this step should permit a subjective comparison of the reconstructed image quality with other published techniques, which is the real motivation for a measurement of quality.

A  $\nabla^2 G$  filtering approach was also motivated by two groups of studies in image compression. The first was the investigation of contour coding [18,82,83], which was considered in the late 1950's and 1960's and then abandoned. This approach was quite forward thinking, in that it recognized that the edges in an image were of fundamental importance to the human viewer. A low-pass version of the image was transmitted at a low data rate, and the sharp edges were synthesized from a two-dimensional code outlining the image contours. Fairly good image reconstruction was obtained, but the computational intensity of the task was prohibitive in its day. Further, the images appeared somewhat artificial, since all the edges in the image were synthesized at the same spatial scale (the concept of multiresolution edges was not present). With reduced computational costs and a multiresolution approach, it was felt that an edge based approach might yield good image reconstruction.

The second group of studies is more recent, and can be classed as sub-band coding. The basic key to this approach is that an image may be easier to code if it is first separated into a set of sub-images. The use of quadrature mirror filtering [97] and Laplacian ( $\nabla^2 G$ ) filtering [5,7,88], together with statistical coding of the sub-band images, have yielded image compression to the current state of the art. These findings encouraged the further consideration of a multiresolution approach to reconstruction from partial information.

## 1.5 Scope of Thesis

This thesis will consider the separation of an image into a family of  $\nabla^2 G$  filtered images, and the merging of this family of images to rebuild the original. Further, it will consider the characterization of a  $\nabla^2 G$  filtered image by its zero crossing locations. Two theorems in support of such a suggestion are considered. It is shown that  $\nabla^2 G$  filtered images cannot, in general, be so characterized.

While not all  $\nabla^2 G$  filtered images can be characterized by their zero crossing locations, the practical recovery from this partial information will be considered. Experimental results will be presented which show that the recovered images suffer significant loss.

From this experimental experience, it is proposed that the  $\nabla^2 G$  filtered images might be practically recovered from their zero crossing locations, with some additional characterization of the image in the vicinity of each zero crossing. A step edge model is used as the basis for such a characterization. This edge model is fitted to the edges found in the  $\nabla^2 G$  filtered image, through a local estimation procedure. The result of this estimation procedure is a scalar model fit estimate at each of the zero crossing sample points. A redrawing algorithm is then developed for recovering the  $\nabla^2 G$  filtered images from this partial information. The performance of this non-iterative approach is demonstrated on natural scene and synthetic images, and the advantages and disadvantages discussed.

A second approach is then developed, which overcomes the primary shortcoming of the local estimation method. This non-iterative approach, which is based on 2-D filtering, is shown to provide good image recovery (although some tone loss is apparent). Again, the performance of this approach is demonstrated on natural scene and synthetic images, and the advantages and disadvantages of this technique discussed.

The application of this recovery from partial information technique is then considered for image compression. A simple coding scheme is developed for representing the zero crossing contours with linear segments. A comparative study is then considered, examining the tradeoffs between compression tuning parameters and the resulting image

quality.

# Chapter 2

## Implementation of $\nabla^2 G$ Filters

In order to study the recovery of  $\nabla^2 G$  filtered images from partial information, it was necessary to implement a number of tools. This chapter will describe the system implemented for separating an image into a pyramid of low pass filtered images, and computing a  $\nabla^2 G$  pyramid of images. It will also consider a method for remerging the  $\nabla^2 G$  filtered images to recover an approximation of the original image, and for detecting the zero crossings in the  $\nabla^2 G$  filtered images.

### 2.1 Pyramid Filtering Technique

Images can present a wide range of spatial frequencies. Examination of an image can occur at varying degrees of resolution, depending on the accuracy required. If fine detail is needed, the full image resolution is suitable. If only crude outlines are required, the examination of the high resolution image may lead to a confusing level of detail. As a result, many studies have considered the use of multiresolution image processing [3,8,50]. The full resolution image is used to create a family of low pass filtered images. If these low pass images are properly generated, they can be subsampled without aliasing or loss of information at their lower resolution. These low pass images are spatially correspondent, and therefore suggest the structure of a pyramid.

As mentioned earlier, studies have found evidence that the human vision system separates the retinal image into spatially correspondent, frequency-tuned channels [22,28,32,96]. These studies have found evidence of the existence of several such discrete channels. In a similar manner, a family of  $\nabla^2 G$  filters can be used to compute this set of frequency-tuned images. The  $\nabla^2 G$  filters are bandpass, with a bandwidth of approximately one octave as defined by their -3 dB limits. A 1-D plot of this filter's frequency response can be seen in Figure 2.3. The one octave bandwidth suggests an arrangement where the center frequency of each bandpass channel is 1/2 that of its higher frequency neighbor. A system was therefore developed where the original image was used to generate a low pass pyramid of images, with the bandwidth of each successive plane 1/2 that of higher resolution neighbor. Each successive plane is subsampled by a factor of two in each dimension. A block diagram of this system can be seen in Figure 2.7. A system similar to this has been used in a number of other studies [8,20,49].

It was assumed that the images under consideration were properly bandlimited before they were sampled, to minimize aliasing. In this study, standard test images from the University of Southern California image data base were used [95]. Synthetic test images were also used. Based on the resolution requirements discussed in Chapter 1, original image sizes of 256 X 256 pixels were used.

## 2.2 Low Pass Filter Design

In order to generate the low pass pyramid from the original image, it was necessary to design a low pass filter to limit an image to 1/2 its original bandwidth. It is impossible to realize an ideal low pass filter, yet desirable to have as close as possible unity gain in the pass band and zero gain in the stop band. It was decided to implement this filter using a finite impulse response (FIR) filter design. This filter design technique provides a stable and easily realized 2-D filter [15].

A common approach to the design of 2-D filters is to design a 1-D filter, and transform it to 2-D. This was the approach taken for the design of the low pass filter. A 1-D low

pass filter was designed with the Remez exchange algorithm, as implemented in the IEEE 'Programs for Digital Signal Processing' (program 'EQFIR'). Briefly, this program designs an optimal Chebyshev linear phase approximation filter, for a given set of filter constraints and a given filter size. The filter constraints are specified by a pass band and a stop band, with an implied transition width. For a  $1/2$  bandwidth low pass filter, a narrow transition band and small pass/stop band ripple are desired. There exists, however, a complex relationship between transition width and pass/stop band ripple. A set of filter constraints suggested by Rabiner, McClellan and Parks [65] were therefore chosen, for a good compromise between filter ripple and transition width. These filters were implemented with FFTs rather than convolutions, because they could be computed more quickly on the array processor available for this study. As a result, a large filter size was chosen to minimize the error introduced by truncating the number of filter coefficients. The designed filter has a maximum ripple in the pass band of 0.074 dB, and the filter is down a minimum of -41.4 dB in the stop band. A plot of this filter can be seen in Figure 2.1. Further filter design details are included in the Appendix.

The 1-D low pass design was then transformed into a 2-D design. A common method used for such a transformation is the McClellan transform [51,52,54], which maps a 1-D filter onto a 2-D plane. This approach preserves the circular symmetry of the 2-D filter near the center of the plane, but introduces progressive distortion towards the edges of the plane. A more direct method of transforming a 1-D filter to 2-D is to rotate the 1-D filter in the frequency domain. This approach preserves the circular symmetry well, but has not been used in the past because of the computational load of computing the 2-D inverse Fourier transform [31,37,98]. This one-time computational load was not considered significant, and therefore the frequency domain rotation method was used.

The rotation of the 1-D filter onto the 2-D plane was done as follows. The 1-D low pass filter impulse response, as generated by the Remez exchange algorithm, was loaded into a 1024 complex string (imaginary part equal to zero). The FFT of this string was computed, so that a closely spaced set of estimates of the frequency response of the filter was available. The 2-D filter frequency response was constructed on a 256 X 256 complex plane. For each point on the plane, the radius was approximated to its closest




Figure 2.1: 1/2 Bandwidth 1-D Low Pass Filter




Figure 2.2: 1/2 Bandwidth 2-D Low Pass Filter. Note: Center of plot corresponds to  $\omega_1 = \omega_2 = 0$

point in the 1024 point estimate. The 1-D response sweeps a circle out on the 2-D plane. Those points outside the circle correspond to regions in the filter stop band, so zero gain was assumed in these regions. With the desired 2-D frequency response constructed, the inverse 2-D FFT was computed. Because of the symmetry of the 2-D filter, the space domain coefficients were all real. The space domain coefficients were then truncated to the same number as in the 1-D design step. A plot of the 2-D low pass filter can be seen in Figure 2.2. All the 2-D filter plots have been arranged so that the center of the plot represents the point  $\omega_1 = \omega_2 = 0$ .

With the low pass filter applied to an image, the image can be subsampled by a factor of two in each dimension. By successively applying the low pass filter and subsampling, a low pass image pyramid can be constructed. The same filter coefficients can be used at each level in the pyramid, since those coefficients effect a 1/2 bandwidth filter at that scale. The scaling of the coefficients was adjusted for each plane size, to compensate for zero padding. At each level of the pyramid, the low pass filter has the same Chebyshev pass band ripple characteristic. As a result, successive application of the filter to generate the low pass image pyramid results in a larger propagated pass band ‘ripple’. This propagated effect was examined, and found to result in only a -0.35 dB maximum variation in the pass band.

## 2.3 $\nabla^2 G$ Filter Design

The  $\nabla^2 G$  filter was designed directly in the space domain from its closed form:

$$\nabla^2 G(r) = \frac{-1}{\pi \sigma^4} \left( 1 - \frac{r^2}{2\sigma^2} \right) e^{-r^2/2\sigma^2} \quad (2.1)$$

where  $r = \sqrt{x^2 + y^2}$ . The sigma value chosen was based on that suggested by Marr and Hildreth [49];  $\sigma = \sqrt{2}$ . Again, the filter was implemented with FFTs, so a large filter size could be used without penalty. Since the  $\nabla^2 G$  filter is bandpass, the sum of its space domain coefficients (or impulse response) should equal zero. The number of coefficients was increased until the sum of the coefficients equaled zero (to within the numerical



resolution of the system used). A filter size of 21 X 21 was used. A plot of the 1-D  $\nabla^2 G$  filter can be seen in Figure 2.3, and a plot of the 2-D filter can be seen in Figure 2.4.

The 2-D  $\nabla^2 G$  filter was applied to each level of the low pass filtered image pyramid. Since each successive plane was subsampled in the low pass step, the same  $\nabla^2 G$  filter coefficients could be used for each plane. A family of  $\nabla^2 G$  filtered images were thereby generated. Each  $\nabla^2 G$  channel had the same one octave characteristic, with a band pass center frequency twice that of its lower frequency neighbor. A 1-D plot of this family of  $\nabla^2 G$  filters can be seen in Figure 2.5.

## 2.4 Recovering an Image From a $\nabla^2 G$ Pyramid

The recovery from partial information strategy discussed in the introduction focused on the recovery of a  $\nabla^2 G$  filtered image plane. The original image is separated into 4 such  $\nabla^2 G$  filtered images, with 4 separate recovery problems. This section will discuss the remerging of a pyramid of  $\nabla^2 G$  filtered images, to recover an approximation of the original image.

The 'optimal' remerging of the  $\nabla^2 G$  channels to approximate the original image must consider the characteristics of the human viewer. Unfortunately, no comprehensive model of the human vision and perception process exists today. Since the development of such a model is a massive task outside our field of study, it was decided not to consider it in this study. In the absence of an adequate model, it was decided to remerge the images with a goal of recovering a frequency spectrum like that of the original image. Restated, a choice was made to remerge the filter channels so that the merged system transfer characteristic had a frequency spectrum as flat as possible (with unity gain).

Since the  $\nabla^2 G$  filtered images are bandpass, each channel represents a portion of the image's spectrum. If these channels were non-overlapping, the images could simply be expanded to the same size and summed to approximate the original. A  $\nabla^2 G$  channel has non-zero gain for all frequencies other than  $f = 0$ , however, and therefore the bandpass




Figure 2.3:  $\nabla^2 G$  1-D Filter




Figure 2.4:  $\nabla^2 G$  2-D Filter. Note: Center of plot corresponds to  $\omega_1 = \omega_2 = 0$




Figure 2.5:  $\nabla^2 G$  Family of Filters

channels overlap. Further, the near-dc portion of the image's frequency spectrum is not contained in any of the bandpass channels. Also, the highest frequency  $\nabla^2 G$  channel has effectively zero gain for the image's maximum spatial frequencies ( $f = f_{s/2}$ ).

A scheme was therefore developed to weight each of the  $\nabla^2 G$  channels, and sum the expanded images to approximate the original. The use of pre-filtering or post-filtering to boost the spectrum near  $f = f_{s/2}$  was considered, but not implemented (primarily due to the author's time constraints). Unity gain was therefore chosen for the highest frequency  $\nabla^2 G$  channel, in order to include as much of the image's high frequency components as possible. A search algorithm was then used to adjust the gains on the lower 3  $\nabla^2 G$  channels, to minimize the mean squared error in the summed spectrum. The channel weights determined through this process are included in the Appendix.

The summed  $\nabla^2 G$  channels did not include the image's near-dc spectrum, and it was found experimentally that the absence of this spectral portion had a high visual impact. Therefore, a matched low pass filter was designed to capture this spectral portion. The near-dc filtered image was sampled on a 16 X 16 grid. A 1-D plot of the summed  $\nabla^2 G$




Figure 2.6: 1-D Sum of  $\nabla^2 G$  and Near-dc Channels

filter channels, with the near-dc low pass channel, can be seen in Figure 2.6.

As each image plane was expanded for summing with its higher frequency neighbor, the image was filtered with a reconstruction low pass filter. The 1/2 bandwidth low pass filter (discussed earlier) was used for this purpose, and served to smooth the block-like characteristic introduced when expanding an image with pixel replication. A diagram illustrating the filtering system is shown in Figure 2.7.

One natural scene image (girl) and one synthetic image (cake) were used extensively in this study. The 'girl' low pass pyramid is shown in Figure 2.8. In this photograph, the upper left shows the original image (256 X 256). The low pass planes generated from this image are shown in the upper right (128 X 128), the lower left (64 X 64), and lower right (32 X 32). The subsampled images were expanded to the same size for presentation, using pixel replication. This pyramid presentation format was used throughout the study. The 'cake' low pass pyramid is shown in Figure 2.9. The  $\nabla^2 G$  pyramids of these images is shown in Figures 2.10 and 2.11. Since the entire spectrum of the original image was not contained in the summed channels, some loss in the remerged image's quality is to

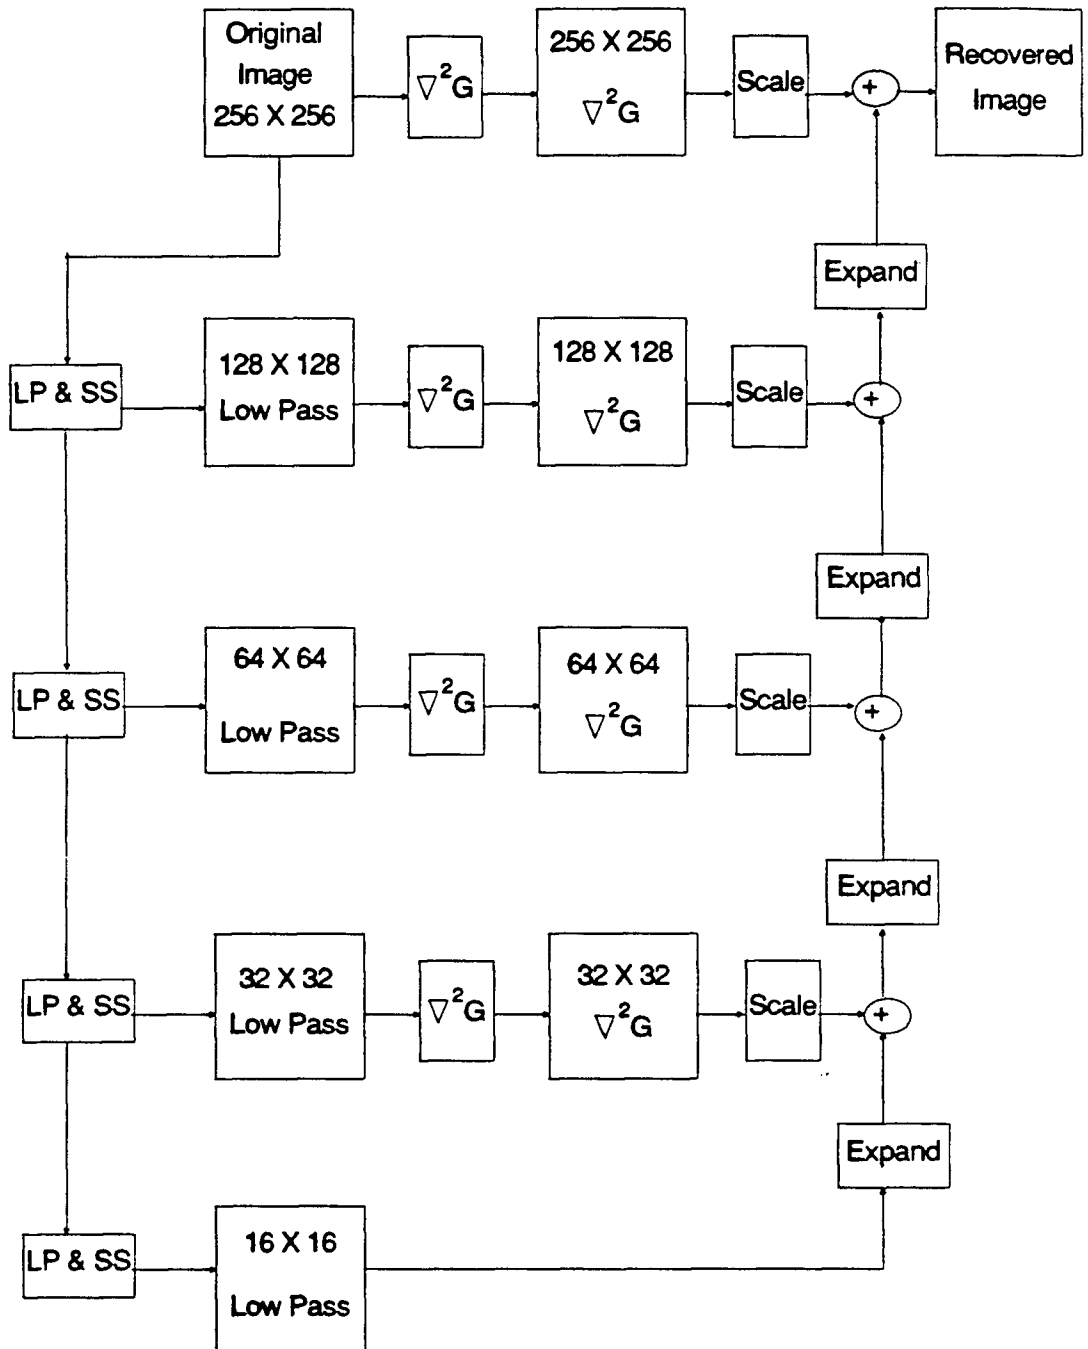


Figure 2.7:  $\nabla^2 G$  Image Generation and Image Recovery

be expected. The impact of this loss can be seen in Figure 2.12, where the original and remerged ‘girl’ images are shown.

