Cellular Logic Image Processing

M. J. B. DUFF

T. J. FOUNTAIN

Cellular Logic Image Processing

Edited by

M. J. B. DUFF

and

T. J. FOUNTAIN

Department of Physics and Astronomy,
University College London

ACADEMIC PRESS

Harcourt Brace Jovanovich, Publishers

London Orlando San Diego New York
Austin Montreal Sydney Tokyo Toronto

ACADEMIC PRESS INC. (LONDON) LTD.
24/28 Oval Road
London NW1 7DX

United States Edition published by
ACADEMIC PRESS INC.
Orlando, Florida 32887

British Library Cataloguing in Publication Data

Cellular logic image processing.
1. Image processing
I. Duff, Michael J. B. II. Fountain, T. J.
006.4'2 TA1632

ISBN 0 12 223330 1

Printed in Great Britain

CONTRIBUTORS

K A CLARKE, IGE Medical Systems Ltd, Colney Street, St Albans, Herts AL2 2ER.

M J B DUFF, Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT.

T J FOUNTAIN, Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT.

H H-S IP, Imperial Cancer Research Fund Laboratories, P O Box 123, Lincoln's Inn Fields, London WC2A 2PX.

G P OTTO, Department of Computer Science, University College London, Gower Street, London WC1E 6BT.

S D PASS, Technology Group, Marconi Command and Control Systems, Chobham Road, Frimley, Camberley, Surrey GU16 5PE.

D J POTTER, Department of Computer Science, Rensselaer Polytechnic Institute, Troy, NY 12181, USA.

D E REYNOLDS, Cambridge Consultants Ltd, Science Park, Milton Road, Cambridge CB4 4DW.

A M WOOD, Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT.

PREFACE

The Image Processing Group in the Department of Physics and Astronomy at University College London traces its origins to a project in 1958 which had as its objective the design and construction of automatic photoelectronic instruments for measurement of tracks in nuclear emulsions. In the following quarter of a century the Group widened its interest in image data to include all types of images and also carried out extensive studies of computer architectures designed specifically for the analysis of two-dimensional data arrays.

As in any other university research group, a large amount of the detailed work has been carried out by post-graduate students in research projects directed towards a doctorate, in this case under the supervision of one of the editors, Professor Michael Duff. However, construction of a full-scale computer involves a considerable engineering effort and the other editor of this volume, Dr Terry Fountain, was responsible for directing the construction programme for the image processing CLIP systems.

It seemed worthwhile to collect together a representative selection of extracts from the many PhD theses which have been written by members of the Group over the past years since these provide an insight into the philosophy which stimulated the research programme and led to the development of one of the world's fastest image analysis systems. Many others contributed to the research but space did not permit the inclusion of more material.

It is only right that our indebtedness to the late Professor Sir Harrie Massey, the previous Head of our Department, should be recorded here. Without his steadfast support the CLIP project would probably have not been able to survive, particularly during a period of very sparse funding. The continuing support from the present Head of

Department, Professor Franz Heymann, has ensured the project's continuation at a time when so many research programmes are facing extinction. Finally, the lasting confidence of the technical assessment committees advising the Science and Engineering Research Council, under whose sponsorship this project has been conducted, is gratefully appreciated.

No research work is carried out in a vacuum and an attempt has been made to acknowledge other research which has influenced the CLIP programme. In the area of processor array design, few projects have actually progressed to the stage of construction. The ICL Distributed Array Processor project was initiated in 1972 but construction of a pilot DAP did not start until 1975, two years after the commissioning of CLIP3. MPP, built for NASA by Goodyear Aerospace, was not completed until 1983. These and other competing systems now being designed have made their impact on the commercial world but only recently has a manufactured version of CLIP4 reached the market. The emphasis in the CLIP programme has been to build systems primarily in order to study the relationships between architectures and algorithms. Except for the ever-present restraints imposed by limited funding, engineering design for cost-effectiveness has not been of major concern; it is gratifying to see that colleagues in industry have been able to meet this challenge.

