INVITED PAPER : ISPRS 16th Congress, Kyoto, Japan, July 1-12, 1988

STEREO MATCHING USING TRANSPUTER ARRAYS

J-P Muller(1)
K A Collins(2)
G P Otto(3)
J B G Roberts(4)

(1) Department of Photogrammetry & Surveying, University College London, Gower Street, London WC1E 6BT, U.K. (ARPANET : jpmuller@cs.ucl.ac.uk)

(2) THORN EMI Central Research Laboratories Limited, Dawley Road, Hayes, Middlesex, UB3 1HH, U.K. (ARPANET : cols%alvey.uk@cs.ucl.ac.uk)

(3) Department of Computer Science, University College London, Gower Street, London WC1E 6BT, U.K. (ARPANET : otto@cs.ucl.ac.uk)

(4) Royal Signals & Radar Establishment, St Andrews Road, Gt. Malvern, Worcs. WR14 3PS, U.K. (ARPANET:jbgr%rsre.mod.uk@cs.ucl.ac.uk)

Commission III

Abstract

The Alvey-funded* project on "Real-time 2.5D Vision Systems" is concerned with the development of a real-time exemplar machine for stereo matching. This machine will be a loosely-coupled multiprocessor, built using INMOS Transputers (32-bit microprocessors with fast built-in serial links). The main characteristics of such machines will be described and contrasted with both SIMD and other MIMD architectures.

The implementation of several different stereo algorithms on a Transputer array will be discussed, with particular attention to efficiency, ease-of- implementation and extensibility.

From this, a set of requirements will be deduced for a machine for real-time stereo ranging. Using these requirements, a comparison will be made between various possible machine architectures and arguments presented to show how and under what circumstances special-purpose hardware will prove to be more effective. Finally, an architectural description will be given of the Alvey exemplar machine.

1 Introduction

Automated accurate measurement of surface shape of objects from the topography of the Earth to the microstructure of man-made materials has long been a goal in photogrammetric research and in the operational application of photogrammetric engineering. Over the last twenty years, extensive efforts in the automation of three-dimensional object recognition by computer scientists, engineers and cognitive scientists has led to the development of many hundreds of so-called stereo matching algorithms whose primary aim is to provide sufficient knowledge of an object's shape in order to enable machines to interact with it and its environment.

* in collaboration with Laser-Scan Laboratories, Cambridge. The UK Alvey programme is funded by the Department of Trade and Industry, the Science and Engineering Research Council and the Ministry of Defence.

These two aims have often seemed incompatible and this is exemplified by the differences in approach between the two cultures, one based on obtaining high spatial density and accuracy, the other based on extracting the minimum amount of information for characterisation of an object's shape for automated recognition.

Whichever approach is adopted for whatever practical application, they are unified by one common problem - computational processing speed. In the case of video refresh rate processing of two 512x512x8-bit input images into one 512x512x32-bit disparity output image, upwards of 1 billion (10^9) arithmetic operations per second is required with an I/O rate of upwards of 200 Mbps. This is approximately equivalent to the processing rate of today's supercomputers (**Hwang, 1987**) of which very few, if any, have sufficiently high-speed I/O to cope with these date-rates.

Supercomputers tend to be inaccessible to most researchers in image processing and certainly do not appear to be well-suited to this class of algorithms as vector processing is far better adapted to solving non-linear Partial Differential Equations than processing pixels. More importantly, the speed quoted for such supercomputers is often misleading when applied to image processing as their processing rates primarily refer to peak performance when inner loops have been vectorised and not average performance for a broad range of processing, storage and I/O tasks.

There have been a number of solutions proposed to solving this computational speed problem. These include special-purpose low-level processors (from vector-array processor add-ons to special-purpose VLSI chips for Digital Signal Processing) to distributed workstations to massively parallel machines.

It is therefore in this context that this Alvey MMI-137 project on "Real-time 2.5D Vision systems" was initiated to address the problem of assessing which machine architecture(s) are best suited to solving the problem of general-purpose stereo matching irrespective of the applications domain. Our objective is to obtain the most accurate disparities possible at the lowest computational cost. Application areas being addressed at the present time include topographic mapping from stereo SPOT images and online digitisation and matching of aerial photographic data as well as shape measurement of anthropometric and industrial objects for direct entry into CAD systems.

In this paper, we attempt to show that coarse-grain MIMD machines constructed of general-purpose 32-bit microprocessors with four fast built-in serial links called transputers appear to offer the best solution to the problems of building future real-time stereo matchers. This will include some preliminary results of parallelising a data-dependent algorithm on an MIMD array of transputer elements.

Other papers in this Congress will describe several other aspects of the Alvey project including results for SPOT camera geometry (see **O'Neill and Dowman, 1988** and **Otto, 1988**); quality assessment of the resultant DEMs (**Day and Muller, 1988**); dynamic visualisation of DEMs coupled with different types of reflectance data (**Muller et al., 1988a**); development of automated GCF extraction system for digital map database (**Stevens et al., 1988**); map-image registration and automated image feature extraction (**Muller et al., 1988b**); the use of graphics workstations for photogrammetric analysis (**Muller and Paramananda, 1988**) and the fractal properties of terrain (**Muller and Saksono, 1988**).

In section 2, we present a review of previous efforts to automate stereo matching in our two areas of photogrammetry and machine vision concentrating on what hardware solutions were adopted. Section 3 describes the computational structure of typical stereo matching algorithms and gives an algorithmic description of the stereo matchers we have implemented to date whilst section 4 gives a short review of special-purpose hardware for image processing. In section 5, results for transputer use in both low-level (e.g. convolutions) and high-level image processing tasks such as stereo matching is given including some preliminary results of parallelising a data-dependent stereo matcher on an MIMD array of transputer elements. Finally section 6 gives some initial estimates on transputer speed for several other different types of stereo matcher and concludes with a look at the importance of extensibility and flexibility when trying to choose between SIMD and MIMD solutions.

2 Previous approaches

2.1 Online stereo matching hardware for aerial photo-based topography

There have been several substantial efforts in the sixties and seventies to devise machines that would recover surface topography automatically from pairs of aerial photographs at sufficiently high speeds to be incorporated into production mapping systems (see **Dowman, 1982**; **Dowman and Muller, 1986**).

These systems have addressed the problem of producing low planimetric quality Digital Elevation Models (DEMs) from primarily small-scale photographs (smaller than 1:50 000 scale) using opto-analogue video correlation. Examples include the Universal Automatic Map Compilation Equipment (UNAMACE) at Defence Mapping Agency, USA (**Bertram, 1969**; **Dowman, 1982**) and the Gestalt Photomapper at Energy, Mines and Resources, Canada (**Kelly et al., 1977**; **Dowman, 1982**). No figures for timings or disparity accuracies are given although **Barrow et al. (1977)** stated that at that time some 40% of the input images were rejected because the UNAMACE equipment would not correlate continuously.

Case (1982) described the design concepts of the Digital Stereo Comparator/Compiler (DSCC) system being developed at that time for the Defence Mapping Agency, USA. The system was said to require a stereo correlator to produce DEM estimates at a rate of 200 points per second with an rms x-disparity of 1 pixel (for a pixel size of 12.5 microns), although whether this digital system ever succeeded in reaching its design goals is unknown.

Since 1982, there have been a number of attempts to produce lower cost systems using computer controlled "analytical plotters" with CCD cameras at one or both of their focal planes and in some cases, digital framegrabbers to enable TV-sized images (up to 640x480x8-bit for RS-170 and 768x575x8-bit for CCIR) to be acquired in a few refresh cycles. These semi-automatic systems include devices from Kern (**Benard et al., 1986**), Wild (**Gruen, 1987**), **Helava (1987)** and **Zeiss (1987)** and rely upon the use of slow serial hosts to perform the stereo matching. Few, if any, figures are available on their performance for stereo matching.

Li(1986) described results of experiments for one such analytical plotter-based system from Kern giving evidence for timings and accuracies. He concluded that a single estimate of parallax (i.e. one correlated point) could be produced in around 7 seconds (on a host PDP11/73) from two input 512x480x8-bit images (at a fixed pixel size of 12.5 microns at the photo-scale) compared with 2-3 seconds for single spot heights derived from manual measurements. **Li(1986)** also showed that whereas manual measurements could obtain typical rms errors of 0.15 per mille of flying height, correlator-based measurements were typically poorer with rms errors of around 0.4 per mille of flying height.

2.2 Online stereo matching systems for Robot vision and Autonomous Land Vehicles

It has long been recognised that vision-based sensory systems need to be developed for telerobotics, path-planning and interactions with a machine's environment (see, for example, **Binford, 1982** and **Shirai, 1987**). The importance of time-critical processing has led some researchers to develop special-purpose or SIMD hardware (see **Flynn, 1966** for the definition of parallel machine classification) whilst many new techniques implicitly rely upon the future development of massive processing resources for their real-world application.

<u>Serial computer solutions</u>

Barnard (1987) and **Bolles et al. (1987)** describe recent work at SRI International within the United States DARPA Autonomous Land Vehicle (ALV) program. Their algorithms for both stereo matching (based on simulated annealing) and obtaining scene structure from spatio-temporal sequences rely upon the future development of massive computational systems. For example, **Barnard (1987)** quotes that their implementation of simulated annealing takes 12hours elapsed time on a Symbolics 3600 machine for two 512x512 images with a disparity range of 50 pixels.