## 2.5 Zero Crossing Detection

A procedure was implemented for finding the zero crossings in the  $\nabla^2 G$  filtered plane. Zero crossings of 2-D analog functions form continuous contours. There are, therefore, an infinite number of zero crossings in a 2-D zero crossing segment. For sampled images, one must decide how to represent the 2-D zero crossings. Almost all the zero crossings in a  $\nabla^2 G$  filtered image occur between adjacent image samples of opposite sign. A choice was made to operate on zero crossings localized to an adjacent sample point on the image sample grid. No attempt was made to localize the zero crossings with subpixel accuracy [33], with the choice rather to place a zero crossing point at the nearest sample to the right or below the detected zero crossing point. A zero crossing was placed wherever a change in sign was detected between adjacent pixels, or where a zero crossing through a sample of value zero was detected. As an example, the zero crossings detected in the ‘girl’  $\nabla^2 G$  pyramid are shown in Figure 2.13.



Figure 2.8: Low Pass Filtered Image Pyramid: 'Girl'




Figure 2.9: Low Pass Filtered Image Pyramid: 'Cake'




Figure 2.10:  $\nabla^2 G$  Filtered Image Pyramid: 'Girl'




Figure 2.11:  $\nabla^2 G$  Filtered Image Pyramid: 'Cake'





Figure 2.12: Comparison: Original Image, and Best Possible Remerged Image: 'Girl'




Figure 2.13: Zero Crossing Image Pyramid: 'Girl'

## Chapter 3

# Characterizing $\nabla^2 G$ Filtered Images by their Zero Crossings

The relationship between a function and its zero crossings has been studied extensively [4,47,69,70,91,92]. A fundamental issue in these studies has been whether a function can be uniquely characterized by its real zero crossings. This problem has two distinct sub-problems; assuring unique characterization, and developing a (practical) recovery method. Further, each of these sub-problems needs to be considered for continuous and sampled functions.

In general, it is not possible to uniquely characterize an arbitrary function by its real zero crossings. Under certain conditions, however, such a characterization is possible. It has been suggested that  $\nabla^2 G$  filtered images might be characterized by their zero crossing locations. In this chapter, two theorems in support of such a suggestion will be considered: Logan's and Curtis'. A sequence of counter examples will also be presented, which will show that not all  $\nabla^2 G$  filtered images, filtered in 1-D or 2-D, can be uniquely characterized by their zero crossing locations. Recovery problems in a sampled data environment will then be considered. Lastly, the practical recovery of a  $\nabla^2 G$  filtered image from its sampled zero crossing locations will be examined.

### 3.1 Logan's Theorem

Logan has shown that the real zero crossings of a 1-D function  $h$  determine  $h$  within a multiplicative constant, if  $h$  has no zeros in common with its Hilbert transform  $\hat{h}$  other than real simple zeros, and  $h$  is strictly bandlimited to a bandwidth of less than one octave [42]. This theorem establishes a sufficient, although not necessary, condition for the characterization of a function by its zero crossing locations.

Marr et al [48] considered the applicability of Logan's theorem to early visual processing, and concluded that Logan's theorem was relevant. The model they used for the early vision system was based on oriented (sharp cut-off) bar-shaped masks. Since Logan's theorem has been shown only in 1-D, they considered the application of these bandpass filters to 1-D scan lines in the image. They stated that the critical aspect of Logan's theorem was the one octave bandwidth requirement. They also stated that the bandwidth of the spatial-frequency-tuned channels in the human vision system were slightly larger than the one octave limit, and that Logan's theorem would apply directly if this one octave limit were met. Further, their stated important point was that "Logan's theorem shows that zero crossings of a one-octave bandpass signal contain complete information and the evidence is that they remain rich in information even when the one-octave condition is relaxed". This statement was linked with a conjecture that Logan's theorem would almost always hold true for bandwidths up to 1.67 octaves ("Rice's formula shows that Logan's conclusion holds up to 1.67 octaves").

Marr and Hildreth [49] extended the application of Logan's theorem to  $\nabla^2 G$  filtered images, based on the findings of Marr et al [48]. Marr and Hildreth stated that the zero crossing locations of a  $\nabla^2 G$  filtered image were "rich in information", and that there were "grounds for believing" these zero crossing locations formed a complete representation. They pointed out that the  $\nabla^2 G$  filter had a half-sensitivity bandwidth of about 1.75 octaves, and that therefore Logan's theorem did not strictly apply. They further stated that if the  $\nabla^2 G$  filter had a half-sensitivity bandwidth of one octave, the  $\nabla^2 G$  filtered images would be completely characterized by their zero crossings. This statement is not supported

by arguments based on Logan's theorem, since Logan's theorem requires strict, one octave bandpassing (i.e. spectra confined to  $[-\beta, -\alpha]$  and  $[\alpha, \beta]$  where  $0 < \alpha < \beta < \infty$  and  $\beta = 2\alpha$ ).

Consider the differences between a function strictly bandpassed to one octave, and a  $\nabla^2 G$  filtered function.  $\nabla^2 G$  filtered functions are bandpassed to approximately one octave as defined by their -3 dB limits, but they are not strictly bandpassed. A 1-D  $\nabla^2 G$  filter is given in the space domain by:

$$\nabla^2 G(x) = \frac{-1}{\sigma^3 \sqrt{2\pi}} \left(1 - \frac{x^2}{\sigma^2}\right) e^{-x^2/2\sigma^2} \quad (3.1)$$

with Fourier transform:

$$\mathcal{F}\{\nabla^2 G(x)\} = \omega^2 e^{-\omega^2/2\sigma^2} \quad (3.2)$$

The filter can be seen to be nonzero for all frequencies except  $\omega = 0$ , and is therefore not strictly bandlimited. This implies that  $\nabla^2 G$  filtered functions will not be assured to be strictly bandpassed.

An ideal, one octave, strictly bandpassed 1-D filter is given by:

$$h(x) = \frac{\sin(\alpha x)}{\alpha x} \cdot \cos(\beta x) \quad (3.3)$$

where  $2\alpha$  is the bandpass width,  $\beta$  is the bandpass center frequency, and  $\beta = 3\alpha$ . Note that the space domain filter contains an infinite number of zero crossings.

The application of these two filters to many natural scene images will result in significantly different sets of zero crossings. Many natural scene image intensity changes occur as steps or smoothed steps. A step input therefore represents a possible and realistic condition. Figure 3.1 shows the step response of an ideal, one octave, strictly bandpassed filter. This response contains an infinite number of zero crossings. Figure 3.2 shows the step response of a  $\nabla^2 G$  filter. This response contains a single zero crossing. It is clear from these two figures that the number and location of zero crossings which result from these two filters differ significantly.  $\nabla^2 G$  filtered images therefore violate Logan's theorem not just in the slight breach of the one octave limit, but more seriously in the fact that the functions are not strictly bandpassed. This violation results in a dramatic

difference in the number and location of the zero crossings (1 .vs.  $\infty$ ). Logan's theorem is therefore not applicable to  $\nabla^2 G$  filtered images.

### 3.2 Curtis' Theorem

More recently, Curtis has shown that under certain conditions, a bandlimited 2-D function can be uniquely characterized by its zero crossing locations [10,11,12,13,14]. These findings are based on the representation of a 2-D function with a polynomial using a Fourier series, and on Bezout's theorem. If the 2-D function is strictly bandlimited, its Fourier series representation will be finite in length. The 2-D function approximated with a polynomial can then be expressed in the form:

$$f(x, y) = \sum_{n_1} \sum_{n_2} \mathcal{F}(n_1, n_2) \omega_1^{n_1} \omega_2^{n_2} \quad (3.4)$$

where:

$f(x, y)$  is the polynomial approximation of the 2-D function

$\mathcal{F}(n_1, n_2)$  are the Fourier series coefficients

$$\omega_1 = e^{j2\pi x/\tau_1}$$

$$\omega_2 = e^{j2\pi y/\tau_2}$$

$\tau_1$  is the period in the  $x$  direction

$\tau_2$  is the period in the  $y$  direction

Since the Fourier coefficients are nonzero for positive and negative values of  $n_1$  and  $n_2$ ,  $f(x, y)$  is not a polynomial in  $\omega_1$  and  $\omega_2$  (strictly speaking). Since  $f(x, y)$  is bandlimited, however, its Fourier coefficients equal zero outside the region  $-N_1 \leq n_1 \leq N_1, -N_2 \leq n_2 \leq N_2$ . The polynomial can therefore be written in a proper form through the following modification:

$$\hat{f}(x, y) = \omega_1^{N_1} \omega_2^{N_2} f(x, y) \quad (3.5)$$

$$= \sum_{n_1=0}^{2N_1} \sum_{n_2=0}^{2N_2} \mathcal{F}(n_1 - N_1, n_2 - N_2) \omega_1^{n_1} \omega_2^{n_2} \quad (3.6)$$




Figure 3.1: Ideal Bandpass Filter Step Response




Figure 3.2:  $\nabla^2 G$  Filter Step Response

where  $\hat{f}(x, y)$  is the modified polynomial expression.

Based on the preceding polynomial representation form, Curtis' has shown that:

*Let  $f(x, y)$  and  $g(x, y)$  be real, two-dimensional, doubly-periodic, bandlimited functions with  $\text{sign } f(x, y) = \text{sign } g(x, y)$ , where  $f(x, y)$  takes on both positive and negative values. If  $f(x, y)$  and  $g(x, y)$  are nonfactorable when expressed as polynomials in the Fourier series representation, then  $f(x, y) = cg(x, y)$ .*

Stated simply, this result shows that the zero crossings of a 2-D function uniquely determine the function to within a scalar if that function is strictly bandlimited and nonfactorable when written in the above polynomial form. Again, this theorem establishes a sufficient, although not necessary, condition for the characterization of a function by its zero crossing locations.

Let us consider each of the conditions of this result, to see if  $\nabla^2 G$  filtered images satisfy them. In considering Logan's theorem, it was shown that  $\nabla^2 G$  filtered functions are not assured to be bandlimited. If an input image of infinite bandwidth were  $\nabla^2 G$  filtered, the result would not be strictly bandlimited. Such an image would violate Curtis' requirement for the function to be strictly bandlimited. In practice, images are usually bandlimited by the systems which manipulate them (e.g. an imaging system's optical bandwidth, or an anti-aliasing filter prior to sampling). Most such images are finite in size, however, and are therefore not strictly bandlimited. For functions of this type, Curtis considers the function as one period (or part of a period) of a periodic function. Under certain conditions, such a periodic function may be bandlimited.

Curtis argues that, for all 2-D functions, "the irreducibility constraint is satisfied with probability one, since it has been shown that the set of reducible  $m$ -dimensional polynomials forms a set of measure zero in the set of all  $m$ -dimensional polynomials (for  $m > 1$ )" [13]. This argument is based on a finding of Hayes and McClellan [27], who showed that "almost all" multidimensional polynomials are nonfactorable. It should

be noted that a set of measure zero does not imply that the set is empty. For example, the set of all rational numbers forms a set of measure zero in the set of all reals, even though there are an infinite number of rational numbers.

As a result, there is no assurance that a  $\nabla^2 G$  filtered image will be nonfactorable. As a simple example, consider a 2-D function created by rotating a 1-D function about its origin. If the 1-D function is not identically zero, the rotated function, written as  $f(x, y)$ , is clearly a 2-D function. This function can be factored, however, through a transformation into polar coordinates (if the 1-D function has a degree greater than 1). One might argue that the 2-D function is in some sense one dimensional, but the example represents an entirely possible 2-D image. It is quite possible that a camera image could be rotationally symmetric. Further, there exist other examples of 2-D functions which can be expressed in 1-D through a coordinate transformation. Therefore, not all  $\nabla^2 G$  filtered images are nonfactorable. This point is important because it shows that possible  $\nabla^2 G$  filtered images do not satisfy the non-factorability condition of Curtis' theorem.

It can be argued that if a factorable 2-D image did occur, it could be perturbed slightly to create a non-factorable 2-D polynomial. On this point, it should be noted that the determination of the factorability of a 2-D polynomial is nontrivial (if not impossible). In practice, it would be important to determine if a given  $\nabla^2 G$  filtered image satisfied the constraints of Curtis' theorem, since meeting these conditions is necessary to assure uniqueness of representation on the basis of Curtis' theorem. If a  $\nabla^2 G$  filtered image were found which did not satisfy the factorability constraint, its perturbed form would also have to be examined to assure non-factorability. Further, this argument is somewhat more problematic when considering sampled images. Once sampled, a finite number of quantized states are represented in the image. A perturbation could occur only in quantized amplitude steps at each sample point. It is not clear what the probability of obtaining a nonfactorable 2-D polynomial would be in this situation.

In addition to uniqueness of representation, one must consider the stability of the representation for practical recovery. A small error in localizing the zero crossings must result in only a small error in the polynomial represented by those zero crossings. If



an arbitrarily small error in determining the zero crossings leads to a large error in the the represented polynomial, the usefulness of this representation is limited in a practical application. Hummel [34] has recently considered the stability of the representation of a polynomial by its zero crossings, and found that “stability of reconstruction is unlikely”. Restated, he has found that “there can be widely different initial data leading to nearly identical zero-crossing data”. This instability limits the usefulness of the representation for practical recovery.

In summary, not all  $\nabla^2 G$  filtered images meet the constraints of Curtis’ theorem. As a result, Curtis’ theorem does not assure the unique characterization of all  $\nabla^2 G$  filtered images by their zero crossings. Further, it is difficult to establish if an image meets the constraints of Curtis’ theorem.

### 3.3 Scale Space Considerations

Recently, the zero crossings found in  $\nabla^2 G$  filtered images, filtered at a continuum of scales, have been considered for characterizing an image. Yuille and Poggio [99,100] have argued that these zero crossings uniquely characterize almost all images. Hummel [34] has argued, however, that this representation is not stable for reconstruction. He states that small errors in measurement of the zero-crossings could lead to arbitrarily large errors in the determination of the image. Hummel also provides an example of two different 2-D functions which have exactly the same zero crossings at all levels of resolution. This example is not a circularly symmetric function.

Further, the zero crossings of a scale space image are not of great benefit for an image compression application. The scale space representation of a 2-D image is a 3-D map. The volume of zero crossing data in a scale space map is an order larger than the data in the original image. It would therefore be very difficult to obtain a compression rate approaching anything close to that currently achievable, using this approach.

### 3.4 Different $\nabla^2 G$ Filtered Images with the Same ZC Locations

Both Logan's and Curtis' theorems establish sufficient, although not necessary, conditions for the unique characterization of a function by its zero crossing locations. It has been shown that Logan's theorem is not applicable to  $\nabla^2 G$  filtered images, and that not all  $\nabla^2 G$  filtered images meet the constraints of Curtis' theorem. A logical question, then, is if any  $\nabla^2 G$  filtered images exist which are not uniquely characterized by their zero crossing locations. In this section, a sequence of different, continuous, bandlimited  $\nabla^2 G$  filtered images will be presented which show, by counter-example, that not all  $\nabla^2 G$  filtered images can be characterized by their zero crossing locations. Although 1-D examples will be considered here, these examples can be extended to 2-D by rotating the 1-D functions about their origins. Figure 3.3 shows a strictly bandlimited image intensity profile (with 2 'edges'), and its  $\nabla^2 G$  filtered result (with 2 zero crossings). The example in Figure 3.4 demonstrates the problem of 'extra' or 'phantom' zero crossings. Successive image intensity changes with the same second derivative sign result in an extra zero crossing in the  $\nabla^2 G$  filtered image. This extra zero crossing is not associated with an intensity change in the original image. If the extra zero crossing is ignored, it is impossible to differentiate between these two  $\nabla^2 G$  filtered images on the basis of zero crossing locations alone.

The example shown in Figure 3.5 demonstrates that the simple consideration of the extra zero crossing does not resolve the ambiguity. This figure shows a third image intensity profile, with its  $\nabla^2 G$  filtered result. If the extra zero crossing in Figure 3.4 is considered, then these two  $\nabla^2 G$  filtered images have the same sets of zero crossing locations. Thus, the knowledge of zero crossing locations alone is inadequate to characterize all  $\nabla^2 G$  filtered functions.




Figure 3.3: Two Edge Image Intensity Profile, and its  $\nabla^2 G$  Filtered Response




Figure 3.4: Staircase Image Intensity Profile, and its  $\nabla^2 G$  Filtered Response




Figure 3.5: Image Intensity Profile, and its  $\nabla^2 G$  Filtered Response

### 3.5 Problems in a Sampled Environment

The preceding arguments are based on continuous or analog images. Curtis has shown that her arguments extend to sampled images as well. Curtis et al [10] have considered the characterization of an image using the zero crossings found in the real portion of the image's Fourier transform. They have found that such a characterization is not preserved unless very accurate estimates of the zero crossing positions are used. For a strictly bandlimited square image with a spectral limit of  $N$  in both dimensions, they have found that not more than  $16N^2$  zero crossing contour samples are required to uniquely characterize the 2-D function. Subsequently, Zakhor and Izraelevitz [102] have argued that a sampling requirement of at most  $8N^2$  zero crossing samples are required. From these estimates, it is clear that the application of this theory for image compression will be difficult.

