VLSI Image Processing

Edited by

R. J. Offen

VLSI Image Processing

Edited by

R. J. Offen

Manager of the Computer Systems Research Laboratory, GEC Research Laboratories, Hirst Research Centre, Wembley, UK

COLLINS

8 Grafton Street, London W1

Collins Professional and Technical Books
William Collins Sons & Co. Ltd
8 Grafton Street, London W1X 3LA

First published in Great Britain by
Collins Professional and Technical Books 1985

Copyright © R. J. Offen 1985

British Library Cataloguing in Publication Data
VLSI image processing.
1. Image processing 2. Integrated circuits—
Very large scale integration
I. Offen, R.J.
621.38'0414 TA1632

ISBN 0–00–383055–1

Typeset by Hope Services, Abingdon, Oxon
Printed and bound in Great Britain by
Mackays of Chatham, Kent

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted, in any form, or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior permission of the publishers.

Contributors' Biographies

B. Arambepola is Leader of the VLSI Signal Processing Group, Very High Performance Integrated Circuit Design Laboratory, GEC Research Laboratories, Hirst Research Centre, Wembley, UK.

S. R. Brooks is Manager of the Marconi Space and Defence Systems Research Laboratory, GEC Research Laboratories, Marconi Research Centre, Chelmsford, UK. He is involved in microwave surveillance and remote sensing, from airborne and spaceborne platforms, including synthetic aperture radar.

H. Buxton holds a lectureship in Information Technology at Queen Mary College, London, in the Department of Computer Science and Statistics. With a background in cognitive psychology, BSc (Bristol) and PhD (Cambridge), she has been researching temporal aspects of visual perception using an Artificial Intelligence approach; that is, developing a computational theory of both human and machine systems. This research involves parallel computer implementations of processing algorithms to interpret image sequences.

P. V. Collins was, until recently, Manager of the Marconi Systems Division, Computer Systems Research Laboratory, GEC Research Laboratories, Hirst Research Centre, Wembley, UK. He is currently with Marconi Underwater Systems Ltd.

M. M. McCabe is Manager of the Computer Science Research Division and Leader of the Computer Vision Group, Computer Systems Research Laboratory, GEC Research Laboratories, Hirst Research Centre, Wembley, UK.

R. J. Offen was previously Senior Lecturer in Physics at the University of Otago, Dunedin, New Zealand. He is currently Manager of the Computer Systems Research Laboratory, GEC Research Laboratories, Hirst Research Centre, Wembley, UK.

I. N. Parker is Manager of the Very High Performance Integrated Circuit Design Laboratory, GEC Research Laboratories, Hirst Research Centre, Wembley, UK.

P. N. Scharbach is Leader of the Formal Methods Group, Computer Systems Research Laboratory, GEC Research Laboratories, Hirst Research Centre, Wembley, UK. The research interests of the Group include the development and application of formal techniques for specifying and verifying highly concurrent systems.

J. S. Wiejak. After obtaining a BA in Mathematics from Christ's College, Cambridge, followed by a PhD in Physics from Imperial College, he joined the Computer Systems Research Laboratory, GEC Research Laboratories, Hirst Research Centre, Wembley, UK. He is currently working at Meta Machines, Oxford, on sensor-based robotics systems.

此为试读，需要完整PDF请访问：www.ertongbook.com

Preface

Image processing is a large, complex and diverse area of endeavour, and is becoming more so. As practised at present, image processing relies to a great extent on two major, effectively distinct and self-contained domains of activity, computational algorithms and processor architectures – especially those architectures facilitated by VLSI technology. This book has grown out of our belief, based on actual experience, that future progress in image processing, especially in the context of highly concurrent processing, is going to be increasingly dependent on the establishment of productive formal links between these two important domains. This book has much to say about both algorithms and advanced computers, often in the context of specific image processing applications, but the underlying, sometimes implicit, theme is the importance of understanding the potential mappings which exist between the two domains.