Our thanks are due to all those who have not only permitted us to use extracts from their theses but who have also not objected to the sometimes extensive editing which has been necessary in order to harmonise the extracts into what we hope is a coherent account. In this respect, the editors must accept responsibility for any errors and omissions which may have resulted as a consequence of their actions. We are particularly grateful to Miss Annette Harris for her careful and patient work in preparing the camera-ready material for this volume.

Michael J B Duff
Terry J Fountain

University College London
November 1985

CONTENTS

INTRODUCTION

It is, perhaps, one of the strange consequences of the manner in which television cameras are constructed that the traditional approach to image analysis has been not only to enter image data into a computer in raster-scan order, but also to process the data in a similar manner. The conventional computer program starts at the top left-hand pixel in the image and proceeds row by row through the image. More enlightened algorithms may divert from a completely regular path in order to follow around object contours but this can, and often does, lead to unwanted complexity.

In the late 1950s electronic computing was still in its infancy; university computers were few and far between and laboratory computers almost non-existent. In this discouraging environment, scientists were generating vast amounts of data so that the experimental bottleneck was often in the data reduction and data analysis phases of the project. The problem of how to cope with a 'data mountain' was probably most severely felt in High Energy Particle Physics where data appeared in the form of particle tracks in nuclear emulsions and cloud, bubble and spark chamber photographs. It was obvious that techniques would have to be developed that would speed up data processing by several orders of magnitude and various projects were initiated in association with particle physics laboratories in the United States, and in the United Kingdom and other European countries, especially in the CERN laboratories.

At University College London, in the (then) Department of Physics, a small research group was set up in collaboration with the already thriving Nuclear Emulsion Group. The new group, originally known as the Automatic Methods Group, was charged with the task of developing a range of devices which could be used to assist microscopists in their time-consuming work of finding and measuring the tracks of charged particles in their nuclear emulsions. The systems produced at that time usually comprised a microscope with one or more motorised movements (on the X-Y stage, on the objective or in the graticule eyepiece), together with some

CELLULAR LOGIC
IMAGE PROCESSING
ISBN 0 12 223330 1

form of digital position encoder on each movement. Most devices were semi-automatic in that an observer was required to set the microscope measuring graticule on the desired part of the particle track. Calculations on the digitized data were performed either by solenoid-controlled mechanical adding machines or else by arithmetic-logic circuits constructed from solenoid-operated relays. As time progressed into the middle 1960s, the cumbersome and not always reliable relays were replaced by diodes and eventually by transistors and, in more advanced systems, the human microscopists could themselves be replaced, for a very limited set of tasks, by television cameras or scanned optical systems and photomultipliers.

Despite a reasonable degree of success in projects such as these [e.g. 1], it soon became apparent that the systems being produced were strictly limited in their capabilities. An inherent weakness ran through them all: although they were well able to remove much of the drudgery from straightforward measuring tasks, particularly on virtually straight tracks, complex images 'confused' them. Also, the systems could not of themselves be of help when it came to finding and recognising tracks in particular configurations or patterns. Systems designed to extract data from track chamber photographs were similarly restricted in their effectiveness, even though by this time large mainframe computers were appearing on the scene and bubble chamber data computations were being performed off-line on data collected from the semi-automatic measuring machines.

The particular difficulty affecting this method of working concerned the enormous proportion of irrelevant data in the photographic images. It was not practicable to record all the data in digital form, using subsequent computation to reject the unwanted portion. Major rejection before recording was essential and had to be the responsibility of human operators. What was needed was some form of built-in pattern recognition capability in the recording or scanning system itself.

At the same time as physicists were trying to find ways of analysing vast amounts of data, biologists and their engineering colleagues were attempting to understand the structure and behaviour of the nervous system by designing neural analogues. Neurons were modelled and assembled into networks (neural nets) which exhibited some of the observable characteristics of their physiological counterparts. Most of the model design was, of necessity, inspired guesswork since detailed anatomical and physiological information was not (and still is not) available. Nevertheless, the

models helped to provide insight to guide further investigations of neural behaviour and, significantly from the point of view of computer science, served to inspire a new approach to data analysis.