Curtis et al [10] considered the spatial sampling rate requirements for representing a 2-D function with its sampled zero crossings. Consider also the impact of quantizing

each image sample. As discussed earlier, a  $\nabla^2 G$  filter is not strictly bandpassed (or bandlimited). The response of this filter to a step input is infinite in extent, and contains only one zero crossing. As the function proceeds away from the zero crossing, its value approaches zero asymptotically. This can be seen in Figure 3.2. Consider the sampling of this function, with arbitrarily closely spaced samples. If this function is sampled (to any finite number of bits of resolution), samples approximated to zero will be taken. Consider, then, a nonfactorable image which contains two lines of intensity change, separated by a wide region of uniform intensity (such as in Figure 3.3). At each line of intensity change, a positive and negative variation will occur in the  $\nabla^2 G$  filtered image. In the wide region of uniform intensity, no such variations will occur. If this region of separation is made arbitrarily wide, samples of value zero will be taken in this region. The two lines of intensity change (edges) thereby become independent, since the width of the value zero region would not affect either edge. As a result, the recovery of the single image plane could not occur to within a single constant scalar, since each of the three independent regions would require its own scaling factor.

Experimentally, a significant number of zero value (non-zero crossing) samples were found to be resolved from  $\nabla^2 G$  filtered natural scene images (with 8 bits/pixel of resolution). Since the grey scale resolution of the human viewer is less than 8 bits/pixel, it seems unlikely that more accurate estimates are required. On close examination of these sampled  $\nabla^2 G$  filtered images, regions of independent variations as described above were found.

### 3.6 Practical Recovery from ZC Locations

In spite of these findings, the cited image compression application motivated a further investigation into the practical recovery of  $\nabla^2 G$  filtered images from their zero crossing locations. A number of studies have considered a similar goal. Logan has considered the practical recovery of one dimensional signals from their zero crossing locations [43,44]. In this study, signals were constrained to a certain class designed for recovery from their

zero crossing locations. Rotem and Zeevi [73] have considered the recovery of images from zero crossing locations, based on one dimensional scan lines in the image. Their study used (approximated) strictly bandpass one octave filters, zero crossing locations quantized to the image sample grid, and an iterative algorithm for recovering the image scan lines. They found that convergence was possible in almost all cases where the bandwidth was less than one octave, and that the images did not converge for a bandwidth greater than one octave.

Consider the number of zero crossing points found through 1-D and 2-D  $\nabla^2 G$  filtering of an image. As discussed in Chapter 2, the 2-D zero crossing detection scheme used in this study localized discrete zero crossing points. One can therefore compare the number of zero crossing points found in one row of a 2-D processed plane, with the number of zero crossing points found in a 1-D filtered scan line taken from a 2-D image. In the 1-D case, each line must be bandpassed. In the 2-D case, the image is filtered with a circularly symmetric bandpass filter. This implies that each scan line of the 2-D filtered image is assured to be low passed, but not bandpassed [48]. A function which is low passed only may contain very few zero crossings, if the function has a dc-bias which elevates the function variations well above the zero line. In contrast, a bandpassed function must have a zero value mean, and will therefore tend to have zero crossings associated with most function value variations. The 1-D scan lines will therefore tend to have more zero crossing points than scan lines taken from the 2-D filtered images. As a result, more zero crossing points will, in general, be found with the 1-D approach. The determination of how many more zero crossings would occur is dependent on the content of a specific image.

A choice was therefore made to consider a 2-D approach to recovery from zero crossings, since the fewer number of zero crossing points would be better for compression. Recently, a 2-D approach has been considered in several studies. As mentioned, Curtis [13] considered the recovery of images from closely spaced zero crossing location estimates found in the 2-D Fourier transform of an image. The recovery methods considered were a linear equations method, and an iterative scheme based on a generalized Gerchberg-Saxton algorithm.

The linear equations method considered by Curtis is very computationally intensive. Curtis et al [10] have reported the recovery of a 16 X 10 pixel image through the solution of 178 equations in 159 unknowns with the zero crossing positions determined to 5 decimal places. For full size images (256 X 256, for example), this approach would lead to a prohibitive computational problem.

The iterative approach considered by Curtis involves the successive application of known problem constraints in two domains. In her study, this involved the application of constraints in the space and frequency domains. A small space domain image is centered on a large, zeroed plane. The Fourier transform of the large plane is computed, providing closely spaced estimates of the image's Fourier transform. The zero crossings of the real portion of the Fourier transform are saved. The iterative procedure then involves successively imposing the correct sign in the frequency domain, and imposing the region of support in the space domain. Curtis has reported convergence for a 64 X 64 image placed on a 256 X 256 plane in 25 iterations. It should be noted that the correct Fourier transform phase was used as the starting point for the iteration cycle.

Sanz [80] has also considered the characterization of a 2-D image by its sampled zero crossings. He has found that sampling the zero crossing contours results in a significant loss of information, and that an infinite number of signals which have a prescribed bandwidth may coincide over a continuum of samples. Very recently, Sanz and Huang [78] have considered the recovery of 2-D  $\nabla^2 G$  filtered images from their sampled zero crossings. In this study, the zero crossings found in 4  $\nabla^2 G$  filtered channels were used to recover the original image. The  $\nabla^2 G$  channels were merged by summing them directly, and applying an "inverse filter" to flatten the image spectrum. An iterative scheme was developed which used all 4  $\nabla^2 G$  channels at once. To start, each  $\nabla^2 G$  channel image was clipped to its correct sign. The 4  $\nabla^2 G$  channels were then summed and inverse filtered. The result was filtered back into 4  $\nabla^2 G$  channels. Each  $\nabla^2 G$  image was corrected where it presented an incorrect sign, and the process was repeated. The reconstructed images were found to retain pictorial features, but appeared as if "an impressionist had painted them".

A choice was made here to consider the recovery of each of the  $\nabla^2 G$  filtered images independently, using the remerging and zero crossing detection scheme described in Chapter 2. Two iterative methods were considered. Based on the success of Rotem and Zeevi [73], an extension of their 1-D iterative method was made to 2-D. Their study was based on an algorithm developed by Voelcker and Requicha [93], and was extended to this application in the following form. Given an original image  $S$ , iterative estimate  $S_{n+1}$  is given by:

$$S_{n+1} = S_n - c[B_p Sgn(S_n) - S_0] \quad (3.7)$$

where:

$S_{n+1}$  is the current iteration result

$S_n$  is the last iteration result

$c$  is a constant

$B_p$  is the bandpass operation

$Sgn$  is the clipping operation

and the constant term  $S_0$  is given by

$$S_0 = B_p Sgn(S) \quad (3.8)$$

This algorithm was extended to 2-D by operating on the whole 2-D image with each iteration, with the  $\nabla^2 G$  filter used as the bandpass operator. It was found that the iterative sequence converged for a number of iterations, and then diverged. This convergence characteristic was fairly consistent, somewhat independent of the constant  $c$ . The iteration was continued until divergence was detected, with the last iterative result taken as the best recovered image. The recovered images retained the basic pictorial features, but suffered significant loss in shading and tone.

The second iterative approach examined was based on the application of constraints in two domains. For each  $\nabla^2 G$  filtered image, the zero crossing locations are known in the space domain, and the image's frequency characteristics have been shaped by the  $\nabla^2 G$  filter. An iterative scheme was implemented which started with a clipped version of the space domain image. With each iteration, the space domain image was  $\nabla^2 G$  filtered to



shape its frequency domain characteristic, and then ‘corrected’ in the space domain where it presented an incorrect sign. This algorithm was also found to converge for a number of iterations, and then diverge. As before, the last converging iteration was taken as the best recovered image. The recovered images were found to be very similar in quality to that obtained with the other iterative approach. The ‘girl’  $\nabla^2 G$  pyramid recovered using this scheme is shown in Figure 3.6. The image remerged from this  $\nabla^2 G$  pyramid can be seen in Figure 3.7.

### 3.7 The Importance of Contrast

Close examination of the iteration results yielded some insight into the recovery problem. The recovered images suffered significant loss in shading and tone. This loss seemed most apparent in the large magnitude discontinuities in image intensity (or edges) which were reconstructed in regions presenting only subtle variations in the original image. This observation led to a closer examination of the original image and reconstructed edges.

Image intensity discontinuities (or edges) result in zero crossings in their corresponding  $\nabla^2 G$  filtered images. The variations in a  $\nabla^2 G$  filtered image (which are generated by a discontinuity in image intensity) are bipolar, about the zero crossing. These variations (in isolation) resemble the function shown in Figure 3.2. The amplitude of the  $\nabla^2 G$  filtered function variations are linearly related to the magnitude of discontinuity in the original image intensity. This discontinuity can be thought of as a local contrast at an edge. Large discontinuities in image intensity (or high contrast edges) result in large amplitude variations in the  $\nabla^2 G$  filtered image. Likewise, small discontinuities in image intensity (or low contrast edges) result in small amplitude variations in the  $\nabla^2 G$  filtered image. One can therefore consider a local contrast at each zero crossing (or edge) sample point. Since the magnitude of the original image edge contrast and the amplitude of the variations appearing in the  $\nabla^2 G$  filtered image are linearly related, the local contrast can be estimated in either image.

The images reconstructed from zero crossing locations alone were found to be recon-




Figure 3.6: 'Girl'  $\nabla^2 G$  Pyramid Recovered From ZC Locations



Figure 3.7: Comparison: Best Possible Remerged Image, and Image Reconstructed From ZC Locations: 'Girl'

structed well in the vicinity of high contrast edges. However, edges of like high contrast were also reconstructed in regions where the actual edges were of low contrast. This situation was especially noticeable in regions where the zero crossings resulted from very low magnitude discontinuities (in a region of light texture, for example). This problem was attributed, in part, to the region independence which resulted from sample quantization (as discussed in Section 3.5).

Based on these findings, we were convinced that sampled zero crossing location information alone was inadequate for the practical characterization and good quality recovery of  $\nabla^2 G$  filtered images.

## Chapter 4

# Recovery From Zero Crossings Using Local Estimation

Evidence from both theoretical and practical considerations have shown that sampled zero crossing locations are inadequate for the practical recovery of good quality  $\nabla^2 G$  filtered images. Based on these findings, the use of some additional information was considered for this task. In this chapter, a simple, non-iterative technique will be developed for the recovery of a  $\nabla^2 G$  filtered image from its sampled zero crossings, with local contrast estimates taken at each zero crossing sample.

### 4.1 Motivation

The recovery of images from partial information has been studied extensively. The majority of this work has focused on the image's frequency domain characteristics, in the recovery of the image from phase only [60,61,81], magnitude only [90,75,76,77], or from magnitude or phase [25,26,79]. A summary view of this work appears in [62]. Shitz and Zeevi [85] have also shown the duality of space and frequency domain approaches to the recovery of a signal from partial information.

These studies have generally not considered the application of their findings to image compression. The image's frequency domain phase or magnitude will generally each

contain as much data as that in the original image. While the use of a statistical coding of either of these images could be examined, a choice was made to further examine an edge-based coding and recovery scheme.

In the last chapter, it was pointed out that a discontinuity in original image intensity results in a zero crossing in the corresponding  $\nabla^2 G$  filtered image. Further, the concept of associating a local contrast with each edge was introduced. Briefly, small magnitude discontinuities in original image intensity correspond to low contrast edges; large magnitude discontinuities correspond to high contrast edges. Since the magnitude of original image intensity discontinuity can vary along an edge, such an edge contrast can also vary along an edge.

In studying the recovery of  $\nabla^2 G$  filtered images from their zero crossing locations, it was found that the contrast of the reconstructed image edges was fairly uniform for all edges. The uniformity of the recovered image edge contrast suggested a reconstruction approach which captured some estimate of the edge contrast at each sample point along the edge. This step can also be viewed as extracting some characterization of each local edge segment. The use of a characterization parameter is usually based on some function or source model. Since the characterization of image edges is being considered, image edge models were examined.

The exact characterization of intensity variations in the vicinity of an edge has been found to be quite complex [40]. Haralick [23] has considered the general properties of edges found in an image, however, and concluded that two basic types can occur; step and ramp edges. A step edge occurs between two regions of significantly different image intensity. A ramp edge occurs at the line of change between regions of increasing and decreasing image intensity (like a roof profile).

These definitions are imprecise, however, and ill posed for sampled images without the use of some surface fitting or image smoothing. Edges can occur over a wide range of spatial frequencies (with corresponding spatial extent). A  $\nabla^2 G$  filter was used to select a range of spatial frequencies (with implied image smoothing), and localize the

edges in that frequency range. On close examination of the  $\nabla^2 G$  filtered image edge characteristics, it was found that the profile of the edge was consistently quite similar to the step response of the  $\nabla^2 G$  filter. Restated, it was found that most of the image edges could be modelled as step edges. In the original image, this image is given (in 1-D) by:

$$f(x) = a, \quad x \leq x_1 \quad (4.1)$$

$$f(x) = a + b, \quad x > x_1 \quad (4.2)$$

$$(4.3)$$

where:

$f(x)$  is the original image intensity

$a$  is the image intensity value on one side of the intensity discontinuity

$a + b$  is the image intensity value on the other side of the intensity discontinuity

$x_1$  is the position of the edge

The corresponding  $\nabla^2 G$  filtered image intensity variations are given by the step response of the  $\nabla^2 G$  filter. This intensity profile can be found by numerically convolving the above step function with a 1-D  $\nabla^2 G$  filter, given in the space domain by:

$$\nabla^2 G(x) = \frac{-1}{\sigma^3 \sqrt{2\pi}} \left( 1 - \frac{x^2}{\sigma^2} \right) e^{-x^2/2\sigma^2} \quad (4.4)$$

This step edge property was observed in all four of the  $\nabla^2 G$  filtered image channels, for several different natural scene images. This finding led us to consider the use of an image model for  $\nabla^2 G$  filtered images, based on the step response of the  $\nabla^2 G$  filter. A choice was therefore made to investigate the recovery of  $\nabla^2 G$  filtered images from their zero crossings (edges), with local contrast estimates taken at each zero crossing contour sample.

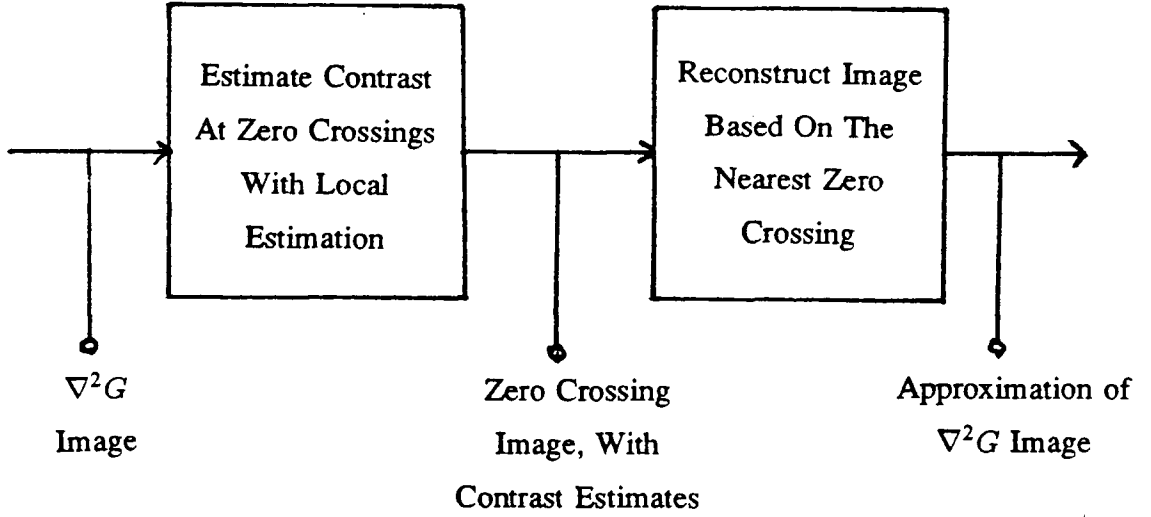


Figure 4.1: Block Diagram: Image Characterization and Reconstruction with Local Estimation

## 4.2 Description of Method

The original images were first filtered into a family of  $\nabla^2 G$  filtered images and their zero crossings localized, using the system described in Chapter 2. The recovery of each of the  $\nabla^2 G$  filtered images was then considered independently. When all four of the  $\nabla^2 G$  filtered images had been reconstructed, these images were remerged to an approximation of the original image (also using the system described in Chapter 2).

Each  $\nabla^2 G$  filtered image was characterized with a set of contrast estimates, found at each zero crossing sample. The  $\nabla^2 G$  filtered image was then reconstructed from this sparse set of data. A block diagram of this sequence of steps is shown in Figure 4.1. As a first approach, a simple method was developed for estimating the contrast at each zero crossing (or edge) sample found in the  $\nabla^2 G$  filtered image. This estimate was found by fitting a template of a  $\nabla^2 G$  filter step response at each zero crossing sample. The procedure for the fitting of this template will be described in the next paragraphs. Con-

sider first, however, the properties of the zero crossings observed in a 2-D  $\nabla^2 G$  filtered image. The zero crossings of 2-D continuous functions form continuous contours [10]. As a result, every point on the continuous contour has an associated spatial orientation within the frame of reference of the image. In this study, the zero crossing contours are represented with discrete zero crossing samples. While these sampled zero crossings do not form continuous contours, they remain connected in adjacency chains (i.e. the next zero crossing sample appears in an adjacent or neighboring sample position). Through these connected chains of zero crossing samples, spatially oriented contours are also approximated. In that sense, each zero crossing sample has an associated spatial orientation in the frame of reference of the image (although this orientation may be ambiguous for isolated or multiply connected zero crossing samples).