This book is not a primer in image processing, in that a certain amount of prior knowledge is assumed. Many excellent introductory and review texts already exist, and it was not our intention to produce yet another. Rather, we have tried to present some new and novel perspectives of image processing in the hope that they are of interest and eventual utility. As a consequence, both the choice of material and its organisation represent something of an experiment. I hope that you, as the reader, feel that in some measure it succeeds. Chapter 1 covers some introductory background material on image processing techniques and computer architectures, as well as setting the scene for what is to follow. Chapters 2 and 3 cover algorithms and architectures respectively, while Chapter 4 attempts an initial synthesis between these two domains. In order to provide some concrete illustrative examples, Chapters 5 and 6 cover two very different image processing applications areas – machine vision and synthetic aperture radar. Finally, Chapter 7 presents one view of what the future might hold.

Many people, in one way or another, have contributed directly and indirectly to the preparation of this book. In particular, various members of the research and support staff of the GEC Research Laboratories, Wembley, UK, and especially Mrs Nell Peters, who has borne much of the brunt of manuscript preparation. To them all, my heartfelt thanks.

R. J. Offen

Contents

Chapter 1

Whence Image Processing?

R. J. OFFEN

1.1 INTRODUCTION

An increasingly accepted contemporary view of image processing is to regard it, in a broad and generic sense, as being nothing more than one aspect of 'information processing' (Machlup and Mansfield, 1983), as opposed to the more *ad hoc* traditional view which tended to embrace one or more, depending on individual preference or prejudice, of various technological niches including, *inter alia*, picture processing, image processing, pattern recognition, scene analysis, image interpretation, optical processing, video processing and image understanding. These niches overlap to greater or lesser extents and each has its own history, peculiarities and ardent supporters.

There have been a number of attempts to produce some measure of unification and demystification by categorising these niches into more reasonably circumscribed areas, such as 'signal processing', 'classification' and 'understanding' (Cohen and Feigenbaum, 1982). Signal processors transform an input image into another image that has desirable properties. For example, the output image may have a better signal-to-noise ratio or have certain features enhanced in some way. Image processing and picture processing are the most common terms for this class of processing. Both digital and optical techniques can be used. Classification techniques classify images into predetermined categories. Character recognition and signature verification are typical examples. Pattern recognition and pattern classification are the most commonly used terms, though the word 'recognition' has mainly a historical context. Image understanding has a definite semantic context, requiring both knowledge about the specific problem domain, often complex and three-dimensional, as well as sophisticated image processing techniques. The inherent complexity of image understanding has resulted in the fact that decomposition of processing into 'early', 'intermediate' and 'late' stages is almost axiomatic.

We will use the term 'image processing' as a catch all for all of these activities, with the implicit understanding that the fundamental underlying activity is that of 'information processing', performed on two-dimensional projections of some property of the physical world.

The most important factors in bringing together the disparate aspects of these diverse yet strongly related fields has been the burgeoning availability

此为试读，需要完整PDF请访问：www.ertongbook.com

and use of advanced computing, storage and display technologies, especially those incorporating VLSI design and fabrication techniques. This intrusion of radically new technology has caused some fundamental rethinking of many aspects of image processing. In fact, we shall adopt the premise that VLSI has become the natural implementation route for realising many of the essentially parallel algorithms which are characteristic of image processing.

Figure 1.1 illustrates, in a simple schematic form, the linkages which exist between the key conceptual aspects of implementing image processing systems. The algorithms–architectures linkage is fundamental, yet it is still not clearly understood. Currently it represents the weak link in the design pathway. It is the exploration of this link, in the context of the very considerable potential of VLSI design and fabrication technology, which provides the underlying theme of this book. This lack of understanding is especially critical in the context of systems with high levels of concurrency. The current state of system implementation is very much a case of machine architectures dictating algorithmic design. The classical von Neumann architecture is so central to the intellectual ethos of modern computing that the architecture–algorithm duality is firmly skewed towards this type of architecture. It is only with the developing interest in highly concurrent non-von Neumann architectures, coupled with the high degree of versatility