Four influential papers were published in the late 1950s and early 1960s: two by Unger [2,3] describing a mesh of simple processors for image data analysis, one by Lettvin et al. [4] proposing a model for the visual system of the frog, and another by Hubel and Wiesel [5] presenting a now classic description of the structure and function of the cat's visual cortex. Herscher and Kelley [6] modelled the frog's visual system and Unger's ideas were embodied into designs for image analysis computers proposed by Slotnick et al. [7] and, later, McCormick [8]. The work by Hubel and Wiesel continues to direct attention to the use of clusters of detectors and processors as a means of implementing local neighbourhood operators for line segment detection and for spatial filtering.

At University College London, attention focussed on the specific problem of automatically locating areas of interest in charged particle tracks. Since a straight track can be assumed to be not interacting, it follows that interaction points are also points at which the track deviates so the problem reduces to one of finding sharp changes of direction in the track, positions which are usually referred to as vertex points. Once such positions could be found, measuring tracks radiating from each point was a process lending itself readily to automation. It was decided to build a parallel processing system, loosely modelled on the retina, to find vertex points automatically and rapidly, as a preprocessing stage prior to measurement.

UCPR1

In 1967 Duff et al. [9] demonstrated a fixed logic, 400-processor system arranged as a 20 x 20 square array. Charged particle track images were projected on to an input array of photodiodes. The outputs from the photodiodes, suitably thresholded, were then summed over three by three windows. A further layer formed sums over a five by five ring of outputs from the first summation.

Finally, a global threshold was applied, reducing from a maximum until at least one processor output exceeded the current threshold level. At this point, one or more miniature indicator lamps lit up in the corresponding positions

in the output array. These lamps indicated the approximate position of a vertex in the input image. In a subsequent development, the same system, operating with an increasing rather than a decreasing threshold (and a small circuit modification to ensure detected points were elements of the input particle track image), was able to locate line ends. By combining the two functions and by partitioning the output into regions, a crude line feature analyser was constructed enabling simple recognition of line figures (such as a subset of the alphanumerics).

The capabilities of UCPR1 were extremely limited although it did serve to illustrate the potential of parallel processing arrays. The similarity to the retina was superficial so that the system could certainly not be described as a retina model. Furthermore, as a device for aiding the analysis of track chamber photographs, it was of no direct value. What, therefore, was learned from this project?

In the first place, the processing speed was impressive by the standards of the day, vertices being found in about 20 ms. Secondly, it was soon clear that a simple four-layer processing array (comprising input, two summations and output) could hardly be applied to more than a few elementary tasks. Thirdly, the technology employed in the array construction, which incorporated hand-assembled multilayer printed circuit boards, was clumsy and inappropriate. The possibility of constructing an array with, say, ten or more layers of logic, each performing a fixed function but allowing data to be moved in either direction between layers, was given serious consideration. However, the technical difficulty of making the necessary connections between the layers was daunting and other lines of thought were pursued in order to avoid this problem.

The Diode Array

In the Laboratorio di Cibernetica, near Naples, Italy, Levialdi [10] had been investigating the processing capability of a simple array of switches, connected so as to detect closed loops in binary figures. This array demonstrated the power of propagation as a means of passing information between processors over long distances in the array. Discussions between one of us (MJBD) and Professor Levialdi, marking the beginning of what was to become a long and fruitful collaborative association, resulted in a study being carried out to determine the intrinsic processing

capability of an array of 'minimal' processing elements. These elements were to contain the smallest set of components that could be thought to have the required capability: an input and an output for image data, paths between adjacent processors, and a means for inverting or switching on or off the outputs from the processors, both the data outputs and the outputs to the neighbouring processors. A processing 'cell' was defined, incorporating a two-pole, double-throw toggle switch, a neon indicator (with a series resistor) and diodes linked to input and output buses surrounding each cell. The series resistor permitted the neon indicator to be illuminated or extinguished by applying a positive voltage to one or other side of the indicator when the resistor itself was earthed. The array was 'programmed' by choosing an appropriate set of connections within each cell, the connection pattern being identical in every cell.