The approach taken in estimating the edge contrast is based on a fit of a 1-D step edge model (or template), at each edge sample. Consider the properties of the  $\nabla^2 G$  filtered image in the vicinity of each zero crossing. About each zero crossing contour is a bipolar response, similar to the step response waveform in Figure 3.2. This 1-D response appears along a line normal to the local 2-D zero crossing contour orientation. This relationship can be seen in Figure 4.2. The magnitude of the edge contrast in the original image is linearly related to the magnitude of the bipolar response appearing at the edge in the corresponding  $\nabla^2 G$  filtered image. An estimate of the edge contrast was therefore found by estimating the peak bipolar values (both positive and negative) associated with the edge in the  $\nabla^2 G$  filtered image. The larger of the two magnitudes was taken as the local contrast figure estimate. The positive and negative peaks of the bipolar response both occur approximately 1.4 pixel units from the actual zero crossing for an ideal step input (although obviously on opposite sides of the zero crossing). The actual locations of the peaks vary slightly, depending on the local characteristics of the edge. A local search was therefore conducted to find these peaks of the local bipolar response, on both sides of the zero crossing point.

To help localize the search, an estimate was taken of the local 2-D contour orientation at each zero crossing sample. The orientation estimates were derived from the local adjacency neighborhood. For every possible orientation in the image, two polarity






Figure 4.2: Step Edge Model Configuration

arrangements of bipolar step response are possible. Consider, for example, a horizontal edge across the image. The bipolar step response could be positive above the edge and negative below, or the opposite arrangement could be true. For each possible orientation, one of two polarity arrangements is possible. In this context, only  $180^\circ$  of unique, bipolar orientations exist in a 2-D plane. The orientations of the zero crossing contours were estimated to one of four possible bipolar orientations. These orientations are shown in Figure 4.3.

The search for the bipolar response peaks at each zero crossing sample was conducted along a line normal to the local estimated contour orientation. From the zero crossing sample, the search was extended until the limit of the monotonically increasing (or decreasing) region was found. A maximum search range limit of 2.5 pixels was also imposed. Recall that the  $\nabla^2 G$  filtered images are stored in a pyramid arrangement. As a result, the spatial extent of the step response is the same (in pixel units) for each plane in the  $\nabla^2 G$  pyramid. Refer to Chapter 2 for a review of this pyramid structure.

A single contrast figure magnitude was estimated at each zero crossing sample point,




Figure 4.3: Quantized Edge Orientations

using the preceding method. For each of the orientations shown in Figure 4.3, one of two polarity arrangements is possible. If the positive portion of the bipolar response occurred toward the right half of the plane (in the direction of the arrow in Figure 4.3), the local contrast figure estimate was stored as a positive value. If the opposite polarity were true, a negative contrast figure was stored. In this way, the polarity and magnitude were captured with a single value. Further, it was not necessary to save the local orientations, since these estimates can be derived again from the adjacency neighborhood at the receiver.

To reconstruct the  $\nabla^2 G$  filtered image from these contrast figure estimates, a 1-D template of the  $\nabla^2 G$  filter step response was used (as in Figure 3.2). This template was computed by numerically convolving a step with the space domain  $\nabla^2 G$  filter. This template was then scaled, so that the magnitude of the peaks associated with this response were of value one. Since a  $\nabla^2 G$  filtered image is bandpass with zero gain at  $\omega = 0$ , the mean image value is zero. This fact implies that only those pixels near an edge take on values other than zero.

Recall again that the  $\nabla^2 G$  filtered images are stored in a pyramid arrangement. There-

fore, the spatial extent of the step response is the same (in pixel units) for each plane of the pyramid. For each pixel in the image, a search was conducted to find the nearest zero crossing pixel. The search radius was limited, since the ‘range of influence’ of a zero crossing is about 6 pixels (with a maximum contrast edge and 8 bits/pixel resolution). For each pixel less than 6 pixel units in distance from a zero crossing sample, an intensity value was computed. First, the distance between the zero crossing and the pixel to be shaded was computed (with simple geometry). The template function value at the corresponding distance from the zero crossing was then found. The contrast value stored at each zero crossing sample represents the peak amplitude of the characteristic bipolar step response, found during the estimation step. This contrast value was used to scale the template value. The sign of this intensity value was then determined, by the positional relation of the pixel being shaded to the orientation and polarity of the contrast estimate found during the estimation step. Stated simply, it was necessary to decide if the pixel being shaded was on the negative or positive side of its nearest edge.

### 4.3 Experimental Results

In spite of its simplicity, fairly good image reconstruction was obtained with this scheme. For sparse images well modelled with the step edge model, the reconstruction was predictably quite good. The reconstructed ‘cake’  $\nabla^2 G$  pyramid can be seen in Figure 4.4. A comparison of the best possible reconstruction and the reconstructed ‘cake’ can be seen in Figure 4.5. This technique also worked fairly well on complex natural scenes. The reconstructed ‘girl’  $\nabla^2 G$  pyramid can be seen in Figure 4.6. While some reconstruction error is apparent in the  $\nabla^2 G$  pyramid, this error seems less apparent in the remerged image (seen in Figure 4.7).




Figure 4.4: Reconstructed  $\nabla^2 G$  Pyramid, Using Local ZC and Local Estimation: 'Cake'




Figure 4.5: Comparison: Best Possible Remerged Image, and Image Reconstructed Using ZC and Local Estimation: 'Cake'




Figure 4.6: Reconstructed  $\nabla^2 G$  Pyramid, Using Local ZC and Local Estimation: 'Girl'



Figure 4.7: Comparison: Best Possible Remerged Image, and Image Reconstructed Using ZC and Local Estimation: 'Girl'

## 4.4 Applicability and Limitations

The described local estimation technique is non-iterative, does not involve the solution of simultaneous equations, and is fairly simple to implement. It therefore avoids the convergence problems and large numbers of computations often faced with the use of iterative methods. Further, it avoids the numerical problems which often occur in the solution of simultaneous equations. It is, however, based on several important assumptions which limit its applicability. The most significant assumption is that the edges are independent. For each pixel being shaded, it was assumed that only the nearest zero crossing pixel influenced its brightness. This assumption is absolutely true only for sparse images, where the image edges are separated by a minimum of about 12 pixels (with 8 bits/pixel resolution). For images with edges spaced more closely, the edges tend to be reconstructed with too much contrast. Error in reconstruction is apparent for this type of image. The reconstructed image quality was still fairly good, however, with much of the image shading and tone preserved.

The assumption of independence could be overcome by computing the superimposed effect from all influencing zero crossings, but this measure requires finding the non-superimposed source values at each zero crossing point. This deconvolution step is formidable, and was therefore not examined further in this portion of the study.

The other significant assumption is that of the step image model. Not all possible edges in an image will be adequately modelled with a step model. For the images considered in this study, however, this assumption was not found to be significant.

In summary, this reconstruction technique provides a stable method for the recovery of  $\nabla^2 G$  filtered images from their sampled zero crossings, with local contrast figure estimates taken at each zero crossing sample. This technique does not involve the solution of simultaneous equations, or the use of iteration. The quality of recovered image depends on the nature of the image.

## Chapter 5

# Recovery From Zero Crossings Using Filtering

In the last chapter, a method was developed for recovering a  $\nabla^2 G$  filtered image from its zero crossings using local estimation. The most significant shortcoming of this method was the assumption that the waveforms from each zero crossing contour did not interfere with waveforms from other, nearby zero crossing contours. For many natural scene images, this assumption is not true. While it would be relatively straightforward to compute the superimposed effect of several edges for any given pixel, it is nontrivial to compute the non-superimposed sources for each zero crossing sample. The most direct method for computing these deconvolved sources would involve solving sets of simultaneous equations for each zero crossing sample. In this chapter, a more elegant method for computing these non-superimposed sources will be developed using inverse 2-D filtering. A 2-D reconstruction filter will also be developed, for recovering a  $\nabla^2 G$  filtered image from these inverse-filter-derived zero crossing estimates.

### 5.1 Description of Method

Again, the motivation here was to consider the recovery of  $\nabla^2 G$  filtered images from their zero crossings, with contrast figure estimates taken at each zero crossing sample

(see Section 4.1). As discussed, the most significant shortcoming of the local estimation technique was the assumption of edge independence. An investigation of this problem led to a solution based on 2-D filtering.

The key to the filtering approach used here stemmed from a desire to examine the retention of only contrast figure estimates, saved at each zero crossing sample. The contrast data can be viewed as a train of (scaled) impulses, standing out of the image plane. A filter was desired that would accept such a contrast impulse image, and would generate a  $\nabla^2 G$  filtered-like output image. This output  $\nabla^2 G$  image would be the same as would be produced by  $\nabla^2 G$  filtering an original image containing a set of step-like edges at each of the impulse sample points.

The image edge model assumed here for all edges is the same as the model used in the local estimation method. In the original image, edges are modelled (in 1-D) by:

$$f(x) = a, \quad x \leq x_1 \quad (5.1)$$

$$f(x) = a + b, \quad x > x_1 \quad (5.2)$$

$$(5.3)$$

where:

$f(x)$  is the original image intensity

$a$  is the image intensity on one side of the edge

$a + b$  is the image intensity on the other side of the edge

$x_1$  is the position of the edge

The corresponding  $\nabla^2 G$  filtered image response  $s(x)$  can be found by numerically convolving the above edge model function with the 1-D  $\nabla^2 G$  expression as follows:

$$\nabla^2 G(x) = \frac{-1}{\sigma^3 \sqrt{2\pi}} \left( 1 - \frac{x^2}{\sigma^2} \right) e^{-x^2/2\sigma^2} \quad (5.4)$$

$$s(x) = f(x) \odot \nabla^2 G(x) \quad (5.5)$$

A processing step is therefore required which will transform the image of scaled impulses, located at the zero crossing sample points, to the required 2-D step response-like waveforms. An impulse can be filtered into a given waveshape by constructing a



filter with an impulse response given by that waveshape. By developing a filter with an impulse response equal to the step response of a  $\nabla^2 G$  filter, a set of impulse values can be filtered into a step response-like waveform. With a linear filter, the superimposed effect of all the impulse input values will be computed. Further, the scale of the step response-like waveforms is regulated by the magnitude of the impulse values. In this way, a set of magnitude values placed at the zero crossing samples can be transformed into a set of scaled and superimposed step response-like waveforms.

The problem remains, however, of how to compute the non-superimposed source estimates at the zero crossing samples. These values can be computed using the inverse of the reconstruction filter already discussed. The zero crossing estimates can then be computed by applying the inverse filter to the input  $\nabla^2 G$  filtered image, and sampling the filtered result at the zero crossing samples of the input  $\nabla^2 G$  image. This sequence of steps is illustrated in Figure 5.1. A block diagram of the sequence of filtering steps is shown in Figure 5.2.

## 5.2 Design of Reconstruction Filter - No Smoothing

The start of the filter design must begin with the one dimensional step response of a  $\nabla^2 G$  filter (given in Equation 5.5 and shown in Figure 5.2), since this is the filter impulse response required. It is not immediately apparent, however, how this 1-D response should be transformed into 2-D. Consider the characteristics of the required 2-D response. The zero crossings of 2-D functions form contours. As a result, an isolated zero crossing point cannot occur (a function zero does not constitute a zero crossing). The 2-D reconstruction filter must therefore produce the required 2-D step response waveform from the combined effects of a train of adjacent impulse values. Restated, the superimposed effect of the application of the reconstruction filter to a train of impulse values must yield the required 2-D step response waveform, through the zero crossing contour.

Again, a choice was made to implement the filters using a FIR filtering approach, implemented with FFTs. A choice was also made to design the 2-D reconstruction filter in

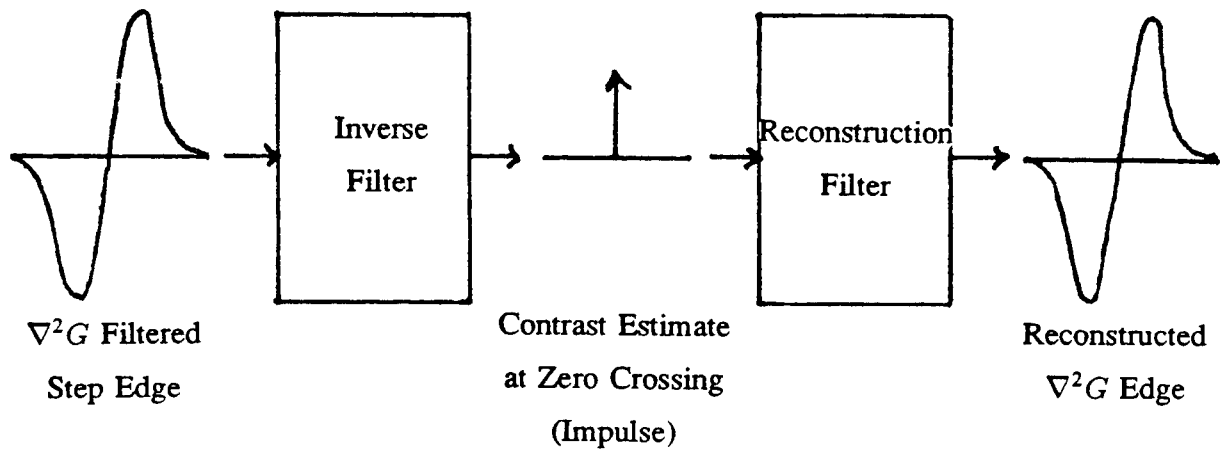


Figure 5.1: Illustration: Inverse and Reconstruction Filtering Steps




Figure 5.2: Block Diagram: Inverse and Reconstruction Filtering Steps




Figure 5.3: 1-D Space Domain Reconstruction Filter

the frequency domain. A reconstruction filter is required which has an impulse response given by the 1-D  $\nabla^2 G$  filter step response. A filter's impulse response is given by its space domain characteristic. The 1-D step response was therefore constructed in the space domain, centered at  $x = 0$  (this function can be seen in Figure 5.3). This step response was computed by numerically convolving a step with the 1-D space domain  $\nabla^2 G$  filter.

This function is symmetric about the origin ( $x = 0$ ), with odd symmetry. The odd symmetry implies that the filter will be purely imaginary in the frequency domain, and non-causal. Since the whole image is available at the time of processing, the causality does not present a problem.

The frequency response of this 1-D filter is symmetrical about  $\omega = f_s/2$ , and of odd symmetry. A plot of the imaginary portion of the frequency response can be seen in Figure 5.4. The real part of the frequency response equals zero. The frequency response was rotated so as to satisfy continuity constraints, and to produce a filter which was symmetrical to produce a 'real-only' space domain filter. Again, the filter's frequency response was constructed on a complex, 256 X 256 plane. A 1-D 1024 point estimate of the reconstruction filter frequency response was computed, to provide a closely spaced set




Figure 5.4: Imaginary Portion of the Frequency Response of a 1-D  $\nabla^2 G$  Filter Step Response

for use in the rotation procedure. The rotation in the frequency domain was accomplished as follows. To clarify this discussion, the 2-D frequency domain plane will be divided into four quadrants and labeled as shown in Figure 5.5. The frequency response shown in Figure 5.4 was placed on the  $\omega_1 = 0$  and  $\omega_2 = 0$  axis. This step is required for the proper filtering of an input signal with frequency components along only  $\omega_1$  or  $\omega_2$ . The response in quadrant 1 was then constructed by rotating the 1-D  $\omega = \{0 \rightarrow f_s/2\}$  response, with  $\omega = 0$  centered at  $\omega_1 = \omega_2 = 0$ . Similarly, the 1-D response from  $\omega = \{f_s/2 \rightarrow f_{s-1}\}$  was rotated about  $\omega_1 = \omega_2 = f_{s-1}$  in quadrant 4, with  $\omega = f_{s-1}$  centered at  $\omega_1 = \omega_2 = f_{s-1}$ . This partial result is illustrated in Figure 5.6.

To realize a real valued space domain filter, the frequency response must exhibit symmetry in the frequency domain. The frequency response must have complex-conjugate symmetry about the  $\{\omega_1 = 0; \omega_2 = f_{s-1}\}$  to  $\{\omega_1 = f_{s-1}; \omega_2 = 0\}$  diagonal. Further, it is desirable for the frequency response to be smooth and continuous. Consider the frequency response in quadrants 2 and 3. The positive halves of the bipolar 1-D frequency




Figure 5.5: Frequency Plane Organization And Labelling

response lie on the  $\omega_1 = 0$  and  $\omega_2 = 0$  axis in these quadrants (as shown in Figure 5.6). The positive rotated response in quadrant 4, however, suggests a negative response over segment *a* and *b* in Figure 5.6. This implies that the responses in quadrants 2 and 3 will be positive on one axis, and negative on the other. As a first implementation, the responses in quadrants 2 and 3 were therefore rotated with a discontinuity along the  $\{\omega_1 = 0; \omega_2 = f_{s-1}\}$  to  $\{\omega_1 = f_{s-1}; \omega_2 = 0\}$  diagonal. This measure satisfies the outer axis continuity, and the complex conjugate symmetry required for realizing a real valued space domain filter.

With the desired 2-D frequency response constructed, the inverse 2-D FFT of the filter was computed. The 2-D space domain coefficients were then examined, to see if the number of coefficients could be truncated for convenient storage. The space domain coefficients were found to be centered about the origin, with most significant variations within a region of about 13 X 13 coefficients about the origin. Because of the discontinuity in the frequency domain, however, some small magnitude ringing was present in the space domain filter. A large number of space domain coefficients were therefore saved; 49 X




Figure 5.6: 2-D Reconstruction Filter (Frequency Domain), Specified in quadrants 1 and 4, and along the  $\omega_1 = 0$  and  $\omega_2 = 0$  axis

49 coefficients for all pyramid planes except the 32 X 32 plane, where 32 X 32 space domain coefficients were saved. Because of the complexity of the 2-D filters, plots of both the frequency domain imaginary portion and magnitude (absolute value of the imaginary portion, since the real part equals zero) are presented. Plots of the reconstruction filter frequency response can be seen in Figures 5.7 and 5.8 respectively. These 2-D plots have been rearranged from the frequency domain layout used in Figure 5.5. Each of the Figure 5.5 quadrants have been reorganized so that the  $\omega_1 = \omega_2 = 0$  is set at the center of the 2-D plot. Further, these plots have been rotated to provide the most illustrative viewpoint.

### 5.3 Design of Inverse Filter - No Smoothing

Having designed the reconstruction filter, it is necessary to consider how to compute the contrast figure estimates to be placed at the zero crossing samples. This deconvolution step, as discussed, is formidable when computed using simultaneous equations. These estimates can be computed, however, using a filter which is the inverse to the reconstruction filter. The successive application of the inverse and reconstruction filters on a  $\nabla^2 G$  filtered image should yield back the same  $\nabla^2 G$  filtered image.