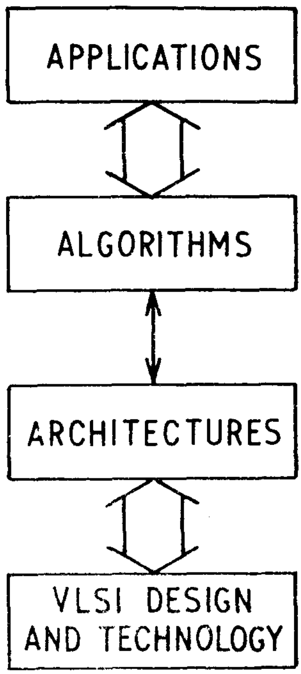

Fig. 1.1 The interrelationships which exist between the key stages which underlie the design and implementation of VLSI-based image processing systems. The width of the arrows represents the extent of our understanding of the implicit mappings.

inherent in VLSI design and fabrication, that researchers have started to question some of the conventional routes for implementing algorithms on computers.

In studying, in a general sense, the interrelationship between computational algorithm and implementation architecture some fundamental issues soon become apparent. These include questions relating to representation, description, data structures, mappings, data and control flow, and performance. These are all issues which will be explored in later chapters. At an even more fundamental level, in the context of this book, is the issue of what one is doing in attempting to process images, especially at the level of image understanding. The central paradox of much of image processing, especially image understanding, is that it is in general underconstrained in the sense that given one or more images of a scene the set of possible real-world scenes is very large indeed. In this sense it is reasonable to regard any image of the real world as a two-dimensional projection of the physical universe. An intrinsic characteristic of image formation is the information loss inherent in the process of capturing the image. The intensity values represented by the individual pixel values are the result of the interaction of many factors, including the intensity, colour, location and nature of the light sources, the position, reflectance, transparency and opaqueness of the objects in the scene, the transmission, refractance, absorption and scattering properties of the transmission media, the optical and electrical properties of the imaging device, etc., etc. The difficulty in interpreting images occurs not because we do not understand these diverse phenomena but because of the confounding of complexities inherent in the mappings of these various processes onto the resulting image. While a given scene gives rise to a unique image in any given direction, that image (and knowledge of the direction of projection) does not uniquely specify the three-dimensional scene that it resulted from. There is, for example, no way to fix the depth of any point seen in the image. Mathematically speaking, a projection operator does not have a unique inverse. The question is how properly to represent and sensibly to resolve these uncertainties? Mackworth (1983a) has proposed a set theoretic approach to this problem, where the imaging process is represented, in the simplest case, by mappings between the 'world', 'projection' and 'image' domains. This is illustrated in Fig. 1.2. This, at least in principle, makes possible the representation of all the solutions, potentially infinite in number, to the image interpretation problem in a finite manner. This is a representation, which though it may be of considerable theoretical interest, needs further exploration if it is to be of any practical use.

As implied above, an 'image' in the context of this book is a two-dimensional, almost invariably Cartesian, array of data resulting from sampling the projected instantiation of a local variable, the scene brightness function, obtained via a sensing device. The function values are either brightness values or vectors of brightness values sensed in different spectral bands, e.g. colour images. In the black-and-white case these values are

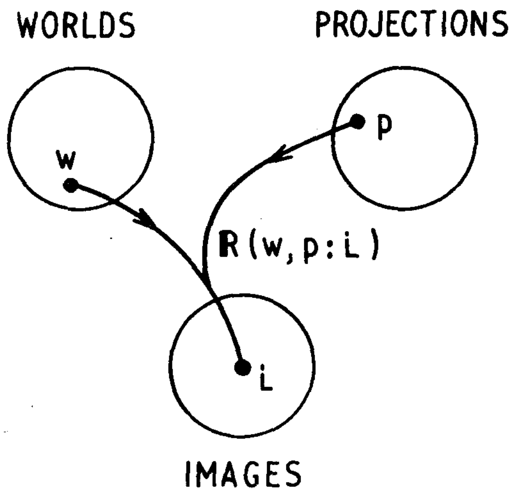

Fig. 1.2 Symbolic interrelationships between the world, projection and image domains (*ibid.*). **R** represents the 'mapping function'.

usually called grey levels. The array values are usually real, non-negative, bounded and implicitly zero outside the field of view bounded by the array dimensions. Digital image arrays can be very large. For example, digitising a fairly standard black and white photographic image can involve a storage requirement of 8 Mbits for an array of 1024 × 1024 elements. In many applications even finer resolution is required, with the attendant increase in storage overheads.