Special-purpose hardware solutions

Gennery et al (1987) describe recent work at the NASA Jet Propulsion Laboratory aimed at developing sensing systems based on multiple-camera stereo for robotic systems aimed at construction in space and satellite capture as well as for autonomous land vehicles for future unmanned missions to Mars. Special-purpose pipeline hardware, called a Programmable Image Feature Extractor (PIFEX) is being developed to produce motion estimates for features extracted from a spinning spacecraft at rates of up to 10 frames per second (**Gennery and Wilcox, 1985**). This system has unknown disparity accuracies or disparity ranges at the time of writing.

Nishihara (1984) described the world's first video-refresh-rate stereo matching system, developed originally at M.I.T. and now being extended at Schlumberger Palo Alto Research Centre (**Nishihara, 1987**). He showed its applications for path-planning and bin-picking (**Nishihara,1984**) as well as microscopic shape measurements (**Nishihara,1987**). The system uses unstructured textured light to provide badly-needed stereo cues for industrial objects. It uses special-purpose convolution hardware to give estimates of the sign of the Difference-of-Gaussian at three different spatial scales and a serial host to perform remaining correlation calculations. Accuracies are claimed to be to 0.1 pixel (**loc.cit.**) for a coarse grid of 36x26 disparities (from a RS-170 source) and for a disparity range of 200 pixels with three stages in the coarse-to-fine matching, the complete image takes 30 to 40 seconds to compute. Non-epipolarity up to 2 lines can be accommodated (**loc.cit.**).

Ohta et al. (1987) described a stereo matching system based on a MC68000-bus based processor board which includes an affine transformation module for resampling input images to epipolar lines. It uses an epipolar-constrained dynamic programming algorithm (see **Ohta and Kanade, 1985**) where the feature-space is reduced by matching only certain types of edge pixels (hereafter referred to as "edgels"). The pipeline nature of the hardware means that its speed is proportional to the number of potential match nodes per scan-line (viz. edgels). Typical processing times for a 256x256 image of an industrial object is 2.5 seconds for some 1300 edgels and accuracies appear to be around 1 pixel disparity.

SIMD solutions

The final example of hardware solutions to stereo matching with robotic/ALV applications comes from M.I.T. where a SIMD parallel machine with 32K 1-bit processing elements (PEs) with 4K RAM/PE machine called the "Connection Machine" has been developed (see **Hillis, 1985**). **Drumheller and Poggio (1986)** describe the application of this machine to the production of stereo-matched edgels of indoor scenes taken by a flexible "eye-hand" system developed in their laboratory. For images of 256x256x8-bits, with an unknown edgel density of stereo matches and an unknown disparity rmse, for a disparity range of 41 were obtained in 1.5

seconds.

2.3 Offline stereo matching systems for remotely-sensed data

The term offline is used in this case to indicate that the data is currently processed at slow rates because of the logistical and technical difficulties obtaining it (e.g. cloud cover for satellite images and time to scan a complete aerial photograph) and because of the massive amounts of digital data that are input. It is usually processed on mini-computers of the Vax 11/750 class with large backing stores.

The first application of offline processing was to small-scale aerial photographs (digitised by an Optronics rotating-drum scanner) by Geospectra (**ATOM™, 1986**). Their unknown epipolar-based stereo matcher can handle disparities of up to 128 pixels but has unknown accuracies at the disparity level (figures are quoted for Large Format Camera with a 50 micron spot size, a 40m grid and a 65m elevation "resolution"). Timings for Vax11/750 correlation are quoted as 800 points per second.

One of the first applications of overlapping satellite images (which occurs preferentially at high latitudes) was the production of DEMs from LANDSAT-MSS 80m resolution data. This early work carried out at the Canadian Centre for Remote Sensing by Simard and co-workers (see, for example, **Simard and Krishna, 1983** and **Grieve and Simard, 1984**) used cross-correlation to obtain disparity (height) rms accuracies of around 0.2 pixels (150m) for the poor B/H ratios (<0.1) of MSS imagery. Timings were not given nor were figures on the amounts of data processed.

Stereo matching has also been carried out for LANDSAT-TM 30m resolution data which has got the same poor B/H (<0.1) as previous MSS data. **Cooper et al., (1986, 1987)** described the Macdonald Dettwiler and Associates (MDA) stereo matching system which is based on dynamic programming of edgels using epipolar lines (based on **Ohta and Kanade, 1985**). Their filtered stereo-matched output gives around 3-3.5% of the total pixels matched with a rms error of around 0.4 pixels (60m) for a range of 20 pixels (500-3367m) for 512x512 image patches taken from either Band 4 or 5 of LANDSAT-TM. Timings were not given.

Recent results on LANDSAT-TM Band 4 data from **Ehlers and Welch (1987)**, based on using a least-squares two-stage correlation system (similar to **Simard and Krishna, 1983**) matched 95.3% of the available pixels with a rms error of 0.3 pixels (42m) for a range of four pixels). Timings were reported as between 0.2 and 10 seconds per pixel on a serial minicomputer (i.e. for the 527x422 pixel DEM these timings correspond to from 0.5 to 25 DAYS).

Dowman (1987) gave a recent summary of attempts at stereo matching of 10m panchromatic SPOT image data. This review includes discussion of results from Geospectra (see **Vincent, 1987**) with rmse of 18.4m (9 check points) which took, for an unknown scene size, 8 hours (on a Vax11/750) for 90% of all points; Digim/EMR Canada which reported 60-82% of all points matched using a hierarchical correlation scheme which took 30hrs on a Vax11/780 for a quarter scene (3000x3000); **Guichard (1987)** who used an epipolar-based dynamic programming technique which gave rmse of 4m but whose run-time was reputed to be 3-4 hours (Vax11/780) for 250x250 image windows.

Ley (1987) reported accuracies for the MDA system of rmse between 0.6-2 pixels for 16 different 100x100 patches (again no timings were available for the MDA figures nor were results on disparities, disparity ranges or percentages of feature pixels extracted or correlated). It should be noted that the systematic errors and the rsme figures quoted for the MDA system in **Ley (1987)** may result from the MDA system use of epipolar pairs (see **Otto, 1988** for a detailed discussion of this point).

Finally, the NASA Goddard Space Flight Centre recently published results (see **Ramapriyan et al., 1986**) on the automated stereo matching of SIR-B SAR data using the SIMD (128x128 PEs with 4K RAM/PE) Massively Parallel Processor (MPP, see **Batcher, 1980**). The authors showed that for 25m SAR pixels with an image patch size of 512x512, the MPP produces disparity estimates using a hierarchical warp stereo technique (see **Quam, 1984**) in just 7 seconds. However, the reported quality of the results (elevation differences up to 1 km) is poor in comparison with the predicted rmse of 120m (**loc.cit.**). This poor accuracy is probably due to the input 6-bit data which had a poor signal-to-noise ratio because of multiplicative noise (speckle) and the lack of sufficient GCFs in the scene for the given exterior orientation.

3 Computational Stereo matching

3.1 Algorithmic classes

Barnard and Fischler (1982) group stereo matching into two main categories :

1) area (patch)-based techniques which rely upon the concept of smooth surfaces so that adjacent pixels will generally represent contiguous points in space;

2) feature-based techniques which are futher sub-divided into :

2.1) semantic features which have known physical properties and/or spatial geometry. These can be sub-divided further into :

- generic (e.g. vertices of linear structures, occlusion edges and prominent surface markings);
- domain-specific types (e.g. corner of building, road surface marking)

2.2) intensity anomaly features which include zero crossings of the Difference of Gaussian type as well as anomalous image patches (e.g. bright corner of field)

The first approach is generally what has been tried over the years in online aerial photogrammetric applications where sub-pixel acuity and density of features is most important whilst the second way has been favoured by online robotic and ALV vision systems owing to their lesser demands for high accuracy or sampling density and their greater demands for speed.

The area-based technique has, in recent years, been enhanced by two significant developments :
a) development of iterative adjustment of the foreshortening correction (e.g. **Gruen, 1985**)
b) hierarchical coarse-to-fine strategies to reduce calculation of correlations in areas with insufficient contrast and texture (e.g. **Quam, 1984**; **Hannah, 1985**)

Attempts to combine different techniques from the two cultures have been extremely limited to date. In addition, attempts to implement stereo matching algorithms at more realistic speeds have been constrained to the development of special-purpose hardware for feature extraction (an ideal pipelined computational problem) and/or the use of SIMD parallel computers to providing fast solutions for algorithmic solutions which make use of simple mappings of pixels to PE.

3.2 Architecture of stereo matchers

An alternative classification scheme to that proposed by **Barnard and Fischler (1982)** considers stereo matching in terms of the generic operators, their type and the order in which they occur. We will adopt this latter scheme here. Details of speed assessment, disparity range and density and modifications will be given in section 5.

We have chosen three stereo matchers out of the many hundreds proposed in the literature as they appear to represent typical examples of the three types :
(a) intensity anomalies (**Barnard and Thompson, 1980**);
(b) adaptive least squares correlation (**Gruen, 1985**)
(c) semantic edge features (**Pollard et al., 1985**);

They can be characterised as strongly non-epipolar (very large search-space), quasi-epipolar (moderate search-space), strictly epipolar (small search-space).

<u>B&T (Barnard and Thompson, 1980)</u>
This algorithm was originally developed at the University of Minnesota, USA. It doesnt require epipolar data and so is generally applicable to most stereo imagery which contains some sort of intensity anomalies. **Gennery (1980)** used the output of their interest operator (with a different matching scheme) to estimate the relative orientation of cameras in a robotic vehicle.