For the successive application of the inverse and reconstruction filters to reproduce the input image, the frequency domain product of these two filters must equal unity for all frequencies. Two problems arise in designing such an inverse filter. Clearly, if the reconstruction filter were zero gain for some frequencies, it would be impossible to recover that spectral portion of the input image using this scheme. Secondly, the designed inverse filter will tend to infinite gain where the reconstruction filter tends to zero. These problems are related, since they both concern portions of the reconstruction filter spectra where the gain tends to zero. Fortunately, the spectral characteristic of the reconstruction filter is quite like that of the  $\nabla^2 G$  filter. Both exhibit a bandpass characteristic, which rolls off smoothly. Since the  $\nabla^2 G$  filtered images have a low energy content approaching  $\omega = 0$  and  $\omega = f_s/2$ , the loss in reconstruction in these regions should have little impact.




Figure 5.7: 2-D Reconstruction Filter Frequency Response, Imaginary Part, No Smoothing. Note: the center of the 2-D plot represents  $\omega_1 = \omega_2 = 0$ .




Figure 5.8: 2-D Reconstruction Filter Frequency Response Magnitude, No Smoothing. Note: the center of the 2-D plot represents  $\omega_1 = \omega_2 = 0$ .



As a result, a choice was made to limit (clip) the inverse filter at the point where the reconstruction filter's gain was down by a factor of 10 from the peak filter gain.

The frequency response of the 2-D reconstruction filter was used to compute the inverse filter frequency response. The reconstruction filter frequency response is imaginary, with its real part equal to zero. The inverse filter was computed so that the frequency domain product of the reconstruction and inverse filters was unity gain. This step was accomplished by the following steps:

$$a = \Im\{f(\omega_1, \omega_2)\} \quad (5.6)$$

$$b = 1.0/a \quad (5.7)$$

$$g(\omega_1, \omega_2) = z(0.0, b) \quad (5.8)$$

where:

$f(\omega_1, \omega_2)$  is the (complex) reconstruction filter frequency response,

$\Im\{f(\omega_1, \omega_2)\}$  is the imaginary portion of this response,

$z(real, imag)$  is an operator to create a complex value, and

$g(\omega_1, \omega_2)$  is the computed inverse filter frequency response.

The reconstruction filter tends to zero near  $\omega = 0$  and  $\omega = f_s/2$ . The inverse filter magnitude  $b$  was clipped at the point where the reconstruction filter gain was down by a factor of 10 (20 dB) from its peak value. Since the peak reconstruction filter gain was scaled to be unity gain, this implies clipping the inverse filter gain at a value of 10. The sign of the clipped value was the same as that of the reconstruction filter. This implies discontinuities at  $\omega_1 = f_s/2$ ,  $\omega_2 = f_s/2$ , and along the diagonal from  $\{\omega_1 = 0; \omega_2 = f_{s-1}\}$  to  $\{\omega_1 = f_{s-1}; \omega_2 = 0\}$ . The 2-D inverse FFT of this frequency response was then computed, to obtain the space domain inverse filter coefficients. Again, the space domain coefficients were examined, to see if the number of space domain filter coefficients could be reduced for convenient storage. 49 X 49 space domain coefficients were used for all planes except the 32 X 32 plane, where 32 X 32 space domain coefficients were used. Imaginary part and magnitude plots of this filter's frequency response are presented in Figures 5.9 and 5.10.




Figure 5.9: 2-D Inverse Filter Frequency Response, Imaginary Part, No Smoothing. Note: the center of the 2-D plot represents  $\omega_1 = \omega_2 = 0$ .




Figure 5.10: 2-D Inverse Filter Frequency Response Magnitude, No Smoothing. Note: the center of the 2-D plot represents  $\omega_1 = \omega_2 = 0$ .

## 5.4 Experimental Results - No Smoothing

The non-smoothed inverse and reconstruction filters designed in the preceding sections were tested on several images. Predictably, some ringing was found in both the inverse and reconstructed images. This ringing occurred primarily along a  $45^\circ$  line across the image, stemming from the discontinuities in the inverse and reconstruction filter frequency responses. In spite of this ringing, fairly good image reconstruction was obtained. The inverse filtered 'girl' pyramid is shown in Figure 5.11. The ringing is most noticeable at high contrast edges oriented at  $45^\circ$ , from upper left to lower right. This can be seen near the peak of the girl's hat. The pyramid reconstructed from these inverse images is shown in Figure 5.12. A comparison of the merged original  $\nabla^2 G$  filtered image (the best possible reconstruction) and the merged reconstructed images is shown in Figure 5.13. While some oriented smearing is apparent, the recovered image is quite an accurate rendering of the original image. Only slight loss in tone and shading are apparent.

The impact of sampling the inverse filtered images only at the zero crossings of the  $\nabla^2 G$  filtered image was then examined. The  $\nabla^2 G$  filtered images will be reconstructed to the degree that the original image can be modelled by the step image model. These reconstructed images were found to retain most of the image information, but some further loss in image shading and tone was apparent. Again, some oriented smearing occurred from high contrast edges oriented at  $45^\circ$  in the the image. A comparison of the best possible merged image and the image recovered from only the zero crossing samples is shown in Figure 5.14.

## 5.5 Design of Reconstruction Filter - Smoothed

The oriented smearing observed in the previous section was examined further, and it was found that the primary cause was the diagonal discontinuity in the inverse and reconstruction filters. Smoothing was therefore applied to these discontinuities. For the reconstruction filter, the discontinuity cuts diagonally across quadrants 2 and 3. To




Figure 5.11: Inverse Filtered Pyramid, No Smoothing: 'Girl'




Figure 5.12: Reconstruction Filtered Pyramid, No Smoothing: 'Girl'



Figure 5.13: Comparison: Best Possible Remerged Image, and the Image Recovered Using Reconstruction Filtering - No Smoothing: 'Girl'



Figure 5.14: Comparison: Best Possible Remerged Image, and the Image Recovered Using Reconstruction Filtering, Using Only the Zero Crossing Samples - No Smoothing: 'Girl'

construct a smooth frequency response, a choice was made to apply sinusoidal smoothing over these quadrants. The filter gain was scaled as a function of quadrant angle, so that zero gain would be applied along the diagonal line. A scale factor of magnitude 1 was used along the  $\omega_1 = 0$  and  $\omega_2 = 0$  axis. With a sinusoidal scaling (over  $\pi/2 \rightarrow 3\pi/2$ ), a smooth frequency response was obtained.

The 1-D reconstruction filter frequency response was rotated and smoothed with the sinusoidal profile, and the 2-D inverse FFT computed. The space domain coefficients were found to exhibit less spread than was apparent with no smoothing. The number of space domain coefficients was not changed from that used in the non-smoothed design. Plots of this filter's frequency response imaginary portion and magnitude can be seen in Figures 5.15 and 5.16 respectively.

## 5.6 Design of Inverse Filter - Smoothed

A new inverse filter was designed to twin the smoothed reconstruction filter. With the smoothing applied to the reconstruction filter, discontinuities occur in the inverse filter along the  $\omega = f_s/2$  boundaries, along the diagonals across quadrants 2 and 3, and at the origin. The diagonal discontinuity was smoothed using a sinusoidal smoothing as a function of the angle in the quadrant, in a manner similar to that used for the reconstruction filter. The discontinuities at  $f_s/2$  were smoothed by rolling off the gain, with a rotationally symmetric cosine profile. As already noted, the source  $\nabla^2 G$  filtered images are bandpass in nature. There is little loss, therefore, in rolling off the inverse image in the region where the  $\nabla^2 G$  signal is very small. A choice was made to roll off the inverse filter from the point where the input image is down by a factor of 10 (20 dB). The reconstruction filter is scaled for unity maximum gain, and this roll off point therefore occurs at the point where the inverse filter gain is of magnitude 10 (at approximately  $0.3125f_s$  in the 1-D frequency response). For portions of the 2-D inverse frequency response derived from the 1-D reconstruction filter defined over  $\omega = \{0 \rightarrow 0.3125f_s\}$  and  $\omega = \{(1.0 - 0.3125)f_s \rightarrow f_{s-1}\}$ , no roll off is applied. For portions of the 2-D inverse




Figure 5.15: 2-D Reconstruction Filter Frequency Response, Imaginary Part, With Smoothing. Note: the center of the 2-D plot represents  $\omega_1 = \omega_2 = 0$ .




Figure 5.16: 2-D Reconstruction Filter Frequency Response Magnitude, With Smoothing. Note: the center of the 2-D plot represents  $\omega_1 = \omega_2 = 0$ .

filter frequency response outside these regions, the following roll off was applied:

$$g(\omega_1, \omega_2) = 0.5(\cos(z) + 1.0), \quad z < \pi \quad (5.9)$$

$$g(\omega_1, \omega_2) = 0.0, \quad z > \pi \quad (5.10)$$

$$z = (\sqrt{x_1^2 + y_1^2} - 0.3125)\pi / (0.5 - 0.3125) \quad (5.11)$$

where  $x_1$  and  $y_1$  are given by:

Quadrant 1:  $x_1 = \omega_1, y_1 = \omega_2$

Quadrant 2:  $x_1 = 1.0 - \omega_1, y_1 = \omega_2$

Quadrant 3:  $x_1 = \omega_1, y_1 = 1.0 - \omega_2$

Quadrant 4:  $x_1 = 1.0 - \omega_1, y_1 = 1.0 - \omega_2$

with  $\omega_1$  and  $\omega_2$  given in normalized frequency units.

The inverse filter frequency response was constructed using the smoothed reconstruction filter frequency response, and the smoothing described in the preceding paragraph. The inverse 2-D FFT was then computed, and the same number of filter coefficients saved in the space domain as that used with the non-smoothed inverse filter. Plots of this filter's frequency response imaginary portion and magnitude can be seen in Figures 5.17 and 5.18 respectively.

## 5.7 Experimental Results - With Smoothing

The smoothed reconstruction and inverse filters were evaluated on several test images. The ringing observed in the non-smoothed case was found to be significantly reduced, although not totally eliminated. A cost of this smoothing, however, was some loss of edge contrast for edges oriented on the 45° diagonal. This loss was to be expected, since the gain of the inverse and reconstruction filters is dampened along this orientation. Overall, good image reconstruction was retained, with only slight loss in image tone and shading. A slight loss in image sharpness was noted from the non-smoothed to smoothed recovered images. The smoothed images, however, exhibited reduced image error in the vicinity of high contrast edges oriented on the diagonal.






Figure 5.17: 2-D Inverse Filter Frequency Response, Imaginary Part, With Smoothing. Note: the center of the 2-D plot represents  $\omega_1 = \omega_2 = 0$ .




Figure 5.18: 2-D Inverse Filter Frequency Response Magnitude, With Smoothing. Note: the center of the 2-D plot represents  $\omega_1 = \omega_2 = 0$ .

The inverse filtered ‘girl’ pyramid is shown in Figure 5.19. The pyramid reconstructed from these inverse images is shown in Figure 5.20. Also shown is the reconstruction filtered ‘cake’ pyramid, in Figure 5.21. This image was included to illustrate the oriented edge dampening in isolation. A comparison of the best possible reconstruction and the merged reconstructed images is shown in Figure 5.22. The oriented smearing can be seen to be significantly reduced from that in the non-smoothed case. Overall, a fairly accurate rendering of the original image is obtained. Only slight loss in tone and shading are apparent.

Again, the impact of sampling the inverse filtered images only at the zero crossings of the  $\nabla^2 G$  filtered image was examined. The images reconstructed from this sparse information were also found to retain most of the image information, but some further loss in image shading and tone was apparent. The oriented smearing was also found to be significantly diminished from the non-smoothed case. A comparison of the best possible merged image and the image recovered from only the zero crossing samples is shown in Figure 5.23.

## 5.8 Applicability and Limitations

The described estimation and recovery method is non-iterative, does not involve the solution of simultaneous equations, and is fairly simple to implement. It avoids the convergence problems and large number of computations often incurred with an iterative approach [73]. Further, it avoids the numerical problems often faced with the solution of simultaneous equations [86]. The approach has been experimentally found to be stable (small change or error in the input leading to a small change in the output), for all images tested (5 widely different images were used as test cases). The most significant short-coming of this techniques stem from the discontinuities in the filter frequency responses. It is impossible to realize a filter with an abrupt frequency domain discontinuity without ringing in the filtered result. Smoothing the discontinuity reduces the ringing, but also tends to diminish the image intensity in the vicinity of the smoothing. The image ringing




Figure 5.19: Inverse Filtered Pyramid, With Smoothing: 'Girl'




Figure 5.20: Reconstruction Filtered Pyramid, With Smoothing: 'Girl'




Figure 5.21: Reconstruction Filtered Pyramid, With Smoothing: 'Cake'



Figure 5.22: Comparison: Best Possible Remerged Image, and the Image Recovered Using Reconstruction Filtering - With Smoothing: 'Girl'

manifests itself as oriented smearing, but is quite predictable, and fairly contained (spatially). For more sparse images, such as the ‘cake’, recovery with the local estimation method was better than that with the filtering method. The local estimation method was better for images of this type, because the isolated edges were reconstructed well and the problems of smearing with the filtering method were avoided. For more complex images, such as the ‘girl’, the reverse situation was true. For these images, the smearing introduced with the filtering method seemed less apparent than for sparse images reconstructed with this technique. Moreover, the edge reconstruction error which occurred for complex images reconstructed with the local estimation technique seemed to have a much higher visual impact than the smearing introduced with the filtering approach. In either case, recovered image quality remained fairly high, with most of the shading and tone preserved.

Sampling the inverse filtered images only at the zero crossings of the  $\nabla^2 G$  filtered input images resulted in some loss in reconstructed image quality. Still, the images remained quite recognizable, with most of the shading preserved. The loss associated with each stage of this processing can be seen for the ‘girl’ image in Figure 5.24. This figure shows the original image, the best remerge of the complete  $\nabla^2 G$  pyramid, the image recovered using smoothed inverse and reconstruction filtering with all the inverse filtered image samples, and the same filtering sequence with the inverse filtered image sampled only at the zero crossings of the  $\nabla^2 G$  filtered image. The same sequence for the ‘cake’ image is shown in Figure 5.25.



Figure 5.23: Comparison: Best Possible Remerged Image, and the Image Recovered Using Reconstruction Filtering, Using Only the Zero Crossing Samples - With Smoothing: 'Girl'





Figure 5.24: Comparison of Images, Showing Progressive Loss By Processing Step For 'Girl': original image (upper left), best remerge of original  $\nabla^2 G$  pyramid (upper right), image recovered with smoothed inverse/reconstruction filtering and all samples (lower left), and image recovered with smoothed inverse/reconstruction filtering sampled only at the zero crossings of the  $\nabla^2 G$  filtered image (lower right)




Figure 5.25: Comparison of Images, Showing Progressive Loss By Processing Step For 'Cake': original image (upper left), best remerge of original  $\nabla^2 G$  pyramid (upper right), image recovered with smoothed inverse/reconstruction filtering and all samples (lower left), and image recovered with smoothed inverse/reconstruction filtering sampled only at the zero crossings of the  $\nabla^2 G$  filtered image (lower right)



# Chapter 6

## Application: Image Compression

Having developed a stable method for recovering a  $\nabla^2 G$  filtered image from its sampled zero crossings (with contrast figure estimates), the applicability of this technique for image compression was examined. The investigation of this application did not form a large part of this study, but is included here to provide some idea of the applicability of the technique.

A full examination of the compression of an image through the coding of its edges is a large and complex problem. In this chapter, a simple edge coding scheme based on linear vector regression will be developed, and the image compression that can be obtained with this approach examined. The effects of two tuning parameters on the recovered image quality and image compression rate will also be examined.

### 6.1 Representation of Edges With Vector Segments

In this study, a method for recovering an image from its zero crossings or edges has been developed. While this zero crossing image is sparse, it is not obvious what the best technique would be for compressing these edges in some coded form. If a lossy form of coding is considered, it is not clear what the impact of approximating these edges would be on reconstructed image quality. Further, effective image compression for a sequence

of images (such as the underwater transmission problem discussed in Chapter 1) requires an examination of in-the-frame, and frame-on-frame compression.

These broad questions will not be fully considered in this study. In order to obtain some idea of the compression possible from these edges, however, a simple representation of the image edges will be considered. Schreiber et al [82] have considered the coding of image contours, and have identified four methods for accomplishing this task:

1. Direct position transmission
2. Run-length coding
3. Direct contour tracing, and
4. Fitting curves to the contours

Direct position transmission involves the sending of each of the contour coordinate pairs. This technique is reasonable only for very sparse images, and was not considered further for this application. Run-length coding involves the transmission of scan lines, with codes representing the distances between contour points. This method was not considered for several reasons. First, the magnitude must be coded at each zero crossing sample. Secondly, the contours are highly connected. It was also found that the zero crossing samples were highly correlated in a connected contour. Since simple run-length coding does not lend itself to capitalize on these properties, it was not considered further. Direct contour tracing involves a coding technique which follows each zero crossing contour. This approach was considered, but not examined further due to time constraints. A choice was made, instead, to examine the fitting of curves to the contours. Curve fitting was chosen over direct contour tracing because the curve fitting lent itself to the examination of contour smoothing, and tuned approximation of the edge contour locations.

Consider the properties of the zero crossing contours which need to be represented. For continuous images, the zero crossing contours always close on themselves (or off the frame). Studies have capitalized on this property in the coding of closed contours

[6,30,101]. Unfortunately, the sampled zero crossing contours observed in the natural scene test images did not close on themselves. This problem can be visualized with an image composed of a ramp-like intensity profile, set in an isolated rectangle. The discontinuities in brightness around the top and sides of the ramp will clearly generate edges. If the bottom of the ramp is smoothed into the background plane, however, the magnitude of the edge appearing in the  $\nabla^2 G$  filtered image will be arbitrarily small. If this  $\nabla^2 G$  filtered image is sampled, the zero crossing over this region may not be detected. As a result, the zero crossings samples localized in a  $\nabla^2 G$  filtered image do not always close on themselves.