A trial 5 x 5 cell array was constructed, using electromagnetic relays to select a variety of connection patterns, primarily to demonstrate visibly the ideas forming the basis of the project. At the same time, a Fortran simulation was written and linked to a Monte Carlo program which explored arbitrary connection schemes by choosing them at random and applying the resulting array to a standard test image. Schemes which generated new functions, as shown by the appearance of a previously unseen output image, were stored for future reference. The random search was moderated by a set of intuitively obvious constraints, such as rejection of schemes which would short out the power supply or would only connect to one side of the toggle switch. A typical result occurred when 3906 trials generated 73 schemes, each with different processing properties, the program taking about 30 minutes of execution time in an IBM 360 computer. The discovery curve was clearly flattening by the end of the run and it was interesting to note that all the ten schemes which had already been devised and wired into the demonstration array had emerged from the Monte Carlo search after 18 minutes of processing time.

The Diode Array served as a useful testbed for developing ideas about arrays of mesh-connected processors and pointed strongly to the potential value of propagation. Attempts to describe propagation algebraically met with the difficulty that the diode paths were bidirectional for signals of opposite polarity [11]. By specifying a processing element in terms of pure logic functions (rather than real circuit elements), this problem could be

circumvented. It then only remained to translate the logic specification into the then newly available small scale integrated circuit components.

CLIP1

Despite encouraging results from simulations, little had been done to investigate the practical problems that might arise when mesh-connected processor structures were actually constructed and operated. A haunting fear for those involved in the 400-processor UCPR1 had been that a short-circuit in one element might propagate throughout the array, leaving a trail of burnt-out components in its path. With the integrated circuit network about to be built, dubbed CLIP (a Cellular Logic Image Processor), it was thought that an added hazard might be that the array would go into oscillation due to the presence of signal feedback paths through chains of neighbouring elements.

The processing element circuit was extremely simple and comprised eight two-input and two four-input NAND gates in three small scale integrated circuit packages. One hundred of these elements were mounted on one large printed circuit board to form a 10 x 10, 4-connected array (i.e. each element connected to its N, E, S and W neighbours). A flying-spot scanner generated a binary image which was then stored in the input memory (a shift register). One of three possible functions was selected by appropriate control lines and the contents of the shift register connected to the array element inputs through 100 parallel output lines. The output data was buffered before re-entry into the shift register for display. The three functions chosen for this trial array were extraction of the contents of closed loops (of 1-elements), extraction of sets of 1-elements connected to the array border, and extraction of the outer edges of objects composed of 1-elements. Images were displayed by modulating the brightness of a 10-line raster on a display oscilloscope. The input image could be altered pixel-by-pixel using a light pen to interact with the display of the input image stored in the shift register.

The CLIP1 project and the subsequent development of CLIP2 and CLIP3 was the doctoral research programme of D M Watson [12], working under the supervision of one of us (MJBD). The fears expressed above were largely allayed by CLIP1, instilling confidence to proceed to a more general processor.

CLIP2

There are only 16 Boolean functions of two Boolean variables and they can be generated by conditionally summing the four minterms. In CLIP2, a 16 x 12 hexagonally-connected array of processing elements, every element comprised two programmable Boolean processors each implementing independently the full set of functions of two inputs. One Boolean processor generated an output which it then transmitted to its six neighbours, whilst the other processor output into an image memory for subsequent further processing or display. The two inputs were either both binary images or else a binary image at one input and the ORed interconnection signals at the other. Additional circuits allowed a second image to be ORed in with the interconnection signals at the second input. Once again, the input image was generated by means of a flying-spot scanner in conjunction with a light pen. Instructions were entered manually (as 12-bit words) and stored in a 32-word memory. The system was completed in 1972.

Apart from the obviously small image size, CLIP2 suffered from one major shortcoming: its operations were completely non-directional. When information was sent from processor to processor, all neighbouring processors were addressed simultaneously so that directional information was not retained or transmitted. Possible functions were therefore limited to those in which direction was not implied, involving properties such as enclosure, expansion and connection. This extremely severe limitation was anticipated from the start; CLIP2 was still seen as a testbed for studying array control strategies and for exploring the propagation of signals through arrays. It is also a fact that, in the early days of the CLIP programme, these arrays were regarded as special-purpose image processors, not as general-purpose computers. It was, perhaps, the study of CLIP2 which suggested that a small amount of sophistication of the processing elements would generalise the array so that it would be inherently capable of performing any image operation.