It consists of four sequential phases :
1) features are extracted using the Moravec interest operator (see

discussion in **Foerstner and Guelch, 1987**);
2) a local match sub-network is generated for each feature to be matched, which contains all possible matches for that feature given previously defined x- and y-disparity limits;
3) the match network is iteratively updated ten times to find the most likely match for each feature point using similarity between raw grey levels in a 5x5 window centred on each interesting feature;
4) correspondences are selected by thresholding of their probabilities, those lying outside certain thresholds are discarded as unreliable.

The first phase of this algorithm may be considered as a low-level data-independent pre-processing operation whilst the other phases involving feature matching are likely to be highly data-dependent. Output accuracies are not better than one pixel without an additional stage to construct a correlation surface around each feature point and interpolate a sub-pixel position (see **Foerstner and Guelch, 1987**). This sub-pixel precision can, however, be accomplished in a post-processing stage using the **Gruen(1985)** algorithm. In section 4 we consider how this algorithm was parallelised on a transputer array (**Collins et al, 1987a,b; Collins and Roberts, 1988**) and present recent results on an assessment of speed.

PMF (Pollard et al., 1985)

The PMF stereo matching algorithm (**Pollard et al., 1985; Pollard, 1987**) developed at the A.I. Vision Research Unit at Sheffield University, U.K. (partly under an Alvey IKBS-025) is an edgel matcher which relies upon four key constraints :
a) epipolarity of scanlines;
b) disparity gradient limit for support (derived from pschophysical observations, see **Pollard, 1987**)
c) support is only exchanged between matches over neighbourhoods

PMF consists of the following four phases :
1) extract edges (this is only briefly discussed in the original paper). **Pollard (1987)** suggests using a sub-pixel Canny operator (see **Canny, 1986**);
2) match edgels within given disparity limit along scanline;
3) match strength update using the disparity gradient (usually set at 1);
4) relaxation labelling based on match strength where relaxation iterations are given

The first phase of this algorithm is, once again, a data-independent operation whilst the other three phases can often be more efficiently implemented in a data-dependent fashion. In addition, the first phase needs to take place as a preprocessing phase whilst the other phases can occur in parallel for different scanline blocks of the images. Output accuracies are claimed by the Sheffield group to be sub-pixel depending on the application (see results for SPOT data in **Day and Muller, 1988**). **Chau (1987)** describes our independent implementation of their algorithm.

Gruen (1985)

The original adaptive least-squares correlation algorithm (**Gruen, 1985**; **Gruen and Baltsavias, 1986**) is a multi-parameter optimisation solution for minimising the sum-of-the-squares-of-the-differences between two image patches. It works in four phases :

1) resample the right image patch transformed by the affine transformation shaping parameters using a bi-linear interpolator;

2) find the grey-level derivatives and form the "design matrix";

3) find the column vector, L, each of whose elements is the difference in grey-levels for the appropriate point in the two image patches;

4) solve the over-constrained linear equations using Cholesky decomposition (see, for example, **Stoer and Bulirsch, 1980**)

This algorithm appears at first sight to be highly data-independent as different image patches can be matched completely independently of one another. However, closer scrutiny shows that the disparity sheet-growing version of this algorithm (see **Chau and Otto, 1988; Otto and Chau, 1988**) is data-dependent as it depends heavily on the order of patch processing. Furthermore, tests for correlation accuracy from, for instance, signal-to-noise variance analysis (see **Ehlers and Welch, 1987**) will increase the algorithmic data-dependence. Output accuracies are sub-pixel (see **Gruen, 1987; Day and Muller, 1988**).

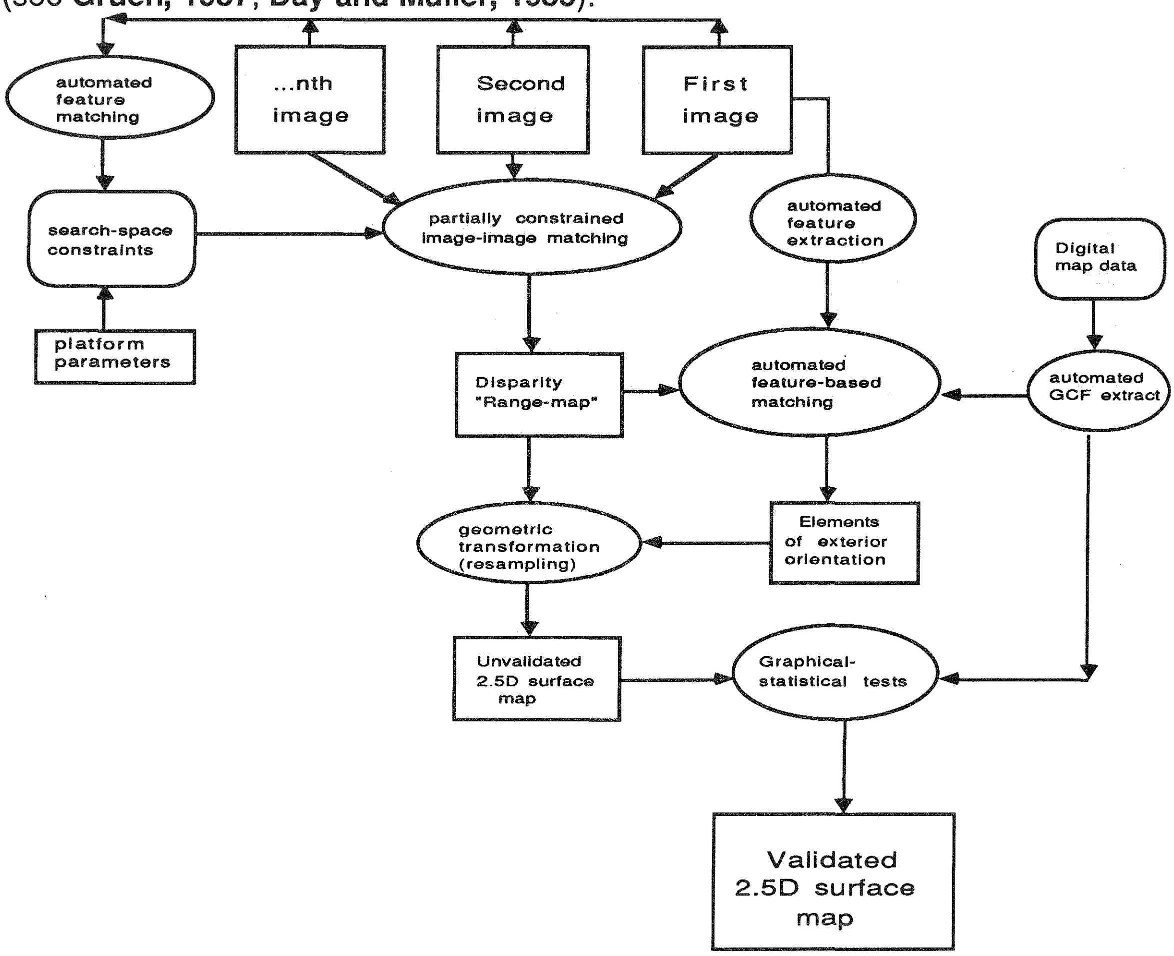

Figure 1 : Schematic diagram of system components for automated SPOT stereo matching

3.3 Integrating different stereo matchers

Figure 1 shows a schematic diagram of the complete system for the production of DEM data from SPOT. In this work, we will only be concerned with describing the parts of the whole process concerned with stereo (multiple) image-image matching. The reader is directed to the other papers previously cited on the Alvey MMI-137 project for further details on other parts.

Two different conceptual schemes for stereo matching have been investigated for SPOT stereo data (other applications are not discussed here). These will be referred to as Scheme A (iterative epipolar resampling) and Scheme B (region-growing Gruen).

For Scheme A, the following sequence of events is envisaged :
1) input raw SPOT images together with header information;
2) calculate approximate x,y disparity ranges from header information;
3) apply B&T with x,y disparity ranges and prune number of matches heavily;
4) calculate relative orientation parameters;
5) resample images to approximate epipolars;
6) apply PMF to obtain edge-based disparities;
7) apply Gruen using PMF edges as seed-points to fill in gaps between edges;
8) apply post-processing to remove blunders based on environmental constraints (see **Grieve and Simard, 1984**);
9) project disparities through absolute orientation model (see **Gugan, 1987; O'Neill and Dowman, 1988**);
10) triangulate using Delauney tessellation (see **McCullagh and Ross 1980**) to obtain either interpolated contours or grid at whatever desired spacing.

The problem with this approach is that we can never obtain perfect epipolars without first having a DEM (see discussion in **Otto, 1988**; **O'Neill and Dowman, 1988**) and hence we need to add at least one further iteration of stages 4 to 7.

For Scheme B, the following alternative sequence of events is envisaged :
1) input raw SPOT images together with header information;
2) calculate approximate x,y disparity ranges from header information;
3) apply feature point-based algorithm with x,y disparity ranges and other geometrical constraints to get a set of consistent matches;
4) use feature points as seed disparities for a region-growing variant of Gruen;
5) apply post-processing to remove blunders based on environmental constraints (see **Grieve and Simard, 1984**);
6) project disparities through absolute orientation model (see **Gugan, 1987**);
7) triangulate using Delauney tessellation (see **McCullagh and Ross,**

1980) to obtain either interpolated contours or grid at whatever desired spacing.

In this case, no (multiple) resampling of the whole image is performed and this scheme appears to be flexible enough to handle almost any applications area.

Estimates of the disparity ranges required for automatically stereo matching SPOT data are a useful aid to assessing the complexity of any single algorithmic approach as well as being a helpful guide to predictions of run-times on arrays of transputer elements. For the B/H ranges expected for SPOT, rotation between images and uncertainties in camera pointing lead to disparity estimates of between 160-260 pixels for x and around 100 pixels for y for expected height ranges within the scene of 1000m (it will be up to ten times larger for rugged areas such as the Himalayas). These disparity ranges are several orders of magnitude larger than the "pull-in" range for most stereo matchers developed to date. **Day and Muller (1988)** show results from our automated processing scheme which can support such vast disparity ranges.