As a result, contour fitting methods for non-closed contours were considered. Since both contour location and contrast figure estimates need to be coded, the variability of the the contrast figure estimates along connected contours was examined. A high degree of variability in contrast figure estimate was found to be possible along a connected contour. An extreme example can be seen with the 'cake' image (see Figure 2.9), where only two connected edges occur. Each of these edges rotates through all possible orientations, representing a constant edge contrast. This implies that the whole dynamic range of possible contrast figure estimates can be presented along such a connected contour.

If a low order model curve fit is imposed on the zero crossing contours, it was felt that it might be possible to represent each contour segment with its location parameters and one average contrast figure for the whole contour segment. If a high order model curve is used, it would be necessary to consider a model for the contrast figure estimates within the contour segment. As a function of simplicity, a choice was made to implement a low order approach.

A linear regression approach was therefore used to fit first order segments through the zero crossing samples. Along each first order segment, the average of the zero crossing sample values was computed and taken as the segment contrast figure. The result of the regression step was a table of extracted vector endpoints and contrast figures, organized into connected contours. The goal of the regression procedure was to fit the longest possible first order segments through the connected zero crossing samples. This was

achieved by successively projecting a first order model over a connected set of zero crossing sample points, and examining the maximum distance between the line model and the underlying zero crossing samples. A linear model with zero distance error can always be found for two connected zero crossing samples. An attempt was made to extend this model over any connected extensions which might exist in the zero crossing samples. An eight neighbor connect rule is used in this search, with the nearest neighbors added to the chain first.

To simplify the computations, a choice was made to represent the linear segments using the form:

$$y = mx + b \quad (6.1)$$

This form simplifies the regression computation, but singularities (vertical lines) must be considered. To address this problem, this form was used with exchanged coordinates when a line with slope greater than  $45^\circ$  was detected. A least squares approximation fit was used to compute these line parameters. A computational procedure for this calculation is given by Strang [86]:

$$m = \frac{n \sum x_i y_i - \sum x_i \sum y_i}{n \sum x_i^2 - (\sum x_i)^2} \quad (6.2)$$

$$b = \frac{\sum y_i \sum x_i^2 - \sum x_i y_i \sum x_i}{n \sum x_i^2 - (\sum x_i)^2} \quad (6.3)$$

This procedure minimizes the squared distance from each zero crossing sample point to the line model, along the  $y$  axis.<sup>1</sup> The impact of this approximation was reduced through the exchanging of the coordinates when a large slope was detected. The minimum distance from the line model to the zero crossing samples was used for the distance check computation, with:

$$d^2 = (x_i - X_0)^2 + (y_i - Y_0)^2 \quad (6.4)$$

where  $X_0, Y_0$  is the point on the model which is the intersection of the line normal to the

---

<sup>1</sup>The actual distance which should be minimized is that between the point in question, and the nearest point on the line. When this algorithm was implemented, the author believed that this calculation involved many more computations than the procedure for minimizing the distance along the  $y$  axis. Subsequently, an algorithm involving few additional computations has been pointed out by members of the supervisory committee.

model through the zero crossing sample  $x_i, y_i$  under consideration.  $X_0, Y_0$  are given by:

$$X_0 = \frac{(x_i/m) + y_i - b}{m + (1/m)} \quad (6.5)$$

$$Y_0 = \frac{(x_i + my_i + (b/m))}{m + (1/m)} \quad (6.6)$$

If a zero denominator is computed in the least squares computation, a vertical line has occurred. In that case, the coordinates are exchanged and the line recomputed. The use of this least squares algorithm was found to yield erroneous results, however, under the following circumstances. If a set of zero crossing samples imply an exactly vertical line, but some distance will exist from the samples to the line, a line normal to the vertical is computed with the least squares procedure. This situation will exist, for example, with a vertical set of samples, where the top and bottom samples are set one unit to the left of the others (as illustrated in Figure 6.1). Clearly, the best fit is a vertical line. The least squares computation does not result in a zero value denominator, and computes a horizontal line through the axis of the vertical group. This problem was overcome by heuristically watching for this situation.

The implemented procedural method for conducting the regression step can be described with the following sequence of instructions:

1. Find a zero crossing segment which has not yet been considered in a linear segment model. When all such samples have been considered, proceed to step 4.
2. From the current (or last) zero crossing sample, find the neighboring zero crossing samples (8 neighbor connectivity), and add the not yet considered samples to the current chain of zero crossing samples (nearest samples first).

*If {no not-yet-considered adjacent zero crossing samples are found from the current sample} Then*

Back search the current chain to find any connected branches not yet considered.

*If {no such branches are found} Then*




Figure 6.1: Illustration of Least Squares Problem Situation

Save the current model and terminate the current chain.

*Else*

Proceed to step 1.

*End If*

*Else*

Continue.

*End If*

3. Compute the model for the current chain of zero crossing samples, and compute the distance from each of the zero crossing chain members to the newly computed model.

*If* {the maximum distance measure has been exceeded by one or more of the recently added zero crossing samples} *Then*

Drop these samples off (one at a time), until an adequate model has been obtained. An adequate model can always be obtained with two samples. Once an adequate model has been obtained, save the current model and terminate the current chain. The last zero crossing from the current chain is used as the starting point for the next chain. Proceed to step 1.

*Else*

Proceed to step 2.

*End If*

4. Examine the connectivity of the extracted models, and attempt to connect the models whose underlying samples were connected. This is done by projecting the segments, and extending them to a junction (if the extension is less than 5 maximum distance measures in length), or through the addition of a connecting segment. The modelling step is then complete.

## 6.2 Approximation Parameters

Using the described regression and coding scheme, some loss will occur even with a zero distance error, since the contrast figure estimates are averaged over each linear segment. With a zero distance error, however, the linear segments will be very short. As a result, this contrast figure averaging should not have too great a visual impact. Allowing spatial error in the regression step results in error in the localization of the image edges. Consider, for example, the 'cake' image. With a zero distance error, a linear segment could be required for each pair of connected zero crossing samples. As the distance error is increased, fewer segments are required, but the image edges are less accurately localized. Fewer linear segments will result in a smaller body of image data, and hence better image compression. Varying the permitted distance error thereby provides a progressive tuning parameter for image compression.

The magnitude of the zero crossing contrast figure is directly related to the contrast of the corresponding edge. Small magnitude samples result in low contrast edges. Low contrast edges tend to have a small visual impact for the human viewer. For lossy coding, the first samples to ignore would therefore be those which have the smallest visual impact in the reconstructed image. Small magnitude zero crossing samples can be eliminated by thresholding the zero crossing contrast figure image. Varying the threshold value provides another progressive tuning parameter for image compression.

The effect of these two tuning parameters on the recovered image quality and compression rate was examined. For each threshold value, the linear regression procedure was conducted for a range of spatial distance errors. The extracted contrast figure estimates were approximated linearly with 8 bits of dynamic range. The vector endpoints were then coding as connected chains. A coordinate pair is required to identify each endpoint. Since the endpoints are connected, only one pair of endpoints and an 8 bit contrast figure are required for each linear segment. One additional coordinate pair is required for each contour. A contrast figure of value zero was set with this coordinate pair to indicate the beginning of a new contour. Further, each plane's coordinate pairs were represented with



the appropriate number of bits. Additionally, the 16 X 16 near-dc low pass image was counted in all compression rate estimates, without coding (8 bits/pixel).

For reconstruction, the coded vector tables were redrawn on the image plane using Bresenham's algorithm [58]. The reconstruction filter was then applied, and the image planes remerged to obtain an approximation of the original image. It was found that the scaling of each image plane was affected by the number of samples which remained after the thresholding step. To obtain a more uniform balance of contrast between the reconstructed  $\nabla^2 G$  planes, a variance figure was extracted from each of the original  $\nabla^2 G$  filtered images. The variance of each reconstructed  $\nabla^2 G$  image was adjusted to more closely resemble that in the original.

The threshold value was varied over the range of 0 to  $\pm 30$  grey scale levels (out of a  $\pm 127$  grey scale level dynamic range), in steps of 5. The distance error value was varied over a range of 0 to 4 pixels for the 256 X 256 plane in steps of 1. The distance errors for each of the subsampled planes was 1/2 that of its larger neighbor, to obtain the same relative effect. Zero crossing samples which were isolated (no adjacency neighbors) after thresholding were ignored.

## 6.3 Experimental Results

Several images were evaluated using the described procedure. A matrix of images, showing reconstruction quality over the considered range of edge thresholding and edge placement error is shown in Figure 6.2. The 'girl' image was selected for this comparison, since it is well known, and represents a complex natural scene with a wide range of spatial frequencies. With no edge thresholding and no edge localization error, some image reconstruction loss was noted over the reconstruction of the images directly from the zero crossing sampled inverse filter output. This loss was attributed to the reduction in contrast figure estimate dynamic range, and the averaging of the contrast figure estimates over each linear segment. This loss was slight, however, and was therefore not considered serious. Thresholding the edges seemed to improve the image quality in some respects,



Figure 6.2: Comparison of Reconstructed Image Quality, With Varying Edge Threshold and Edge Localization: 'Girl'

Thres- hold Vector Error	0	5	10	15	20	25	30
0	5.54	3.44	2.26	1.65	1.28	1.01	0.79
1	2.46	1.77	1.19	0.92	0.70	0.57	0.45
2	1.53	1.27	0.88	0.68	0.55	0.45	0.37
3	1.15	1.09	0.76	0.60	0.49	0.41	0.34
4	0.95	0.99	0.68	0.56	0.46	0.40	0.34

Table 6.1: Image Compression Rates, Corresponding to the ‘Girl’ Image Matrix in Figure 6.1 (in bits/pixel)

since it smoothed out some of the subtle texture variations (which appear somewhat like noise in the image). As the edge threshold increased, however, some low contrast sections of visually important edges were washed out. This can be seen in the girl’s hat and chin. Still, the image remained quite recognizable, and much of the image shading was preserved. Also, the degradation with this tuning parameter was quite smooth and progressive.

The impact of increasing edge localization error was found to be great. Even with a small error, significant loss was noted in the girl’s facial features. Surprisingly, the impact of this error seemed to be reduced with increasing edge threshold level. This observation was attributed to an alignment of the edges on more high contrast lines.

For the ‘girl’ image, the compression rates corresponding to the images in Figure 6.2 are presented in Table 6.1. For the ‘girl’ with zero edge localization error and no thresh-

olding, a data rate of 5.54 bits/pixel (1.44:1) was obtained. This figure clearly does not represent good compression for the recovered image quality. The data is reduced significantly, however, with the application of some edge thresholding and/or permitting some edge localization error. Note, for example, that the most compressed image was reduced to only 0.34 bits/pixel (23:1). At that rate, the image is clearly reduced in quality but remains quite recognizable with most of its prominent edges preserved. Details of the image compression and coding for all of these images are included in the appendix.

# Chapter 7

## Conclusions and Recommendations

This thesis has considered the recovery of a sampled image from partial information. Specifically, the recovery of images from the ‘edges’ or zero crossings found in  $\nabla^2 G$  filtered versions of the image has been considered. A scheme was developed for separating an image into a family of  $\nabla^2 G$  filtered images, and recovering an approximation of the original image from this  $\nabla^2 G$  pyramid. It has been shown that the original image can be recovered from this  $\nabla^2 G$  family, with only slight loss in image crispness.

The recovery of each of the  $\nabla^2 G$  filtered image planes was then considered independently. Two theorems which suggest that  $\nabla^2 G$  filtered images might be recovered from their zero crossing locations alone were considered; Logan’s and Curtis’ theorems. It was shown that Logan’s theorem is not applicable to  $\nabla^2 G$  filtered images. It was also shown that not all  $\nabla^2 G$  filtered images meet the requirements of Curtis’ theorem, and that it is difficult in practice to determine if a given image satisfies these requirements. Difficulties introduced with sampling the zero crossing locations were also discussed.

While it was shown that not all  $\nabla^2 G$  filtered images can be uniquely characterized by their zero crossings, the practical recovery of  $\nabla^2 G$  filtered images from their sampled zero crossings locations was considered. Two iterative algorithms were examined. The images reconstructed with both of these techniques were found to be quite similar, and exhibited

significant loss in image shading and tone. Based on these theoretical and practical investigations, it seems unlikely that the full recovery of  $\nabla^2 G$  filtered images from their sampled zero crossing locations is practical for an image compression application.

The use of some additional information at each zero crossing location was therefore considered. In examining a number of natural scene images, it was found that a step edge model might be appropriate for characterizing the  $\nabla^2 G$  filtered image edges. A simple, non-iterative technique was developed for recovering a  $\nabla^2 G$  filtered image from its zero crossings, with contrast figure estimates at each zero crossing sample. These contrast figure estimates were found through a local examination of the  $\nabla^2 G$  filtered image in the vicinity of each zero crossing sample. This technique was applied to both complex natural-scene and synthetic images. The recovered image quality was found to be excellent for sparse images, and good for more complex images. In both cases, the recovered images had good recovered image shading. The technique was based on an assumption of edge independence, which was found not to be true for complex images. This assumption was found to be the most serious shortcoming of the technique, but the straight-forward computational solution was found to be formidable.

A more elegant computation solution was therefore developed, through the use of 2-D inverse and reconstruction filtering. This non-iterative approach was found to overcome the local estimation technique's shortcomings, but exhibited some ringing, which manifested itself as oriented smearing. Still, the recovered image quality remained quite high, with good recovery of image shading. The filter ringing was addressed through smoothing applied to the filter's frequency response. The oriented smearing was significantly reduced through this measure, but some loss was introduced for edges oriented at  $45^\circ$  (from upper left to lower right).

The application of these techniques for image compression was then examined. A simple approach for coding the zero crossing contours was developed, using linear regression. This regression step produced a table of vectors, which represented the sampled zero crossing contours and contrast figure estimates. Two tuning parameters were identified for the image compression; edge suppression based on contrast, and spatial approxima-

tion of the regression vectors. The direct and cross effects of these tuning parameters were examined against the compression obtained. Compressions from 8 bits/pixel to 0.34 bits/pixel (23:1) were shown, with fairly good image reconstruction. This compression was obtained without iteration, the solution of simultaneous equations, or the use of a code book.

In summary, this study has considered the recovery of  $\nabla^2 G$  filtered images from their zero crossing locations alone, and found that the use of this approach for image compression is not practical. Two new methods have been developed for recovering  $\nabla^2 G$  filtered images from their zero crossing samples with contrast figure estimates. These methods are computationally attractive because that they are non-iterative, and do not involve the solution of simultaneous equations. Further, the application of this technique to image compression has been examined. The compression obtained is comparable to current state-of-the-art methods. The recovery from partial information technique is computationally attractive, and does not involve the use of a code book. While the compression obtained did not meet the goal outlined for the underwater transmission problem, a significant potential exists for improvement in compression with the further development of this technique. Some of these potentials are discussed in the following recommendations.

## 7.1 Recommendations For Further Study

A number of potentials for further examination were found in this study. These fall primarily into two categories - recovery from partial information, and image compression.

In the area of recovery from partial information, several extensions of this work remain open. Based on the experience gained through this study, it is felt that the inverse/reconstruction filtering approach holds more potential than the local estimation technique. As discussed earlier, the discontinuities in these filter's frequency response presented problems in inverse and reconstruction filtered image smearing. One measure has been considered to address this filtering ringing, but found some ill effects resulted

from these measures. Alternative measures of addressing this ringing should be explored, which might contain the ringing without the dampening which experienced in this study.

Further, it was found that the zero crossing samples, as found by the detection algorithm used in this study, were not always well placed over the peaks of the inverse filtered image. This problem may stem from zero crossing localization error, or subtle phase shifts occurring in the inverse filtered image. The relationship between the peaks of the inverse filtered image and the zero crossing samples of the  $\nabla^2 G$  filtered image should be examined further. One alternative might be to disregard the zero crossings samples of the  $\nabla^2 G$  filtered images, and sample only those inverse filtered image values which exceed some threshold. Since the peaks of the inverse filtered image should be localized over the zero crossing contours, this step seems reasonable.

The use of pre-filtering and/or post-filtering should be examined for recovering more of the higher frequency portion of the remerged image. The  $\nabla^2 G$  filter has non-zero gain for all frequencies other than  $\omega = 0$ . This fact implies that an attenuated replica of the whole image frequency spectrum is contained in each  $\nabla^2 G$  filtered plane (less the dc component). Practically, however, the extent to which this situation can be used is constrained by the dynamic range and resolution of the sampled images. The effectiveness of using such an approach should be examined, and weighed against disadvantages which might exist (in a noisy image environment, for example). Further, the family of  $\nabla^2 G$  channels used in this study overlap each other. This implies that portions of the spectrum are being recovered in more than one  $\nabla^2 G$  channel. The use of corrective filtering may permit the recovery of the whole image with fewer  $\nabla^2 G$  channels, which should help in compression. The use of post-filtering could also be used to enhance the subjective quality of the recovered image [67].

The relationships between the  $\nabla^2 G$  filtered image planes might also be used in the recovery and/or compression steps. In this study, the recovery of each of the  $\nabla^2 G$  filtered image planes has been considered independently. These  $\nabla^2 G$  images are inter-related, since they have spectral overlap and spatial correspondence. The use of these interrelationships should be explored.



The use of other edge models might be considered. As discussed earlier, several edge models exist; these have not all been considered in this study. Further, these models are based on the mathematical properties observed in images. The use of models based on human visual perception (such as Cornsweet edges) may offer another avenue for consideration.

In the application of these techniques to image compression, many alternatives for study exist. As indicated, the application of this technique to image compression did not form a significant part of this study, and the developed coding scheme is quite simple. Alternate forms of coding the zero crossing samples should be considered. Further, the use of frame on frame coding should be examined for a sequence of tele-operation images. These images will be highly correlated frame on frame. An edge based coding technique holds much promise for capitalizing on time domain correlation. Many coding techniques based on image statistics loose the image structure in their coding step. Since an edge coding technique is highly related to image contents, this technique should lend itself well to simultaneous in the frame and frame on frame coding.

Within one frame, several techniques could be examined to reduce the data rate. The used of a foveal style presentation may provide an effective means of reducing the data rate with a low visual impact. In the human retina, high resolution is presented only on a circular fovea in the center of the field of view. Outside the fovea, the spatial resolution is progressively reduced as a function of radius. Since this primary area of interest is typically centered in the field of view, the impact of the low perimeter spatial resolution is not too great. A large fraction of the image data rate stems from the highest spatial frequency components. This measure may therefore provide a significant reduction in image data.