Image sensors as such will not be discussed here but in essence such sensors may be regarded either as 'array' devices or 'scanning' devices. An array sensor provides a discrete array of samples of the scene brightness function, while a scanning sensor provides a set of cross-sections of this function along the path of the scanning pattern. We will be concerned in this book only with the situation where the sample array is digitised subsequent to being sensed. These digitised array elements are often called 'picture elements' or 'pixels'. Obviously, in areas such as optical processing the samples are captured and stored on photographic film, or other image transducers, in an analogue form. The images and the processing tasks used to illustrate the various themes and concepts raised in this book will be drawn from areas as diverse as medical imaging, synthetic aperture radar and stereoscopy.

In the sense that much contemporary computer programming is concerned with the manipulation of relatively simple word-oriented data structures, the situation is starting to arise where it makes more sense to move towards programming with actual images, as opposed to their decompositions. This will require the development of new techniques, environments and computing facilities so that users can easily understand, specify and solve a wide range of image-specific problems. Programming, be it with very simple data or with complex images, is essentially a modelling process involving abstraction, decomposition and simplification. It is already clear that many image processing tasks are going to require comprehensive yet tractable computational models if much in the nature of practical applications are to be realised. In this context it is easy and sensible to regard image processing as a generic activity in the sense that much of its application is going to be dependent on common representations, models and underlying theory. This book has a lot to say about this theoretical support and the consequent design and implementation tools, as well as the underlying technologies which are necessary to support such a unifying concept.

In subsequent chapters various different types of image processing operation will be discussed in the context of a range of different applications. Many of them can be classified as either *point, local* or *geometric* operations. Point operations involve transforming single pixels in a way that does not depend on any neighbouring pixels. Local operations involve transformations on pixels so that the new value of each pixel depends also on the values of pixels in some neighbourhood. Geometric operations involve pixel values at some other point, defined by a geometric transformation, or in a neighbourhood of that point.

The following sections of this chapter briefly introduce some of the central concepts necessary for a proper understanding of subsequent chapters. The reader well versed in the field of image processing should skip Section 1.2 and resume at Section 1.3.

1.2 THEMES AND TECHNIQUES: SOME BASIC CONCEPTS

The history of the processing of images by digital methods is relatively circumscribed, in that it is essentially no older than the digital computer itself.

Computer vision began in the 1950s with statistical pattern recognition, the goal of which was to assign an input image to one of a small number of predetermined classes. Optical character recognition is a representative application. Digital image processing technology for the enhancement, restoration, coding and transmission of images soon started to appear at about the same time, and now constitutes a large and sophisticated field that incorporates many recent computer vision techniques. True computer vision, with the goal of 'understanding' images of complex three-dimensional scenes, was first attempted in the early 1960s. The immense computational complexity inherent in mimicking the human visual process soon became apparent. Intuitively appealing detectors for visual features, such as object boundaries and schemes to control processing proved unreliable and inadequate. Devoting massive amounts of processing at the early stages of vision was technically and economically impossible, so in the 1970s a cognitive approach to computer vision emerged which conveniently minimised image-level computation and emphasised the symbolic manipulations to which computers are well adapted. In such 'knowledge directed' vision, computational effort is directed by processes that use facts about such phenomena as gravity, support, occlusion, or the likely spatial relations between objects in the scene. Research turned toward representing and manipulating facts about particular domains, such as polyhedral blocks or terrain features, and exploiting the domain-specific knowledge in vision. The representation and application of knowledge is itself, however, a very difficult branch of Artificial Intelligence, and the available techniques proved inadequate to bridge the gap between the input image and the desired symbolic descriptions of it. In the 1980s the consensus of the computer vision community is that the gap is best bridged by a varied and redundant set of visual data representations arranged in a hierarchy of increasing abstraction. The production of many intermediate representations requires a huge amount of computation, but the human vision system seems to do it, albeit with neural structures that operate quite differently from today's digital computers.