4 Role of special-purpose hardware in image processing

4.1 Display and array processors

Since the early 1960s, there have been a number of developments in special-purpose hardware for image processing. These devices include Look-up Tables (LUTs) for grey-scale manipulation; pixel shifters for scrolling and panning; pixel replicators for zooming and Arithmetic Logic Units (ALUs) for simple bit or byte-level arithmetic. They can all be characterised by their high cost, difficulties in programming and low production rates and their emphasis on matching processing-rates to integer multiples of the video refresh-rates. Recently graphics workstations have caught up with their processing speeds and appear to have comparable facilities at drastically reduced cost.

These image processing "display processors" can all be characterised by their small pipelines which allow global tasks to be performed at video refresh rates (e.g. histogram calculation). However, although many of their ALUs use internal 16-bit calculations, their output is frequently compressed to 12-bits or less. Their role, if any, is limited in higher-level image processing tasks where floating-point and non-raster data-structures are common.

Several years ago, scientists and engineers turned to special-purpose programmable (floating-point) array processors for specialised operations such as convolutions and unitary transforms (e.g. FFTs). **Foerstner and Guelch (1987)** describes the theoretical application of one of these add-on boards, called a Sky Warrior for low-level feature extraction and demonstrated that significant speed-ups should be possible. However, their floating-point performance is easily outstripped by most current Reduced Instruction Set (RISC) microprocessors such as the Weitek chip-set used in FPS T-series (see **Frenkel, 1986**).

4.2 DSP VLSI chips

The increasing role of two-dimensional binary image processing in industrial inspection (particularly in PCBs and lithography) has led to a concomitant increase in the availability of specialised Digital Signal Processing chips for specialised tasks such as small-area convolutions.

Table 1 shows a comparison of timings for a 3x3 spatial filter on 256x256x8-bit input images. Inspection of this table shows that the more specialised the design of the VLSI chip, the faster the speed. However, these are the maximum speeds for the special-purpose devices whereas the transputer speeds can be increased incrementally by increasing their number. Furthermore the high processing speed of special-purpose silicon comes as a result of their inflexibility due to their lack of programmability.

The difficulties in interfacing many of these VLSI devices to parallel machines may also exclude their usage as pre-processing elements, although in one case cited here (Inmos A100 and the transputer) they have been designed as an integral chip-set. It is likely to be many years yet until special-purpose VLSI chips are able to address the accuracy and sampling density requirements of stereo matching for photogrammetric applications although some progress has been reached in special purpose circuits for edge-based matching in very highly epipolar machine vision systems (see **Ohta et al., 1987**). In the meantime, the majority of effort should be directed towards finding computational solutions for stereo matching using SIMD and MIMD parallel machines.

TABLE 1: Estimated performance for edge detection (256x256x8-bit input with a 3x3 spatial filter)

Name	Description	h/w specification	estimated performance
Toshiba	Programmable DSP	10Mhz,32-bit fixed point	65.5ms
INMOS A100	Cascadable Transversal Filter	32 multiply/accumulators 16x4 bits at 2.5Mhz 16x16 bits at 10Mhz	8ms
Plessey PDSP16401	2-D Video Signal Processor	13-bit edge magnitude 3-bit edge direction	4ms
Mil-DAP	Fine Grain SIMD Machine	32x32 1-bit PEs connected in mesh	2.9ms
INMOS 16xT414	Coarse Grain MIMD building block	15Mhz, 7.5MIPS 32-bit fixed point 2Kbytes RAM, 4 x 10Mbps serial links	144ms

4.3 SIMD engines for image processing

There have been a number of attempts to build SIMD engines optimised for low-level image processing (see review by **Fountain, 1986**), including the Massively Parallel Processor (MPP) at the NASA Goddard Space Flight Centre (**Batcher, 1980**); the Cellular Logic Image Processor (CLIP) series at University College London (**Duff, 1983**) and the Distributed Array Processor (DAP) at ICL (**Reddaway, 1973**). Their characteristics are given in Table 2.

Table 2 : SIMD systems for low-level image processing (after **Fountain, 1986**)

System	DAP	CLIP4	MPP
Originators	ICL	UCL	NASA
Manufacturers	ICL	Stonefield	Goodyear
Date	1980	1980	1983
Array size	64x64	96x96	128x128
Memory per PE (bits)	4k	32	1k
Language	DAP-FORTRAN	IPC	Parallel PASCAL
Host	ICL 2900	PDP-11/34	Vax-11/780
Current system	DAP3	CLIP7	-
3x3 erosion 10^9 results/sec (see **Gerritsen, 1983**)	2.276	0.369	8.192

Inspection of this table shows that array sizes are smaller than most CCD or satellite images to be processed so that memory per PE is one of the chief factors in assessing their utility for higher-level processing such as stereo matching. Results have only been published to date on the use of a data-independent stereo matching algorithm on the MPP (see **Rampriyan et al., 1986**).

Gerritsen (1983) lists a number of advantages of SIMD machines for image processing :
1) no special addressing is required to access nearest neighbours;
2) local memory storage in PEs eliminates memory fetch overheads;
3) bit-oriented PEs offer great flexibility in the precision and repesentation of data.

SIMD machines also have a number of disadvantages including :
4) limited local memories restrict the class and breadth of algorithms which can be implemented;
5) table look-ups are not efficient on such machines;
6) certain classes of image processing operations (e.g. geometric transformations) are not efficient on such machines;
7) although other architectures can be accommodated (see **Fountain, 1987**; **Forshaw, 1987**) this appears to require substantial re-design work to be done;
8) algorithms which appear to be best suited are data-independent.

Hypercube architectures were prompted by the Cosmic Cube machine at Caltech and have led to the development of massively parallel SIMD machines which can be used for image processing. The most well-publicised of these originated from work at the MIT AI Laboratory (see **Hillis, 1985**). With up to 64k processors in a single machine, redundancy can be skillfully used (see **Little, 1986**) to handle a variety of medium-level image processing tasks at high speeds but at a comparatively high cost (see discussion in final section).

5 Stereo matching on transputer arrays

5.1 Transputers for image processing

The INMOS transputer is a VLSI chip using CMOS 1.5μ silicon fabrication technology containing a 32-bit processor, 2-4K of on-chip RAM and four bidirectional asynchronous communication links with a capacity of 20 Mbps. There are currently two primary versions of this chip, the first is a 32-bit integer-only processor (T414, see Table 1) whilst the second one is a floating-point chip with some 10 Mips of integer performance and 1.5Mflops of floating-point (see Figure 2 for a schematic diagram of this T800 floating-point transputer). It requires only very minimal glue logic to fabricate large arrays of such chips and given suitable switching circuitry such parallel arrays can be configured (statically or dynamically) in any 4-connected topology (see **Harp et al., 1987**).