As discussed, it was found that the sampled zero crossing contours did not always close onto themselves (or the plane boundary). These contours could be closed, through the use of some image interpretation processing. If the objects in the scene are segmented, the edges can be traced to form closed boundaries. This procedure is typically under determined for a single image frame, without some *a priori* knowledge about the image

contents or imaging geometry. With a set of closed boundaries, techniques for coding closed boundaries could be applied. Further, the digitized contours could be smoothed [84]. These measures present another tunable parameter for image compression.

# References

- [1] John A. Adam, "Probing Beneath the Sea," *IEEE Spectrum*, vol. 22, no. 4, pp. 55-64, 1985.
- [2] D. L. Arkin et. al., "Space Applications of Automation, Robotics and Machine Intelligence Systems (ARAMIS) - Phase II," *NASA Scientific and Technical Branch Contractor Report*, Report # 3734 & 3735, Contract # NAS8-34381, 1983.
- [3] Dana H. Ballard and Christopher M. Brown, *Computer Vision*, Prentice Hall, Englewood Cliffs, 1982.
- [4] F. E. Bond and C. R. Cahn, "On Sampling the Zeros of Bandwidth Limited Signals," *IRE Trans. Inform. Theory*, pp. 110-113, September, 1959.
- [5] Peter J. Burt and Edward H. Adelson, "The Laplacian Pyramid as a Compact Image Code," *IEEE Trans. Commun.*, vol. COM-31, no. 4, pp. 532-540, 1983.
- [6] R. Chellappa and R. Bagdazain, "Fourier Coding of Image Boundries," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-6, no. 1, pp. 102-105, 1984.
- [7] C. H. Chen, "Laplacian Pyramid Image Data Compression," in *Proc. IEEE ICASSP*, pp. 737-739, May 1987.
- [8] J. J. Clark and P. D. Lawrence, "A Hierarchical Image Analysis System Based Upon Oriented Zero Crossings of Bandpassed Images," in *Multiresolution Image Processing and Analysis*, pp. 148-168, A. Rosenfield, ed., Springer-Verlag, Berlin, 1984.
- [9] Jim Collins, Royal Roads Military College, *personal communication*, 1984.
- [10] Susan R. Curtis and Alan V. Oppenheim, "Signal Reconstruction From One Bit of Fourier Transform Phase," in *Proc. IEEE ICASSP*, pp. 12A.5.1-12A.5.4, May 1984.
- [11] Susan R. Curtis, Alan V. Oppenheim, and Jae S. Lim, "Signal Reconstruction from Fourier Transform Sign Information," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. ASSP-33, no. 3, pp. 643-637, 1985.
- [12] Susan R. Curtis, Alan V. Oppenheim, and Jae S. Lim, "Reconstruction of Two-Dimensional Signals From Threshold Crossings," in *Proc. IEEE ICASSP*, pp. 28.2.1-28.2.4, May 1985.

- [13] Susan R. Curtis, "Reconstruction of Multidimensional Signals from Zero Crossings," *Ph.D. Thesis*, MIT, June, 1985.
- [14] Susan R. Curtis, Shlomo Shitz, and Alan V. Oppenheim, "Reconstruction of Non-periodic Two-Dimensional Signals from Zero Crossings," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. ASSP-35, no. 6, pp. 890-893, 1987.
- [15] Dan E. Dudgeon and Russell M. Mersereau, *Multidimensional Digital Signal Processing*, Prentice-Hall, Englewood Cliffs, 1984.
- [16] Jimmie D. Eggerton and Mandyam D. Srinath, "A Visually Weighted Quantization Scheme for Image Bandwidth Compression at Low Data Rates," *IEEE Trans. Commun.*, vol. COM-34, no. 8, pp. 840-847, 1986.
- [17] Staffan Ericsson, "Motion-Compensated Hybrid Coding at 50 KB/S," in *Proc. IEEE ICASSP*, pp. 10.8.1-10.8.4, May 1985.
- [18] Donald Norman Graham, "Image Transmission by Two-Dimensional Contour Coding," *Proc. of the IEEE*, vol. 55, no. 3, pp. 336-346, 1967.
- [19] Robert M. Gray, "Vector Quantization," *IEEE ASSP Magazine*, pp. 4-29, April, 1984.
- [20] W. E. L. Grimson and E. C. Hildreth, "Comments on 'Digital Step Edges from Zero Crossings of Second Directional Derivatives'," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-7, no.1, pp. 121-129, 1985.
- [21] William Eric Leifur Grimson, *From Images to Surfaces: A Computational Study of the Human Early Vision System*, MIT Press, Cambridge, 1981.
- [22] Rafael C. Gonzalez and Paul Wintz, *Digital Image Processing*, Addison-Wesley, Reading, 1983.
- [23] Robert M. Haralick, "Digital Step Edges from Zero Crossing of Second Derivatives," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-6, no. 1, pp. 58-68, 1984.
- [24] Barry G. Haskell and Raymond Steele, "Audio and Video Bit-Rate Reduction," *Proceedings of the IEEE*, vol. 69, no. 2, pp. 252-262, 1981.
- [25] Monson H. Hayes, Jae S. Lim, and Alan V. Oppenheim, "Signal Reconstruction from Phase or Magnitude," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. ASSP-28, no. 6, pp. 672-680, 1980.
- [26] Monson H. Hayes, "The Reconstruction of a Multidimensional Sequence from the Phase or Magnitude of Its Fourier Transform," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. ASSP-30, no. 2, pp. 140-154, 1982.
- [27] Monson H. Hayes and James H. McClellan, "Reducible Polynomials in More Than One Variable," *Proceedings of the IEEE*, vol. 70, no. 2, pp. 197-198, 1982.
- [28] I. E. Holliday, K. H. Ruddock, and P. Skinner, "Spatial Filtering of Retinal Images by the Human Visual System and its Consequences for Visual Thresholds," *SPIE*, vol. 467, pp. 2-6, 1983.

- [29] Berthold Klaus Paul Horn, *Robot Vision*, MIT Press, Cambridge, 1986.
- [30] B. K. P. Horn and E. J. Weldon, Jr., "Filtering Closed Curves," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-8, no. 5, pp. 665-668, 1986.
- [31] T. S. Huang, "Two-Dimensional Windows," *IEEE Trans. Audio Electroacoust.*, vol. AU-21, pp. 88-89, March 1972.
- [32] David H. Hubel and Torsen N. Wiesel, "Brain Mechanisms of Vision," *Scientific American*, vol. 241, pp. 150-162, Sept., 1979.
- [33] Andres Huertas and Gerard Medioni, "Detection of Intensity Changes with Sub-pixel Accuracy Using Laplacian-Gaussian Masks," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-8, no. 5, pp. 651-664, 1986.
- [34] Robert A. Hummel, "Representations based on Zero Crossings in Scale Space," *Proc. IEEE Comp. Vision & Patt. Recog. Conf.*, pp. 204-209, 1986.
- [35] Anil K. Jain, "Image Data Compression: A Review," *Proceedings of the IEEE*, vol. 69- no. 3, pp. 349-389, 1981.
- [36] C. N. Judice and D. LeGall, "Telematic Services and Terminals: Are We Ready?," *IEEE Communications Magazine*, vol. 25, no. 7, pp. 19-29, 1987.
- [37] Haruo Kato and Tomozo Furukawa, "Two-Dimensional Type-Preserving Circular Windows," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. ASSP-29, no. 4, pp. 926-928, 1981.
- [38] Murat Kunt, Athanassios Ikonomopoulos, and Michel Kocher, "Second-Generation Image-Coding Techniques," *Proceedings of the IEEE*, vol. 73, no. 4, pp. 549-574, 1985.
- [39] Peter D. Lawrence, University of British Columbia, *personal communication*, 1987.
- [40] Yvan G. Leclerc and Steven W. Zucker, "The Local Structure of Image Discontinuities in One Dimension," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-9, no. 3, pp. 341-355, 1987.
- [41] Philippe Letellier, Morton Nadler, and Jean-Francois Abramatic, "The Telesign Project," *Proceedings of the IEEE*, vol. 73, no. 4, pp. 813-827, 1985.
- [42] B. F. Logan, "Information in the Zero Crossings of Bandpass Signals," *Bell Sys. Tech. J.*, vol. 56, no. 4, pp. 487-510, 1977.
- [43] B. F. Logan, "Signals Designed for Recovery After Clipping - I. Localization of Infinite Products," *AT&T Bell Lab. Tech. Journal*, vol. 63, no. 2, pp. 261-285, 1984.
- [44] B. F. Logan, "Signals Designed for Recovery After Clipping - II. Fourier Transform Theory of Recovery," *AT&T Bell Lab. Tech. Journal*, vol. 63, no. 2, pp. 287-306, 1984.

- [45] Petros A. Maragos, Ronald W. Schafer, and Russell W. Mersereau, "Two-Dimensional Linear Prediction and Its Application to Adaptive Predictive Coding of Images," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. ASSP-32, no. 6, pp. 1213-1229, 1984.
- [46] Michael Maragoudakis and Jerry D. Gibson, "Experiments on Video Teleconferencing Algorithms at 56 Kilobits/Sec," in *Proc. IEEE ICASSP*, pp. 1071-1074, May 1987.
- [47] Jacques de Mare, "Reconstruction of a Stationary Gaussian Process from its Sign-Changes," *J. Appl. Prob.*, vol. 14, pp. 38-57, 1977.
- [48] D. Marr, S. Ullman, and T. Poggio, "Bandpass Channels, Zero-crossings, and Early Visual Information Processing," *J. Opt. Soc. Am.*, vol. 69, no. 6, pp. 914-916, 1979.
- [49] D. Marr and E. Hildreth, "Theory of Edge Detection," *Proc. R. Soc. Lond. B*, vol. 207, pp. 187-217, 1980.
- [50] David Marr, *Vision*, Freeman, San Francisco, 1982.
- [51] James H. McClellan, "The Design of Two-Dimensional Digital Filters by Transformation," *Proc. 7th Annu. Princeton Conf. on Inform. Sci. and Syst.*, Mar. 1973, pp. 247-251.
- [52] Wolfgang F. G. Mecklenbrauker and Russell M. Mersereau, "McClellan Transformations for Two-Dimensional Digital Filtering: II - Implementation," *IEEE Trans. Circuits and Systems*, vol. CAS-23, no. 7, pp. 414-422, 1976.
- [53] A. Zvi Meiri and Eitan Yudilevich, "A Pinned Sine Transform Image Coder," *IEEE Trans Commun.*, vol. COM-29, no. 12, pp. 1728-1740, 1981.
- [54] Russell M. Mersereau, Wolfgang F. G. Mecklenbrauker, and Thomas F. Quatieri, "McClellan Transformations for Two-Dimensional Digital Filtering: I - Design," *IEEE Trans. Circuits and Systems*, vol. CAS-23, no. 7, pp. 405-414, 1976.
- [55] Hans Georg Musmann and Jurgen Klie, "TV-Transmission Using a 64 KBit/S Transmission Rate," *Proc. IEEE Int. Commun. Conf.*, pp. 23.3.1-23.3.5, 1979.
- [56] Hans Georg Musmann, Peter Pirsch, and Hans-Joachim Grallert, "Advances in Picture Coding," *Proceedings of the IEEE*, vol. 73, no. 4, pp. 523-548, 1985.
- [57] Arun N. Netravali and John O. Limb, "Picture Coding: A Review," *Proceedings of the IEEE*, vol. 68, no. 3, pp. 366-406, 1980.
- [58] William M. Newmann and Robert F. Sproull, *Principles of Interactive Computer Graphics*, McGraw-Hill Book Company, New York, 1979.
- [59] Ramin A. Nobakht and Sarah A. Rajala, "An Image Coding Technique Using A Human Visual System Model and Image Analysis Technique," in *Proc. IEEE ICASSP*, pp. 1358-1361, May 1987.
- [60] Alan V. Oppenheim and Jae S. Lim, "The Importance of Phase in Signals," *Proceedings of the IEEE*, vol. 69, no. 5, pp. 529-541, 1981.

- [61] Alan V. Oppenheim, Monson H. Hayes, and Jae S. Lim, "Iterative Procedures for Signal Reconstruction from Phase," *Proc. Int. Optical Computing Conf.*, SPIE vol. 231, pp. 121-129, 1980.
- [62] A. V. Oppenheim, J. S. Lim, and S. R. Curtis, "Signal Synthesis and Reconstruction from Partial Fourier-Domain Information," *J. Opt. Soc. Am.*, vol. 73, no. 11, pp. 1413-1420, 1983.
- [63] Don E. Pearson and John A. Robinson, "Visual Communication at Very Low Data Rates," *Proceedings of the IEEE*, vol. 73, no. 4, pp. 795-812, 1985.
- [64] Didier Perny, "Compression of Underwater Images for their Transmission on a Low Bit Rate Acoustic Channel," in *Proc. IEEE ICASSP*, pp. 29.1.1-29.1.4, May 1984.
- [65] Lawrence R. Rabiner, James H. McClellan, and Thomas W. Parks, "FIR Digital Filter Design Techniques Using Weighted Chebyshev Approximation," *Proceedings of the IEEE*, vol. 63, pp. 595-610, 1975.
- [66] Sarah A. Rajala, Mehmet R. Civanlar, Wonrae M. Lee, "A Second Generation Image Coding Technique Using Human Visual System Based Segmentation," in *Proc. IEEE ICASSP*, pp. 1362-1365, May 1987.
- [67] Bhaskar Ramamurthi and Allen Gersho, "Edge-Oriented Spatial Filtering of Images with Application to Post-Processing of Vector Quantized Images," in *Proc. IEEE ICASSP*, pp. 48.10.1-48.10.4, May 1984.
- [68] Vivek Randive and Thomas B. Sheridan, "Video Framerate, Resolution, and Greyscale Tradeoffs for Undersea Telemanipulation," *Proc. Int. Conf. Cybern. Soc.*, pp. 77-88, Denver, 1979.
- [69] J. A. Reimer and P. D. Lawrence, "Characterizing  $\nabla^2 G$  Filtered Images By Their Zero Crossings," in *Proc. IEEE ICASSP*, pp. 253-256, May 1987.
- [70] Aristides A. G. Requicha, "The Zeros of Entire Functions: Theory and Engineering Applications," *Proceedings of the IEEE*, vol. 68, no. 3, pp. 308-327, 1980.
- [71] Azriel Rosenfeld and Avinash C. Kak, *Digital Picture Processing, Volume 1*, Academic Press, Orlando, 1982.
- [72] Robert Ross, University of British Columbia, *personal communication*, 1987.
- [73] Doron Rotem and Yehoshua Y. Zeevi, "Image Reconstruction from Zero Crossings," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. ASSP-34, no. 5, pp. 1269-1277, 1986.
- [74] S. Sabri, K. Cuffing, and B. Prasada, "Coding of Video Signals at 50 Kb/S Using Motion Compensation Techniques," *1983 IEEE Military Commun. Conf.*, pp. 809-816.
- [75] Jorge L. C. Sanz and Thomas S. Huang, "On the Stability and Sensitivity of Multidimensional Signal Reconstruction from Fourier Transform Magnitude," in *Proc. IEEE ICASSP*, pp. 1065-1068, May 1985.

- [76] Jorge L. C. Sanz and Thomas S. Huang, "Polynomial System of Equations and Its Application to the Study of the Effect of Noise on Multidimensional Fourier Transform Phase Retrieval from Magnitude," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. ASSP-33, no. 4, pp. 997-1004, 1985.
- [77] Jorge L. C. Sanz, Thomas S. Huang, and Fernando Cukiernan, "Stability of Unique Fourier-transform Phase Reconstruction," *J. Opt. Soc. Am.*, vol. 73, no. 11, pp. 1442-1445, 1983.
- [78] Jorge L. C. Sanz and Thomas T. Huang, "Mathematics and Computer Vision: Image Representation by Zero Crossings," *in press*.
- [79] Jorge L. C. Sanz and Thomas S. Huang, "Unique Reconstruction of a Band-limited Multidimensional Signal from its Phase or Magnitude," *J. Opt. Soc. Am.*, vol. 73, no. 11, pp. 1446-1450, 1983.
- [80] Jorge L. C. Sanz, "On the Reconstruction of Band-Limited Multidimensional Signals from Algebraic Sampling Contours," *Proceedings of the IEEE*, vol. 73, no. 8, pp. 1334-1336, 1985.
- [81] M. S. Scivier and M. A. Fiddy, "Phase Ambiguities and the Zeros of Multidimensional Band-limited Functions," *J. Opt. Soc. Am. A*, vol. 2, no. 5, pp. 693-697, 1985.
- [82] W. F. Schrieber, T. S. Huang, and O. J. Tretiak, "Contour Coding of Images", appearing in *Picture Bandwidth Reduction*, ed. T. S. Huang and O. J. Tretiak, Gordon and Breach Science Publishers, New York, 1972.
- [83] W. F. Schrieber, C. F. Knapp, and N. D. Kay, "Synthetic Highs - An Experimental TV Bandwidth Reduction System," *Journal of the SMPTE*, vol. 68, pp. 525-537, 1957.
- [84] Behzad Shahraray and David J. Anderson, "Optimal Smoothing of Digitized Contours," *1986 IEEE Int. Conf. Comp. Vision Pattern Recog.*, pp. 210-218, June, 1986.
- [85] Shlomo Shitz and Yehoshua Y. Zeevi, "On the Duality of Time and Frequency Domain Signal Reconstruction from Partial Information," *IEEE Trans. Acoust. Speech Signal Process.*, vol. ASSP-33, no. 6, pp. 1486-1498, 1985.
- [86] Gilbert Strang, *Linear Algebra and its Applications*, Academic Press, New York, 1980.
- [87] V. Torre and T. Poggio, "On Edge Detection," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-8, no. 2, pp. 147-163, 1986.
- [88] Anh Tran, Kwun-Min Liu, Kou-Hu Tzou, and Eileen B. Vogel, "An Efficient Pyramid Image Coding System," in *Proc IEEE ICASSP*, pp. 744-747, May 1987.
- [89] D. Jacques Vaissey and Allen Gersho, "Variable Block-Size Image Coding," in *Proc. IEEE ICASSP*, pp. 1051-1054, May 1987.