CLIP3

The required processor enhancement appeared in CLIP3 in which the interconnection structure between processors was improved by individually gating each direction (now 6 or 8, providing hexagonal or square connectivity) and by replacing

the OR gate by a threshold gate. Strictly speaking, the threshold gate was not a necessity although it reduced substantially the number of instructions needed to implement certain valuable local neighbourhood operations. The instruction word was increased to 24 bits and the instruction memory to 256 words. Sixteen bits of local image memory were provided at every processor and various modes for loading data into this memory were allowed.

Many programs were written for CLIP3, ranging from conventional image processing to non-image operations such as maze solving, electrostatic field calculation and lay planning (fitting together garment pattern parts so as to use cloth economically). The small image area limitation was partially overcome by constructing a scanning system which covered a 96 x 96 pixel image with the 16 x 12 processor array. This development made it feasible to write programs to process grey-level images, using bit-serial algorithms; the smaller image size had been too small to allow grey-level images to be represented meaningfully.

The next phase in the CLIP programme was devoted to refining the processor specification, including only those parts which seemed essential but adding a few features which were judged likely to produce a significant improvement in performance. The motivation was to reach a design which could be fabricated as a large scale integrated circuit, principally so that larger arrays could be built at a reasonable cost. The vastly slower performance of the scanned version of CLIP3 (about 3000 times slower than CLIP3 itself) stimulated the desire to build a 96 x 96 processor array which would recover most of the speed lost in the scanning operation.

CLIP4

Despite the usefulness of the CLIP3 threshold gate, it was an expensive luxury in terms of the number of transistors needed to implement it. On the other hand, the need to double the local memory was apparent and a useful increase in processing power could be achieved at little cost by incorporating the few gates needed to convert the processor optionally into a full binary adder. The only readily available integrated circuit technology (available to the UK university community, that is) was metal gate NMOS. Preliminary calculations indicated that it should be possible to design a circuit comprising eight processors, each with 32 bits of local memory, the resulting chip being of a size

which might be expected to give a satisfactory yield. The target price for the eight-processor packaged device was £6 (1974 prices).

Sadly, the initial optimism in this project was soon to be dampened. A full array of 96 x 96 CLIP4 circuits was not forthcoming until early in 1980, some six and a half years after the design contract was placed. In the intervening period, three design houses had applied themselves to the project and, at the end, the circuits produced (usually referred to as CLIP4A) could only be used at 40% of the design clock frequency, 1 MHz rather than 2.5 MHz. Undoubtedly, this shortcoming was the result of an attempt to salvage the aborted earlier work when the last company involved took the project to conclusion. Subsequently, the same organisation, Swindon Silicon Systems Ltd, carried out a complete redesign from first principles and eventually produced a version (CLIP4D) which fully met the original specification. In the Chapters which follow, the work described relates to the CLIP4A circuit but applies equally to CLIP4D in all respects other than performance times.

Details of the operation of the CLIP4 system appear in subsequent Chapters. The system comprises an array of 9216 bit-serial processors, configurable as either a square or hexagonal 96 x 96 array, interfaced to two framestores holding 6-bit images. The input store can be loaded from a variety of television camera optical workstations via an A/D converter which samples the central part of a non-interlaced standard television frame, producing the required digital image. The output memory offloads through a D/A converter for display on a standard television monitor, and facilities are provided for mixing, at variable intensities, the original analogue image and one or other of the input or output digital images. The user communicates with CLIP4 via a DEC PDP-11 computer operating under UNIX. The serial host is time-shared but CLIP4 itself can only accommodate one user at a time. However, control can be switched to another user in a matter of seconds and the comprehensive file system provided by UNIX enables users to keep track of the whereabouts of both their programs and their image data.

In the following Chapters various aspects of the system will be described by some of the postgraduate students who have used it in the five years since it was commissioned. Many other users have worked with the system and CLIP4, although a prototype machine, has given excellent service. Two similar machines have been constructed by the Science and Engineering Research Council (Rutherford Appleton Laboratory) so that others could gain experience of CLIP