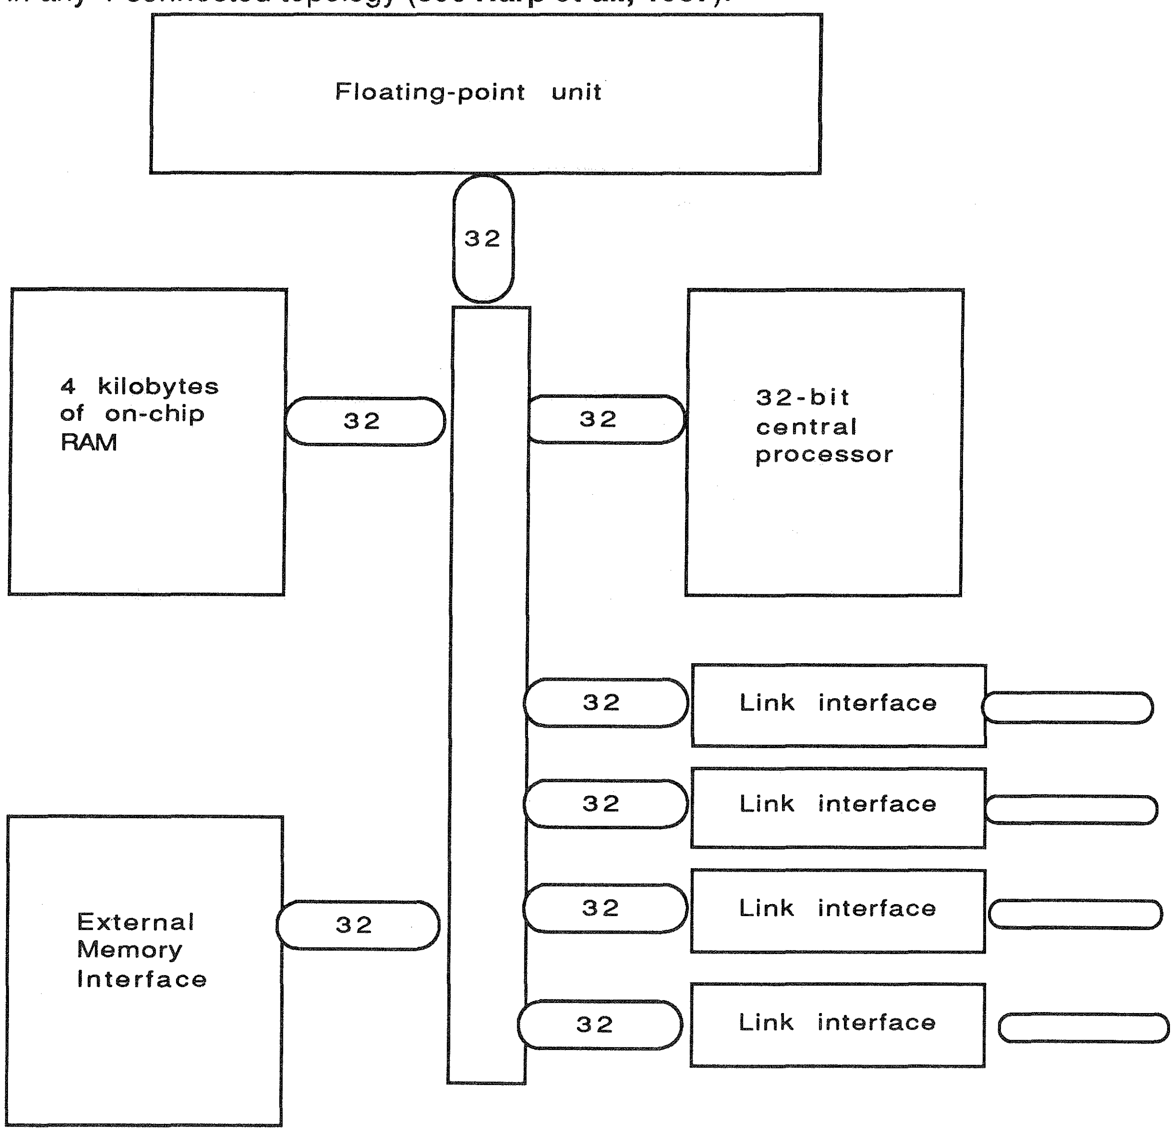

Figure 2 : Schematic architecture of T800 floating-point transputer

Several commercial transputer arrays are available including one from

Meiko™ called a Computing Surface™ (see **Bowler et al., 1987** for a description of this system and its evaluation for numerical simulation experiments). A 20-worker transputer version of this machine (currently with T414 chips) was installed several years ago at RSRE Malvern and a schematic diagram of its configuration is shown in Figure 3. This machine is attached to a Sun-3 workstation which has recently been networked using a 64 Kbps point-to-point link to a Sun-3 workstation at University College London. In this way, experiments can be carried out remotely on this parallel machine and results displayed on workstations some 150 kms distant.

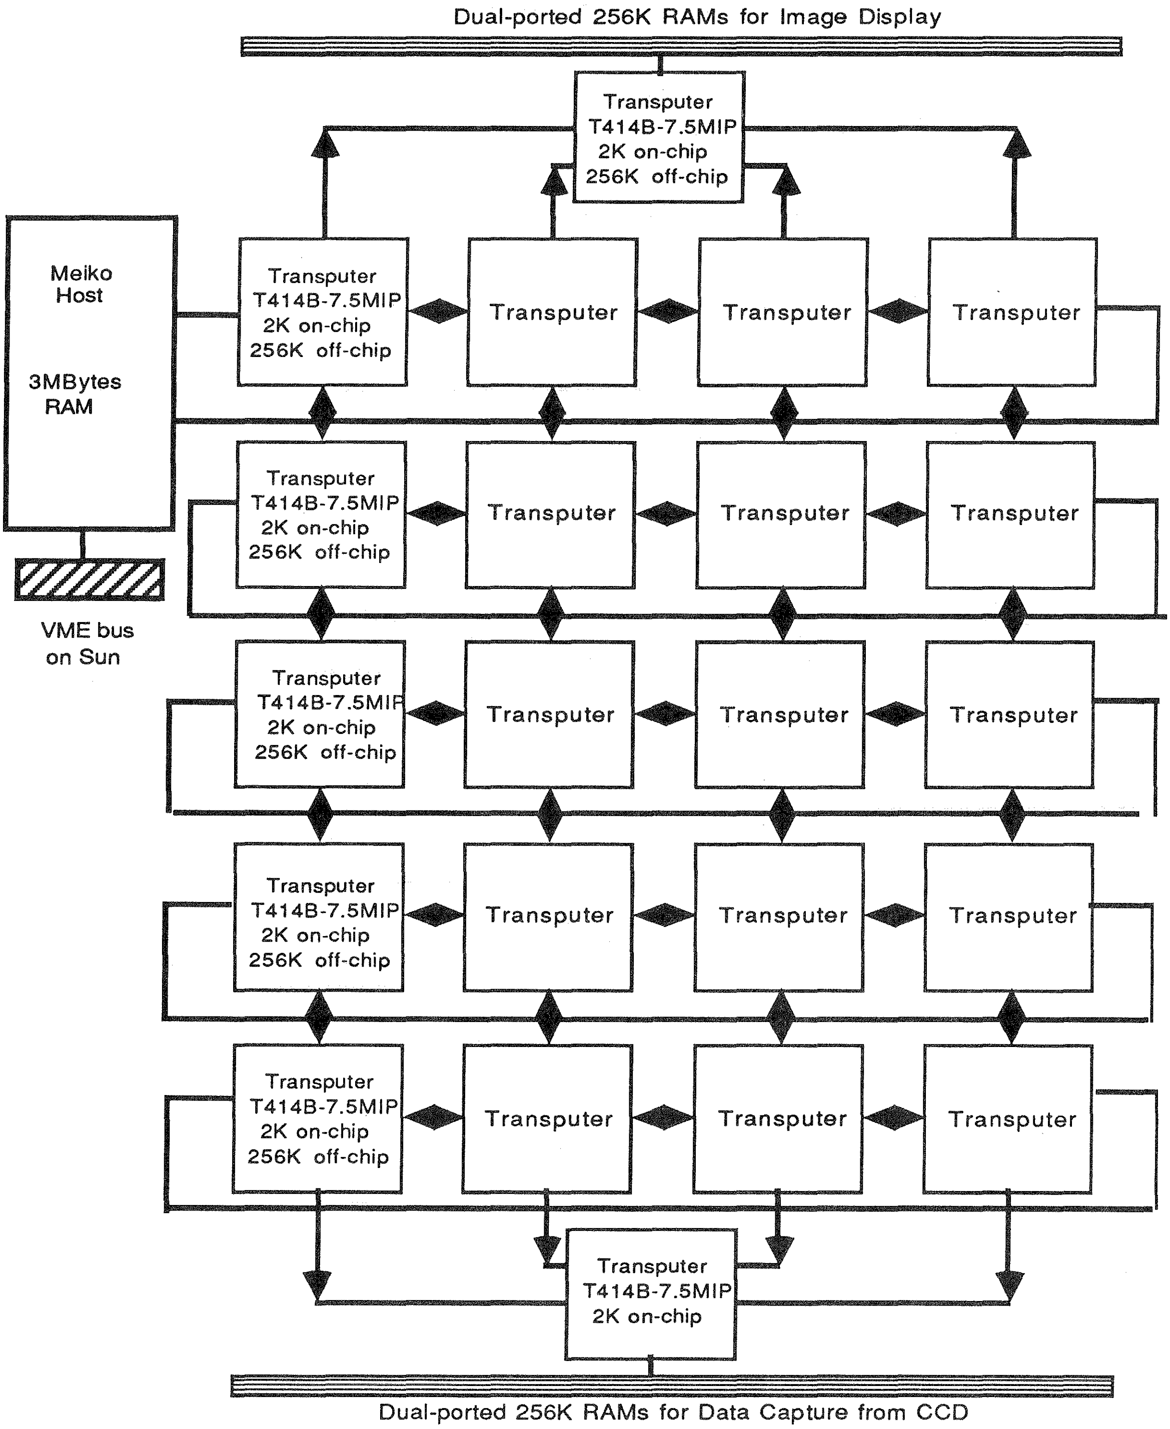

Figure 3 : Schematic diagram of Meiko transputer array at RSRE Malvern

Such transputer-based systems may be programmed in occam (see **Hoare, 1984**) and conventional programming languages (with and without occam harness) such as C, Fortran77, and Pascal. Run-time environments include several versions of Unix™ lookalikes, a Transputer Development System (TDS) based on a novel "folding editor" and an increasingly larger number of diagnostic and debugging aids based on graphical representations.

Studies involving the applications of transputers in image processing have tended to concentrate on evaluating their performance for low-level image processing tasks (see **Harp et al., 1987**; **Manning et al., 1987**), although **Bowler et al. (1987)** showed results from a simulated annealing-based image restoration experiment.

There are several strong reasons for choosing transputer elements in preference to either special-purpose hardware or SIMD engines for high-level image processing tasks such as stereo matching.

These advantages include the following :
1) data-dependent algorithms can be more readily mapped onto a transputer array than an SIMD system;
2) different load-balancing strategies on transputers can be rapidly tested by reconfiguring the network using silicon switches (see **Harp et al., 1987**);
3) although pipelined special-purpose silicon is often faster than MIMD systems, the algorithms are difficult to design and have a strong relationship between the number of processors and problem size. This makes it difficult to map a range of algorithms onto a particular array of processors.
4) Special-purpose silicon usually has a low maximum precision (e.g. 12-bits for most silicon- see Table 1). Transputers can accommodate the full range of floating-point sizes (up to double precision) and communicate using bit-serial links.
5) transputer-based MIMD systems are extensible so that if there is demand for greater computational power, this can be provided by increasing the number of PEs. Only a few SIMD systems will allow this type of extensibility.
6) their cost/performance ratio is very high (e.g. for 1 Megaflop, a Cray X-MP/48 costs $25,000 whilst a T800 transputer costs $1500), see **Bowler et al. (1987)** and **Gelberg and Stephenson (1987)** for recent analyses.