- [90] Patrick L. Van Hove, Monson H. Hayes, Jae S. Lim, and Alan V. Oppenheim, "Signal Reconstruction from Signed Fourier Transform Magnitude," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. ASSP-31, no. 5, pp. 1286-1293, 1983.
- [91] Herbert B. Voelcker, "Toward a Unified Theory of Modulation - Part II: Zero Manipulation," *Proceedings of the IEEE*, vol. 54, no. 5, pp. 735-755, 1966.
- [92] Herbert B. Voelcker and Aristides A. G. Requicha, "Band-Limited Random-Real-Zero Signals," *IEEE Trans. Commun.*, vol. 934, pp. 933-936, August, 1973.
- [93] Herbert B. Voelcker and Aristides A. G. Requicha, "Clipping and Signal Determinism: Two Algorithms Requiring Validation," *IEEE Trans. Commun.*, pp. 738-744, June, 1973.
- [94] Robert Wallis and William K. Pratt, "Video Tele-Conferencing at 9600 Baud," *Proc. IEEE Int. Commun. Conf.*, pp. 22.2.1-22.2.3, 1981.
- [95] Allan Weber, "University of Southern California Image Data Base," *Image Processing Institute USCIP Report 1070*, University of Southern California, March 1983.
- [96] H. R. Wilson, "Psychophysical Evidence for Spatial Channels," from *Physical and Biological Processing of Images*, pp. 88-99, O. J. Braddick and A. C. Sleight, ed., Springer-Verlag, Berlin, 1983.
- [97] John W. Woods and Sean D. O'Neil, "Subband Coding of Images," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. ASSP-34, no. 5, pp. 1278-1288, 1986.
- [98] Tian-Hu Yu and Sanjit K. Mitra, "A New Two-Dimensional Window," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. ASSP-33, no. 4, pp. 1058-1061, 1985.
- [99] A. L. Yuille and T. Poggio, "Fingerprint Theorems for Zero Crossings," *J. Opt. Soc. Am. A*, vol. 2, no. 5, pp. 683-692, 1985.
- [100] Alan Y. Yuille and Tomaso A. Poggio, "Scaling Theorems for Zero Crossings," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-8, no. 1, pp. 15-25, 1986.
- [101] G. Stephan Zabele and Jack Koplowitz, "Fourier Encoding of Closed Planar Boundaries," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-7, no. 1, pp. 98-102, 1985.
- [102] Avidesh Zakhori and David Izraelevitz, "A Note on the Sampling of Zero Crossings of Two-Dimensional Signals," *Proceedings of the IEEE*, vol-74, no. 9, pp. 1285-1287, 1986.

# Appendix

## 1-D LOW PASS FILTER DESIGN DETAILS

FINITE IMPULSE RESPONSE (FIR)  
LINEAR PHASE DIGITAL FILTER DESIGN  
REMEZ EXCHANGE ALGORITHM

BANDPASS FILTER

FILTER LENGTH = 49

\*\*\*\*\* IMPULSE RESPONSE \*\*\*\*\*

H( 1) = -0.46052222E-03 = H( 49)  
H( 2) = 0.62446175E-02 = H( 48)  
H( 3) = 0.14550016E-02 = H( 47)  
H( 4) = -0.45280769E-02 = H( 46)  
H( 5) = -0.31475630E-02 = H( 45)  
H( 6) = 0.53713918E-02 = H( 44)  
H( 7) = 0.59338543E-02 = H( 43)  
H( 8) = -0.55162548E-02 = H( 42)  
H( 9) = -0.96231885E-02 = H( 41)  
H(10) = 0.45008804E-02 = H( 40)  
H(11) = 0.14104469E-01 = H( 39)  
H(12) = -0.17834451E-02 = H( 38)  
H(13) = -0.19167205E-01 = H( 37)  
H(14) = -0.33179231E-02 = H( 36)  
H(15) = 0.24491096E-01 = H( 35)  
H(16) = 0.11761160E-01 = H( 34)  
H(17) = -0.29677479E-01 = H( 33)  
H(18) = -0.25277123E-01 = H( 32)  
H(19) = 0.34297951E-01 = H( 31)  
H(20) = 0.48166532E-01 = H( 30)  
H(21) = -0.37945330E-01 = H( 29)  
H(22) = -0.96351117E-01 = H( 28)  
H(23) = 0.40283095E-01 = H( 27)  
H(24) = 0.31498048E+00 = H( 26)  
H(25) = 0.45891163E+00 = H( 25)

	BAND 1	BAND 2
LOWER BAND EDGE	0.0000000	0.2513000
UPPER BAND EDGE	0.2076000	0.5000000
DESIRED VALUE	1.0000000	0.0000000
WEIGHTING	1.0000000	1.0000000
DEVIATION	0.0085022	0.0085022
DEVIATION IN DB	0.0735374	-41.4093704

EXTREMAL FREQUENCIES--MAXIMA OF THE ERROR CURVE

0.0000000	0.0212500	0.0425000	0.0637500	0.0850000
0.1062499	0.1274999	0.1487499	0.1687499	0.1862499
0.2012499	0.2076000	0.2513000	0.2575500	0.2725500
0.2900500	0.3100500	0.3312999	0.3512999	0.3725499
0.3937999	0.4150499	0.4362998	0.4575498	0.4787998
0.5000000				

## CHANNEL WEIGHTS USED IN REMERGING

256 X 256 plane: 1.00  
 128 X 128 plane: 0.30  
 64 X 64 plane: 0.59  
 32 X 32 plane: 0.55

## Image Compression Details

Note: Cited threshold values were used for all 4 image planes. This value was used to threshold the contrast figure estimates, and is cited in grey scale absolute value units (over a dynamic range of 256). The cited vector error (approx) applies to the 256 X 256 plane. The error applied to each subsequent plane is 1/2 that of its higher frequency neighbour (eg. 1.0 got 256 X 256, 0.5 for 128 X 128, 0.25 for 64 X 64, and 0.125 for 32 X 32). This value is cited in pixel units. The 16 X 16 near-dc low pass plane was not coded, and therefore contributed 2048 bits (0.03 bits/pixel) to the total bits/pixel value. This value is included in each cited total bits/pixel value.

	Thres=0, Approx=0					Thres=0, Approx=1.0				
Plane	256	128	64	32	All	256	128	64	32	All
# zc pts	21178	4788	1247	395		21178	4788	1247	395	
# contours drawn	1356	269	59	20		1116	228	59	20	
# segments drawn	12039	2643	698	226		3031	1341	698	226	
# joins drawn	432	72	15	5		599	110	7	5	
# isolated zc pts	413	61	8	6		205	25	8	6	
# zero value pts	329	46	7	4		51	16	7	4	
# bits	285896	57802	13668	3923		105944	35824	13508	3932	
# bits/pixel	4.36	0.88	0.20	0.05	5.54	1.61	0.54	0.20	0.05	2.46

	Thres=0, Approx=2.0					Thres=0, Approx=3.0				
Plane	256	128	64	32	All	256	128	64	32	All
# zc pts	21178	4788	1247	395		21178	4788	1247	395	
# contours drawn	934	214	51	20		848	204	47	21	
# segments drawn	1685	618	366	226		1239	454	191	211	
# joins drawn	429	101	35	1		253	103	15	7	
# isolated zc pts	150	24	4	6		137	20	2	1	
# zero value pts	25	9	0	4		24	9	0	3	
# bits	65648	19830	8948	3860		48776	15474	4988	3986	
# bits/pixel	1.00	0.30	0.13	0.05	1.53	0.74	0.23	0.07	0.06	1.14

	Thres=0, Approx=4.0					Thres=5, Approx=0				
Plane	256	128	64	32	All	256	128	64	32	All
# zc pts	21178	4788	1247	395		10727	3323	1034	339	
# contours drawn	769	183	47	18		1189	315	79	26	
# segments drawn	1033	343	156	130		6015	1862	580	193	
# joins drawn	172	80	27	11		119	32	13	3	
# isolated zc pts	119	18	3	0		726	127	16	7	
# zero value pts	19	5	0	0		47	13	4	1	
# bits	40248	12526	4568	2852		162440	45064	12408	3680	
# bits/pixel	0.61	0.19	0.06	0.04	0.94	2.47	0.68	0.18	0.05	3.44
	Thres=5, Approx=1.0					Thres=5, Approx=2.0				
Plane	256	128	64	32	All	256	128	64	32	All
# zc pts	10727	3323	1034	339		10727	3323	1034	339	
# contours drawn	1047	289	79	26		1004	289	75	26	
# segments drawn	1681	958	580	193		1215	481	296	193	
# joins drawn	197	62	5	3		116	65	21	2	
# isolated zc pts	600	100	16	7		567	95	13	7	
# zero value pts	8	7	4	1		5	2	0	1	
# bits	69272	28586	12248	3680		51968	18070	7828	3662	
# bits/pixel	1.05	0.43	0.18	0.05	1.76	0.79	0.27	0.11	0.05	1.27
	Thres=5, Approx=3.0					Thres=5, Approx=4.0				
Plane	256	128	64	32	All	256	128	64	32	All
# zc pts	10727	3323	1034	339		10727	3323	1034	339	
# contours drawn	975	276	66	30		957	271	65	29	
# segments drawn	1088	386	170	172		1035	337	138	113	
# joins drawn	64	55	18	9		52	33	22	7	
# isolated zc pts	564	88	14	3		558	87	13	1	
# zero value pts	2	0	1	0		1	0	0	0	
# bits	45288	15416	5028	3734		42528	12932	4428	2672	
# bits/pixel	0.69	0.23	0.07	0.05	1.09	0.64	0.19	0.06	0.04	0.98

	Thres=10, Approx=0					Thres=10, Approx=1.0				
Plane	256	128	64	32	All	256	128	64	32	All
# zc pts	6683	2425	886	292		6683	2425	886	292	
# contours drawn	765	255	73	25		688	241	73	25	
# segments drawn	3621	1269	468	160		992	643	468	160	
# joins drawn	57	15	9	2		113	28	9	2	
# isolated zc pts	433	95	27	14		363	78	27	14	
# zero value pts	20	6	2	1		0	2	2	1	
# bits	100448	31820	10368	3140		42724	19984	10368	3140	
# bits/pixel	1.53	0.48	0.15	0.04	2.25	0.65	0.30	0.15	0.04	1.19
	Thres=10, Approx=2.0					Thres=10, Approx=3.0				
Plane	256	128	64	32	All	256	128	64	32	All
# zc pts	6683	2425	886	292		6683	2425	886	292	
# contours drawn	651	236	72	25		632	226	70	25	
# segments drawn	754	348	247	160		687	296	145	134	
# joins drawn	59	36	17	2		35	34	11	7	
# isolated zc pts	351	73	18	14		351	73	18	9	
# zero value pts	0	0	0	1		1	0	0	0	
# bits	32128	13560	6708	3140		28528	11672	4448	2960	
# bits/pixel	0.49	0.20	0.10	0.04	0.87	0.43	0.17	0.06	0.04	0.75
	Thres=10, Approx=4.0					Thres=15, Approx=0				
Plane	256	128	64	32	All	256	128	64	32	All
# zc pts	6683	2425	886	292		4745	1872	759	246	
# contours drawn	662	227	68	25		578	209	87	28	
# segments drawn	650	267	121	83		2498	891	382	124	
# joins drawn	23	22	14	4		28	5	4	0	
# isolated zc pts	346	74	16	7		331	85	28	15	
# zero value pts	0	1	0	0		11	2	2	0	
# bits	26608	10088	4008	2024		71504	23306	8908	2654	
# bits/pixel	0.40	0.15	0.06	0.03	0.68	1.09	0.35	0.13	0.04	1.65

	Thres=15, Approx=1.0					Thres=15, Approx=2.0				
Plane	256	128	64	32	All	256	128	64	32	All
# zc pts	4745	1872	759	246		4745	1872	759	246	
# contours drawn	535	206	87	28		513	198	83	28	
# segments drawn	714	477	382	124		574	254	206	124	
# joins drawn	70	18	10	0		38	13	11	0	
# isolated zc pts	280	69	28	15		279	68	19	15	
# zero value pts	0	1	2	0		0	1	0	0	
# bits	30928	15320	9028	2654		24408	9584	6008	2654	
# bits/pixel	0.47	0.23	0.13	0.04	0.91	0.37	0.14	0.09	0.04	0.68
	Thres=15, Approx=3.0					Thres=15, Approx=4.0				
Plane	256	128	64	32	All	256	128	64	32	All
# zc pts	4745	1872	759	246		4745	1872	759	246	
# contours drawn	504	197	78	28		498	199	75	27	
# segments drawn	529	230	125	105		511	222	109	66	
# joins drawn	21	15	11	4		8	14	12	5	
# isolated zc pts	271	66	19	11		270	66	18	12	
# zero value pts	0	1	0	0		0	1	0	0	
# bits	22008	8720	4268	2456		20728	8360	3848	1772	
# bits/pixel	0.33	0.13	0.06	0.03	0.60	0.31	0.12	0.05	0.02	0.56
	Thres=20, Approx=0					Thres=20, Approx=1.0				
Plane	256	128	64	32	All	256	128	64	32	All
# zc pts	3464	1472	632	214		3464	1472	632	214	
# contours drawn	465	179	77	28		434	173	77	28	
# segments drawn	1809	673	303	93		544	337	303	93	
# joins drawn	20	6	4	0		40	14	9	0	
# isolated zc pts	266	86	40	21		236	78	40	21	
# zero value pts	8	1	1	0		0	0	1	0	
# bits	53720	18576	7308	2132		23168	11514	7408	2132	
# bits/pixel	0.81	0.28	0.11	0.03	1.27	0.35	0.17	0.11	0.03	0.70

	Thres=20, Approx=2.0					Thres=20, Approx=3.0				
Plane	256	128	64	32	All	256	128	64	32	All
# zc pts	3564	1472	632	214		3564	1472	632	214	
# contours drawn	420	169	73	28		412	169	72	27	
# segments drawn	457	202	165	93		429	189	103	80	
# joins drawn	24	11	11	2		17	6	9	3	
# isolated zc pts	230	80	34	21		229	77	34	17	
# zero value pts	0	0	0	0		0	0	0	0	
# bits	19048	7676	4988	2162		17728	6992	3592	1988	
# bits/pixel	0.29	0.11	0.07	0.03	0.54	0.27	0.10	0.05	0.03	0.49
	Thres=20, Approx=4.0					Thres=25, Approx=0				
Plane	256	128	64	32	All	256	128	64	32	All
# zc pts	3564	1472	632	214		2738	1181	509	174	
# contours drawn	409	169	69	27		374	170	66	26	
# segments drawn	419	181	89	55		1371	531	233	73	
# joins drawn	5	6	12	1		8	6	3	1	
# isolated zc pts	228	77	32	17		222	83	32	16	
# zero value pts	0	0	0	0		2	2	0	0	
# bits	16888	6740	3240	1502		41264	15342	5928	1772	
# bits/pixel	0.25	0.10	0.04	0.02	0.46	0.62	0.23	0.09	0.02	1.01
	Thres=25, Approx=1.0					Thres=25, Approx=2.0				
Plane	256	128	64	32	All	256	128	64	32	All
# zc pts	2732	1181	509	174		2732	1181	509	174	
# contours drawn	352	166	66	26		346	164	62	26	
# segments drawn	416	272	233	73		364	185	126	73	
# joins drawn	22	12	9	1		14	7	4	0	
# isolated zc pts	205	76	32	16		200	77	28	16	
# zero value pts	0	0	0	0		0	0	0	0	
# bits	17368	9908	6048	1772		15088	6920	3848	1754	
# bits/pixel	0.26	0.15	0.09	0.02	0.56	0.23	0.10	0.05	0.02	0.45



	Thres=25, Approx=3.0					Thres=25, Approx=4.0				
Plane	256	128	64	32	All	256	128	64	32	All
# zc pts	2738	1181	509	174		2738	1181	509	174	
# contours drawn	342	163	61	25		342	163	60	25	
# segments drawn	349	175	81	61		347	170	72	43	
# joins drawn	4	3	5	3		5	3	10	4	
# isolated zc pts	199	74	28	14		199	74	27	14	
# zero value pts	0	0	0	0		0	0	0	0	
# bits	14088	6416	2760	1610		14048	6236	2632	1304	
# bits/pixel	0.21	0.09	0.04	0.02	0.41	0.21	0.09	0.04	0.01	0.40
	Thres=30, Approx=0					Thres=30, Approx=1.0				
Plane	256	128	64	32	All	256	128	64	32	All
# zc pts	2102	885	435	143		2102	885	435	143	
# contours drawn	326	121	56	23		311	118	56	23	
# segments drawn	1029	388	180	54		344	191	180	54	
# joins drawn	3	2	4	1		11	10	7	1	
# isolated zc pts	195	88	37	17		177	85	37	17	
# zero value pts	0	0	0	0		0	0	0	0	
# bits	32168	11162	4768	1394		14208	7026	4828	1394	
# bits/pixel	0.49	0.17	0.07	0.02	0.78	0.21	0.10	0.07	0.02	0.45
	Thres=30, Approx=2.0					Thres=30, Approx=3.0				
Plane	256	128	64	32	All	256	128	64	32	All
# zc pts	2102	885	435	143		2102	885	435	143	
# contours drawn	308	118	55	23		306	116	55	22	
# segments drawn	317	132	96	54		311	124	68	43	
# joins drawn	7	3	0	0		3	1	1	3	
# isolated zc pts	175	86	33	17		175	85	32	15	
# zero value pts	0	0	0	0		0	0	0	0	
# bits	12968	4868	3028	1376		12568	4508	2216	1232	
# bits/pixel	0.19	0.07	0.04	0.02	0.37	0.19	0.06	0.03	0.01	0.34

	Thres=30, Approx=4.0				
Plane	256	128	64	32	All
# zc pts	2102	885	435	143	
# contours drawn	306	116	55	22	
# segments drawn	310	122	63	35	
# joins drawn	4	2	5	3	
# isolated zc pts	175	85	31	15	
# zero value pts	0	0	0	0	
# bits	12568	4472	2184	1072	
# bits/pixel	0.19	0.06	0.03	0.01	0.34