Certainly, the most explosive growth has been during the past decade, with the rapidly increasing use of image processing in a multiplicity of new applications, some of which are discussed later in the book. This growth has been the concomitant of considerable improvements in the speed, size, versatility and cost-effectiveness of digital computers and their peripherals. This, in turn, has been the consequence of significant advances in the areas of LSI and VLSI technology. This is a topic to which we will return in later sections. Certain other key events and activities have contributed in a seminal manner to this explosive growth of the application of image processing techniques.

Without the popularisation of the fast Fourier transform algorithm in the mid sixties much of the current activity in the transform processing of images would, in all probability, still be in the realm of optics.

The space programmes of a number of countries have provided a major impetus for computer picture processing. Television coverage of the lunar landings involved major programmes in image coding, transmission, correction and restoration. The capture, transmission and analysis of images captured during planetary 'fly-pasts' imposed even more stringent conditions and so further extended the 'state of the art'.

Attempts to automate the extraction of data from complex images arising, *inter alia,* in medicine, nuclear physics, industry, remote sensing and astronomy have also contributed significantly to our knowledge of image processing. Some of these are discussed later.

Inevitably, the endless search for military superiority by the major powers has provided an engine for progress in such areas as coding, restoration, target recognition, ranging and scene understanding, especially in non-visible portions of the electromagnetic spectrum.

The development and commercialisation of computers more architecturally suited to image processing than the traditional von Neumann machine has been a crucial energising process. Specifically, a number of parallel-array machines and very heavily pipelined special purpose conventional machines have been developed which have performance figures in the Giga operations per second range. It is only by using machines with this level of performance that significant real-time image processing becomes a realistic proposition.

The emergence of Artificial Intelligence (AI) as a discipline in its own right has been a significant factor. The majority of researchers in AI tend to view perception as a form of problem solving, at least at some stages of visual processing and perception. Thus, computer vision and its many potential applications have become major topics of research within the AI community, world wide. An early and highly fruitful consequence of this interest was the decomposition of visual processing into passive, domain-independent, data-driven image analysis (or low level vision) and domain-dependent goal-driven scene analysis (or high level vision): a highly pragmatic decomposition which has persisted. It should be pointed out that since the early sixties there has been a marked divergence between the pattern recognition and AI approaches to the machine understanding of images. The former approach has continued to stress the use of *ad hoc* image features in combination with statistical classification techniques. More recently, use has been made of 'syntactic' methods, where images are 'parsed' in terms of a hierarchical set of primitive constructs. In contrast, the AI approach has continued vigorously to pursue and expand the problem solving metaphor.

The rapid growth of computer graphics and digital scene simulation techniques has provided a tangential yet important metaphor for image processing.

Much of current research centres around the production of intermediate representations formed before object recognition is attempted. These image-like representations are registered with the input image and contain values of physical parameters of scene points such as the distance from a sensor to the point, the albedo of surfaces, the direction of motion of objects, the orientation of surfaces, and so forth. It is usual to assume that the processes that produce these images are part of 'early' or low level vision in that they do not require domain-dependent facts. In fact, these processes cannot be regarded as completely general and reliable, since so much information is projected away in the two-dimensional input image. The fact that they so often work correctly in animal vision seems to imply that they rely on natural constraints or assumptions about the world to derive unambiguous output. A goal of modern computer vision research is the identification and use of such constraints. This, in turn, calls for seeking out properties of the physical world that could help a visual process do useful work, making mathematical models of their interaction with visual phenomena, and implementing the mathematics in computer programs. Currently, attention is centred on the designing of processes that can be implemented on parallel computational architectures, since only in this way is the sophisticated performance of human vision remotely attainable.

此为试读，需要完整PDF请访问：www.ertongbook.com