Disadvantages of transputers include :
1) they are more costly than single-bit PEs used in SIMD for low-level image processing;
2) the bit-serial link-speed is a bottleneck for certain types of algorithmic parallelism;
3) the on-chip memory is small (only 4k maximum) and the factor of three degradation in speed to access off-chip memory can be restrictive.

16MB/PE is the current maximum offered commercially.;
4) they have been very difficult to program until recently because of the prevalence of TDS as a runtime environment (see discussion of programming issues in **McGraw and Axelrod, 1988**). The existence of several parallel-C compilers is beginning to tackle this problem;
5) I/O has been slow by being limited to at most three links. This has recently been improved through the development of high-speed I/O devices (>200 Mbps) with up to 64 links.

5.2 Application of transputers for stereo matching

Types of parallelism and connectivities

Three types of parallelism are usually discussed by users of MIMD systems (see **Bowler et al., 1987**) :
1) geometric (spatial) parallelism in which subsets of data are allocated to each PE and inter-PE communications usually only occurs for data at boundaries;
2) algorithmic parallelism where different parts of the processor array perform different tasks and data is "flowed" through the system;
3) task (event) parallelism where a "master" PE "farms" out work to "slave/worker" PEs.

In addition to these different types of control for the array, the PEs may be configured in any 4-connected topology with the addition of suitable silicon switches.

Topologies relevant to stereo matching include :
1) chain-connected (this is the simplest and the one discussed later on in this section);
2) mesh-connected (nearest North-East-West-South neighbours) for geometric parallelism;
3) pyramids (if sufficient memory is present) for coarse-to-fine matching;
4) trees for cascaded "farming" architectures.

Metrics for performance evaluation

There are a number of possible criteria which may be used to assess speed-ups from the use of a transputer array for stereo matching. These include :
1) comparison of communications time vs. computation time;
2) processor utilisation (see nest sub-section for an example of this);
3) memory utilisation efficiency;
4) efficiency (usually defined as the speed-up factor on N transputers);
5) I/O time;
6) throughput (output disparities to input pixels);
7) processing speed per output disparity value.

A few of these metrics are summarised in Table 3 below.

Table 3 : Comparison of stereo matching speeds on different machine

types

System	x-disp	y-disp	Input	Throughput	Speed	rms	Reference
	search space		image size	o/p+i/p	o/p/sec	pixels	
Minicomputers							
Kern	?	0	512x480	1/245760	7	~1	Li(1986)
Symbolic	50	0	512x512	1	43,200	~1	Barnard(1987)
ATOM 3.0	128	0	?	?	(1)	?	ATOM™ (1986)
CCRS	4	~0	?	1/16	?	~0.1	Grieve&Simard (1984)
MDA	20	0	512x512	3/100	?	~0.4	Cooper et al(1987)
MDA	?	0	100x100	1/6.25	?	0.6-2.2	Ley(1987)
Georgia	10	2	527x422	1	0.2-10	0.3	Ehlers & Welch (1987)
Special-purpose VLSI							
PRISM	200	2	512x512	0.004	23	0.1	Nishihara(1984)
Ohta	?	0	256x256	0.02	520	~1	Ohta et al(1987)
SIMD							
CM	41	0	256x256	~0.01	0.002(2)	1-2	Drumheller&Poggio(1986)
MPP	148	0	512x512	1	0.00016(3)	8?	Ramapriyan et al (1986)
MIMD							
11 Sun-3	260(4)	100	1400**2	0.014	0.01	0.1	Otto & Chau (1988)
Meiko-20	45	17	240x240	0.0026	0.025	~1	Collins&Roberts(1988)

(1) quoted as 800 point elevations/second for x-disparity range of +65 on Vax11/750
(2) quoted as 1.5 seconds for a guesstimate of 1% output density
(3) quoted as 7 seconds for a estimated output of the same size as input
(4) any disparity range can be handled provided at least one initial seed-point is given to within 2 pixels disparity

5.3 Barnard & Thompson : Speed assessment on transputer array

In considering the design of a parallel implementation of this algorithm, it is clear that the feature extraction stage involves predetermined separable operations on pixels suited to simple geometric partitioning of the data. The processing requirements of the feature matching stage (phases 2, 3 and 4 - see discussion above) are dependent upon the images and input parameters used. These produce a unique spatial distribution of feature points which presents a load balancing problem. The strategies adopted, both of which use a linear chain of Transputers, are described below.

Geometric Parallelism

This divides the image into non-overlapping horizontal strips, each strip consisting of several scanlines. Each processor is allocated the corresponding strips for the image pair, such that the first in the chain receives the top strip of each image, and so on for each processor in the chain. This strategy maintains the position of pixels within the image relative to the position of the processor within the chain. Processor communications, used to access any data not resident in the local processor, are restricted to nearest neighbour by ensuring the number of rasters ineach processor is greater than, or equal to, both the maximum vertical disparity and the radius of the update neighbourhood.

Feature Balancing

This uses a control processor to decide which features should be resident on each processor, and to perform simple buffering operations. The crucial design criterion used here was to ensure that the search operations required during the construction and update of the local match networks are distributed across the processing network. Initially input images are split into overlapping blocks and allocated to a processor in the chain. The returned data consists of a simple list of x,y feature point positions and 5 by 5 pixel values centred on x,y. The controller redistributes the data from Phase 1 by farming out roughly equal numbers of feature points from the left image to each processor in the chain.

Next, the controller broadcasts all right image feature points to the chain. Each processor buffers a right image feature if it lies within the disparity space of one or more of the left image features held locally. Phase 2 is completed by each processor calculating the initial match probabilities , and forwarding the results to the controller. During phase 3, the controller broadcasts the probabilities at the start of each iteration, and waits to collect an updated set. Each processor buffers the results from the previous iteration that are required for updating its local match networks, then performs the next update operation using the locally held data, then forwards the updated results to the controller. During phase 4, the match probabilities are thresholded to select the output matches, which are then returned to the controller.

Workload Balancing

This technique applies knowledge of the algorithm complexity to distribute the relaxation workload. In this case the computational complexity of theupdating operation for each left image feature point is 0(n3) where nis the number of its candidate right image features. The controller uses this factor to redistribute the left image features to processors prior to the broadcast of right image feature points.

Figure 4 shows processor utilisation profiles obtained for six processing runs made with the three different strategies for sharing the processing load between 20 worker T414 15MHz Transputers, with and without pruning out low probability (<1%) matches before each of the ten relaxation iterations. In every case, the final results were almost identical, approximately 150 matches being found from an original (left image) set of 324 feature points (see Table 4). A display of the location of the B&T points on the SPOT sub-image can be found in **Collins and Roberts (1987)**.

The processor utilisation profiles are one of many different attempts to monitor the performance of transputer arrays using a graphical representation of transputer activity (in this case how load-balanced the array is). The x-axis represents time and there are three time series shown for each plot. The y-axis is the processor number in the chain (from 0 to 19) whilst the grey-scale value is the processor utilisation expreseed as a percentage (white indicates 100% utilisation and black that the processor is idle). This graph is obtained by executing a supervisor process on each Transputer which samples how busy the processor has been approximately every 100 ms. The network construction phase and

10 relaxation cycles are clearly visible. The time axis is split into 3 adjacent sections, each 25s long. Figure 4 only shows procesor activity for phases 2-4.

Two main points emerge from Figure 4. Firstly, the pruning of low probability matches drastically speeds up the processing without any loss of integrity. Secondly, the highest utilisation and shortest run times are obtained with the workload balancing strategy, allocating estimated workload 0(n3) for n candidate match points, rather than equalising the allocation of reference feature points and considerably better than the simple geometric partitioning of image space. Table 4 confirms these graphical impressions quantitatively. In addition to the approximately 20% speedup due to optimised load balancing, the added flexibility, the ease of programming and more efficient memory utilisation makes this approach a clear winner when image features are densely distributed.

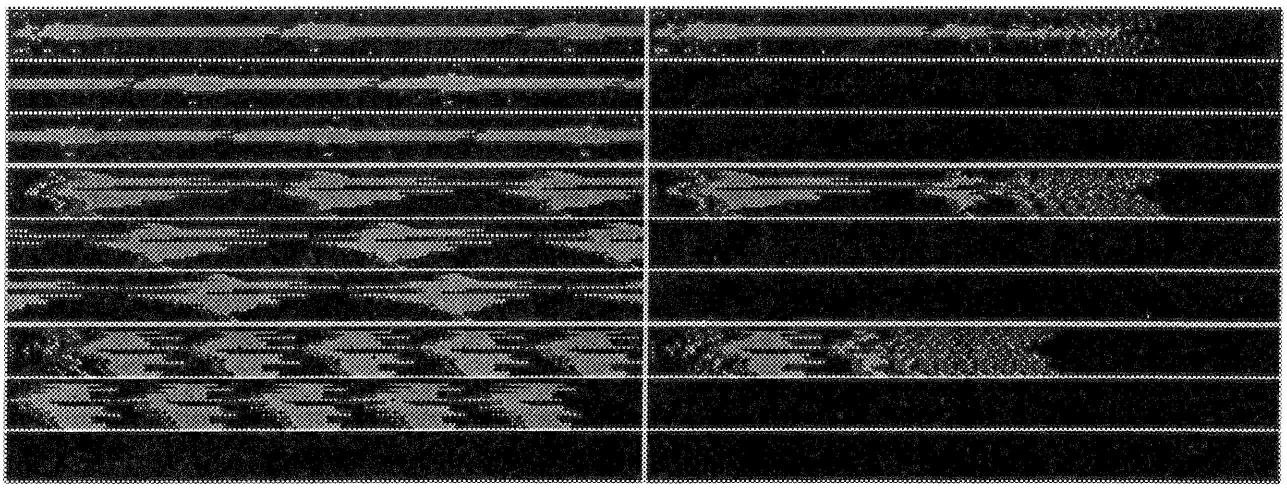

Figure 4 : Processor utilisation profiles for Meiko 20 transputer array
Upper panel : geometric parallelism without (left) and with 1% pruning (right)
Middle panel : feature balanced without (left) and with 1% pruning (right)
Lower panel : workload balanced without (left) and with 1% pruning (right)

Table 4 : Feature point timings in seconds for input images of 240x240x8-bit

Pruning Level (%probability)	Geometric	Feature	Workload
0	105.759	97.192	52.637
1	22.315	22.229	17.900

N.B.
Disparity limits: -35 < x < 10 and -11 < y < 6
Update radius = 11
Moravec operator took only 2.7 seconds and can therefore be considered a negligible part of the total processing time.

It is clear that for the case in which the spatial density of feature points is dense, the overheads involved in setting up and distributing work packages are fully justified and load balanced process farming is the best strategy. In contrast, if the spatial density of features is very low, simple geometric parallelism is slightly faster. However, unless the small speed advantage is critical (eg in a truly real-time application) we prefer process farming because it is easier to program (eg no explicit synchronisations are needed to avoid deadlock) and because it requires less compile-time tailoring of the program to specific parameters.

5.4 PMF : Theoretical Speed assessment

Our independent implementations of PMF (which includes several important changes) run on a Sun workstation in the C programming language (see **Chau, 1987**). The UNIX™ operating system permits timing profiles to be generated frequently without resource to hand-coding. **Recce et al. (1987)** discuss the speed assessment experiments performed on PMF on one of our Sun workstations.

5.5 Gruen:Theoretical Speed assessment and results for LAN of workstations

The speed of the algorithm depends partly on the implementation, and partly on the parameters used (e.g. patchsize and grid spacing). Since we are still in the process of evaluating and tuning it for real data, we cannot yet give a definitive answer. Furthermore, space precludes a full discussion. Thus, this section will just outline our knowledge so far (see **Otto and Chau, 1988** for more details).

On a single SUN 3/180, using a MC68881 floating point coprocessor, our current implementation takes just under 1/2 sec to do one iteration of Gruen's algorithm, using a patch size of 21 x 21, or about 1/4 sec for a patch size of 15 x 15. (The time is approximately proportional to the area of the patches used.) Recoding to use integer arithmetic wherever possible would probably speed this up by a factor of about 4. However, a 20 MHz T800 Transputer (1-1.5Mflops) should be able to reduce these times to about 0.06 and 0.03 seconds respectively.

For any given pair of patches, these times need to be multiplied by the number of iterations required to get accurate convergence - our initial experiments indicate that typically only one or two iterations are needed when Gruen's algorithm is used as part of the region-growing algorithm.

Overall, then, matching one pair of patches on a Transputer will typically take between 0.05s and 0.2s, depending on the parameters used, and the image characteristics.

Since Gruen's algorithm requires significant computation (at least 50 ms worth) on comparatively little data (a few kbytes), an easy and efficient way of parallelising the region-growing algorithm is to have one central "master", which manages the priority queue & the image data, and many "workers", which just apply Gruen's algorithm to whatever patches they are given, and return their results to the "master". It is efficient because all of the "workers" are CPU-bound, since Transputer links are fast enough to transfer the data faster than the CPU's can process it. Furthermore, the "workers" do not need much memory - 100-200kbytes is ample.

If we have enough "workers", then the one "master" will become a bottleneck. Initial analysis has shown that this will not happen until there are at least 30 "workers", and probably more. Even then, it would be relatively straight-forward to partition the master over two or more processors.

We have already parallelised our algorithm on a network of SUN workstations, and speedup is, as predicted, linear up to 15 workstations (which was all we could get our hands on up until now). A Transputer implementation is just about to begin at the time of writing.

Using 30 T800 Transputers, and the algorithm above, it should be possible to produce a high-quality (see Table 3; **Otto and Chau, 1988; Day and Muller, 1988**), dense disparity map from a pair of SPOT images in about 2 hours. (The exact time depends upon the image chacteristics and the algorithm parameters; this time is a "best guess", not a "most optimistic guess").

7 Discussion and Conclusions

The initial implementation of a data-dependent stereo matching algorithm on a transputer array has demonstrated the flexibility of MIMD machines based on transputers for prototyping real-time stereo matcher exemplars.

Their utility for SPOT data processing where huge search spaces need to be accommodated has been demonstrated and it is difficult to see how SIMD machines could address these problems (except possibly using coarse-to-fine strategies such as the PRISM system, see **Nishihara,1984**).

MIMD machines based on transputers appear to offer the best vehicle currently for developing real-time stereo matching machines because :
1) they can be easily assembled using minimal glue logic;
2) they are potentially upgradeable as new GaAs technologies emerge as Inmos has established pin-compatibility between different processors;
3) they are incrementally extensible so that larger numbers of PEs directly contribute to increasing the speed of processing;
4) they form adaptable networks of PEs so that algorithms and topologies may be tuned to the new problem domain;
5) they can be simply interfaced to high-speed I/O devices for entry of multiple cameras;
6) their cost/performance ratio for floating-point operations is very strong in comparison to other microprocessors.

A stereo matching exemplar, based on T800 transputers is currently being designed and built at Thorn EMI Central Research Laboratories and will soon be installed at UCL. This will enable tests to be carried out on prototypical environments for future real-time stereo matching.

At Thorn EMI another transputer exemplar will be used to capture stereo data from a Close Range Vision Cell using frame transfer CCD cameras. In 1986, a Kern DSR-11 based system (with a CCD-framestore system) was installed to enable the current stereo-SPOT image to be modified for applications using digitised photographs. Design work on interfacing this system to the stereo matching exemplar is currently being done.

Finally a supercomputer constructed from transputers being built at RSRE Malvern will also be used to test whether linear speed-ups can be obtained without re-programming. It is hoped that in these ways, the extensibility and adaptability of transputer arrays for stereo matching can be proven and prototypes developed for future use in a variety of different production environments.

Acknowledgements
This Alvey MMI-137 project is in collaboration with Laser-Scan Laboratories, Cambridge. The UK Alvey programme is funded by the Department of Trade of Industry, the Science and Engineering Research Council and the Ministry of Defence. We would like to thank Kevin Palmer (RSRE), Tim Day (UCL-P&S) and Tony Chau (UCL-CS) for all their help with this paper.

References

ATOM™ (1987)
Geospectra sales literature, Geospectra Corporation, P.O. Box 1387, 333 Parkland Plaza, Ann Arbor, Michigan 48106-1387, USA

Barnard, S.T. (1987)
"Stereo matching by hierarchical microcanonical annealing"*SRI Technical Note* 414, February 6, 1987, 14pp

Barnard, S.T., Thompson, W.B. (1980)
"Disparity analysis of images" *IEEE Trans. Patt. Anal. Mach. Intell.*, vol. PAMI-2(4), pp.333-340

Barnard, S.T., Fischler, M.A. (1982)
"Computational Stereo" *ACM Computing Surveys*, vol.14(4), pp.553-572

Barrow, H.G., Garvey, T.D., Kremers, J., Tenenbaum, J.M., Wolf, H.C. (1977)
"Interactive aids for cartography and photo-interpretation" *SRI Technical Note* 137, 46pp

Batcher, K. (1980)
"Design of a massively parallel processor" *IEEE Trans. Comput.* vol.C29-9, pp.836-840

Bernard, M., Boutaleb, A.K., Koelbl, O., Penis, C. (1986)
"Automatic stereophotogrammetry : implementation and comparison of classical correlation methods and dynamic programming-based techniques" *Proc. ISPRS Commission III Symposium* on "From analytical to digital" Rovaniemi, Finland, August 19-22. IAPRS, vol.26-3/1, 10pp

Bertram, S. (1969)
"The UNAMACE and the Automatic Photomapper" *Photogrammetric Engineering*, vol.35(6), pp.569-576

Duff, M.J.B. (1983)
"Computing structures for image processing" Edited by M.J.B. Duff, Academic Press.

Ehlers, M., Welch, R. (1987)
"Stereocorrelation of LANDSAT-TM images" *Photogrammetric Engineering and Remote Sensing*, vol.53(9), pp.1231-1237

Flynn, M.J. (1966)
"Very high-speed computing systems" *Proc. IEEE*, vol. 54, pp.1901-1909

Foerstner, W., Guelch, E. (1987)
"A fast operator for detection and precise location of distinct points, corners and centres of circular features" *Proc. Intercommission Conf.* on "Fast Processing of Photogrammetric Data" held at Interlaken, Switzerland, June 2-4, 1987, pp.281-317

Forshaw, M.R.B. (1987)
"Processor array requirements for advanced image processing : theory and experiment" *Proc. 3rd Alvey Vision Club*, Cambridge, 15-17 September, pp.131-137

Fountain, T.J. (1986)
"A review of SIMD architectures" in "Image Processing System architectures" Edited by J. Kitler & M.J.B.Duff, Research Studies Press.

Fountain, T.J. (1987)
"The development of array architectures embodying partial local autonomy" *Proc. 3rd Alvey Vision Club*, Cambridge, 15-17 September, pp.125-129

Frenkel, K.A. (1986)
"Evaluating two massively parallel machines" *Communications of the ACM*, vol.29(8), pp.752-758

Gelberg, L.M., Stephenson, T.P. (1987)
"Supercomputing and Graphics in the Earth and Planetray Sciences" *IEEE CG&A*, July 1987, 26-33pp

Gennery (1980)
"Modelling the environment of an exploring vehicle by means of stereo vision" *Stanford A.I. Laboratory Memo* AIM-339 (Ph.D. thesis)

Gennery, D.B., Wilcox, B. (1985)
"PIFEX: An advanced programmable pipelined-image processor" *NASA JPL Publication* 84-97, 42pp

Gennery, D.B., Litwin, T., Wilcox, B., Bon, B. (1987)
"Sensing and perception research for space telerobotics at the Jet Propulsion Laboratory" *Proc. IEEE Int. Conf. Robotics and Automation*, held at Raleigh, North Carolina, March 31-April 3, pp.311-317

Gerritsen, F.A. (1983)
"A comparison of the CLIP4, DAP, and MPP processor-array implementations" in Duff (1983)

Grieve, G.R., Simard, R. (1984)
"Digital Elevation Modelling : progress towards automatic procedures with LANDSAT-MSS geocoded data" *Proc. 9th Canadian Symposium on Remote Sensing* held in St. Johns, Newfoundland, Ottawa, pp.461-469

Gruen, A.W. (1985)
"Adaptive least squares correlation : A powerful image matching technique" *S.Afr.J. of Photogramm., Rem. Sens. and Cart.*, vol.14(3), pp.175-187

Gruen, A.W. (1987)
"Towards real-time photogrammetry" Invited Paper at the *41st Photogrammetric Week*, Stuttgart, September 14-19, 1987, 33pp

Gruen, A.W., Baltsavias, E.P. (1986)
"High precision image matching for Digital Terrain Model generation" *International Archives of Photogrammetry and Remote Sensing*, vol. 25(3):254-276

Gugan, D. J. (1987)
"Practical apsects of topographic mapping from SPOT imagery" *Photogrammetric Record*, vol.12, no.69, 349-355pp

Guichard (1987)
"Use of epipolar data for correlation" *Proc. Int. Conf.* on "SPOT1: Image utilisation, assessment, results", held in Paris, France, 23-27 November 1987, 10pp

Hannah, M.J. (1985)
"SRI's Baseline Stereo system" *Proc. Image Understanding Workshop*, Miami Beach, Florida, December 9-10, 7pp

Harp, J.G. , Palmer, K.J., Webber, H.C. (1987)
"Image processing on the Reconfigurable Transputer Processor" *Proc. Occam User Group, September 1987, Grenoble, France,* 9pp

Helava, U.V. (1987)
"Digital Comparator Correlator System" *Proc. Intercommission Conf.* on "Fast Processing of Photogrammetric Data" held at Interlaken, Switzerland, June 2-4, 1987, pp.404-418

Hillis, D. (1985)
"The Connection Machine" MIT Press, Cambridge, MA, USA

Hoare, C.A.R. (1984)
"Communicating Sequential Processes" Prentice-Hall, Englewod Cliffs, N.J.

Hwang, K. (1987)
"Advanced parallel processing with supercomputer architectures" *Proc. IEEE*, vol. 75(10), pp1348-1379

Kelly, R.E., McConnell, P.R.H., Mildenberger, S.J. (1977)
"The Gestalt Photomapping System" *Photogrammetric Engineering and Remote Sensing*, vol.43(11), pp.1407-1417

Ley (1987)
"Some aspects of height extraction from SPOT imagery" *Photogrammetric Record* (in press)

Li, M. (1986)
"A comparative test of DEM measurement using the Kern DSR-11 Analytical Plotter" *Internal Report*, Department of Photogrammetry, The Royal Institute of Technology, Stockholm

Little, J.J. (1986)
"Parallel algorithms for computer vision on the Connection Machine" *MIT AI Memo* 918, November 1986, 30pp

Manning, L.J., Dew, P.M., Wang, H. (1987)
"Programming models for VLSI array processors with application to low-level vision processing" *Proc. BCS Parallel Processing Group Workshop* on "Parallel architectures and computer vision", Oxford, UK, 30-31 March 1987, 15pp

McCullagh, M.J., Ross, C.G. (1980)
"Delauney triangulation of a random data set for isarithmic mapping" *Cartographic Journal*, vol.17(2), pp.93-99

McGraw, J.R., Axelrod, T.S. (1988)
"Exploiting multiprocessors: Issues and Options" in "Programming parallel processors" Edited by R.G. Babb II, Addison-Wesley.

Muller, J-P., Day, T., Dalton, M., Kolbusz, J., Pearson, J.C., Richards, S. (1988a)
"Visualisation of topographic data using video animation" *Proceedings of the XVIth International Congress of ISPRS*, Kyoto, Japan, *IAPRS* 27-A4

Muller, J-P., Dalton, N., Gugan, D., Morris, Peacegood, G., A.C., Stevens, A.C. (1988b)
"Map-image registration and automated image feature extraction" *Proceedings of the XVIth International Congress of ISPRS*, Kyoto, Japan, *IAPRS* 27-A2 (Supplement)

Muller, J-P., Paramananda, V. (1988)
"Graphics workstations for integrated data collection, simulation and analysis" *Proceedings of the XVIth International Congress of ISPRS*, Kyoto, Japan, *IAPRS* 27-A2 (Supplement)

Muller, J-P., Saksono, T. (1988)
"Fractal properties of terrain" *Proceedings of the XVIth International Congress of ISPRS*, Kyoto, Japan, *IAPRS* 27-A3 (Supplement)

Nishihara, H.K. (1984)
"Practical real-time imaging stereo matcher" *Optical Engineering*, vol.23(5), pp. 536-545

Nishihara, H.K. (1987) Private communication

Ohta, Y., Kanade, T. (1985)
"Stereo by intra- and inter-scanline search using dynamic programming" *IEEE Trans. on Patt. Anal. Mach. Intell.*, vol. PAMI-7(2), pp.139-154

Ohta, Y., Takano, K., Ikeda, K. (1987)
"A highspeed stereo matching system based on dynamic programming" *Proc. First Int. Conf. on Computer Vision (ICCV87)*, held in London, U.K., June 8-11, pp.335-342

O'Neill, M.A., Dowman, I.J. (1988)
"The generation of epipolar synthetic stereo mates for SPOT images using a DEM" *Proceedings of the XVIth International Congress of ISPRS*, Kyoto, Japan, *IAPRS* 27-A1

Otto, G.P. (1988)
"Rectification of SPOT data for stereo image matching" *Proceedings of the XVIth International Congress of ISPRS*, Kyoto, Japan, *IAPRS* 27-A3

Otto, G.P., Chau, T.C. (1988)
"A 'region-growing' algorithm for matching of terrain images" In preparation.

Pollard, S.B., Mayhew, J.E.W., Frisby, J.P. (1985)
"PMF: A stereo correspondence algorithm using a disparity gradient limit" *Perception*, vol.14, pp.449-470

Pollard, S.B. (1987)
"The PMF stereo algorithm : theory and implementation"
Proc. BCS Parallel Processing Group Workshop on "Parallel architectures and computer vision", Oxford, UK, 30-31 March 1987, 10pp

Quam,L.H. (1984)
"Hierarchical Warp Stereo" *Proc. Image Understanding Workshop*, DARPA (New Orleans, Oct. 1984), pp.80-86

Ramapriyan, H.K., Strong, J.P., Hung, Y., Murray, C.W. (1986)
"Automated matching of pairs of SIR-B images for elevation mapping" *IEEE Trans. Geosci. Rem. Sens.*, vol.GE-24(4), pp.462-472

Recce, M., Collins, K., Chau, T., Otto, G.P. (1987)
"Speed assessment of MMI-137 First pass stereo algorithms" *MMI-137(UCL-CS) Internal Report*, 27pp

Reddaway, S. (1973)
"DAP- a distributed array processor" in *Proc. of First Annual Symposium on Computer Architectures*, Florida, USA, pp.61-65

Reddaway, S. (1987)
"Mapping images onto processor array hardware" *Proc. BCS Parallel Processing Group Workshop* on "Parallel architectures and computer vision", Oxford, UK, 30-31 March 1987, 14pp

Simard, R., Krishna, V.G. (1983)
"A successful approach in three-dimensional perception of stereo LANDSAT-MSS images over Cordilleran relief" *Proc. 9th Int. Symp. on Machine Processing of Remotely-sensed data*, Purdue University, IN, USA, June 21-23, pp.31-40

Simard, R., Toutin, Th., Leclerc, A., Haja, S.R., Allam, M., Boudreau, R., Slaney, R. (1987)
"Digital Terrain Modelling with SPOT data and geological applications" *Proc. Int. Conf.* on "SPOT1: Image utilisation, assessment, results", held in Paris, France, 23-27 November 1987, 8pp

Shirai, Y. (1987)
"Three-dimensional Computer Vision" Springer-Verlag, London ISBN 0-387-15119-2

Stevens, A., Morris, A.C., Ibbs, T.J., Jackson, M.J., Muller, J-P (1988)
"Automatic generation of image ground control features from a digital map database" *Proceedings of the XVIth International Congress of ISPRS,* Kyoto, Japan, *IAPRS* 27-A2

Stoer, J., Bulirsch, R. (1980)
"Introduction to Numerical Analysis" Springer-Verlay, New York, 1980

Vincent, 1987
"Automatic extraction of high elevation resolution data from SPOT stereo images" *Proc. Int. Conf.* on "SPOT1: Image utilisation, assessment, results", held in Paris, France, 23-27 November 1987

Zeiss (1987) "Indosurf" Sales literature.