A Multi-Resolution Approach to Parallel Stereo Matching of Airborne Imagery

M. J. Zemerly, M. Holden and J-P. Muller

Department of Photogrammetry and Surveying
University College London
Gower Street
London WC1E 6BT
INTERNET: jamal@ps.ucl.ac.uk

Commission II

Abstract

A geometrically parallel area-based stereo matcher (GPSM) has been developed at UCL, originally for SPOT satellite imagery. It uses an adaptive least squares correlation (ALSC) combined with a region growing control strategy. This matcher works well for small-scale imagery, but has difficulty matching aerial images due to their higher resolution (0.6m-2m/pixel). At this resolution large depth discontinuities and low and periodic texture are present due to man-made structures. These can cause the stereo matcher to converge to false matches and impede region growing leading to a reduction in match coverage, this makes it unsatisfactory for the production of dense Digital Elevation Models (DEMs). A multi-resolution (MR) approach to stereo matching was investigated to try to mitigate against the effect of discontinuities by reducing their resolution to that encountered in SPOT imagery. As the first step in the investigation, a stereo image pyramid was formed with the lower resolution pairs matched serially and the original resolution ones matched in parallel on a MIMD architecture (a Transputer array). The results obtained with 5 levels of resolution demonstrated an increase in match coverage of 16% in comparison with 1 level and reduced apparent blundering.

KEYWORDS: Adaptive Least Squares Correlation, Area-Based Matching, Multi-Resolution, General Purpose MIMD, Geometric Parallelism, Aerial Imagery.

1 INTRODUCTION

The Alvey project MMI-137 (Real-time 2.5 Vision Systems) [Muller, 89][1] was concerned with the development of an automatic system for extracting 3D coordinates from 2D stereo images so that dense range maps could be generated from remote passively sensed stereo pairs. This system uses a stereo matching method which is based on Adaptive Least Squares Correlation (ALSC) [Gruen, 85]. It has been successfully applied to derive DEMs from SPOT satellite stereo imagery with typical accuracies of 4-10m rms height (for 10m pixels from 800Km altitude) [Day and Muller, 89]. Airborne stereo imagery has been used for a number of applications such as environmental monitoring [Muller et al. 91] and automated cartography [Hsieh et al. 90]. There are also a number of applications for DEMs derived from stereo matched aerial photographs apart from the obvious ones associated with mapping applications. These include the use of DEMs in orthoimage formation and their subsequent use in visualisation [Muller et al. 88b] and geometric correction of airborne scanner data [Allison et al. 91, Allison and Muller, 92]. Compared to SPOT, aerial imagery proved to be much more difficult to match because of its higher resolution (0.6-2m/pixel) hence low and periodic texture (e.g. fields of crops), and man-made structures with depth discontinuities (e.g. buildings, hedges, roads). These can lead to reduced match coverage, increased false convergence (blunder) rate and in some cases complete failure to sheet grow especially where severe depth discontinuities are present [Muller et al. 91, Hamid and Muller, 91]. This work is particularly concerned with applying automated processing techniques of aircraft data for monitoring environmental change.

By far the most time consuming process in the production of DEMs is the stereo matching of conjugate points within stereo imagery. Because large amounts of data, intensive floating point calculations and high-level programming techniques are needed, SIMD type fine-grain machines are not suitable. Instead, a coarse-grain MIMD machine, in particular a Transputer array, was chosen as the exemplar machine for the implementation. In particular, a Supernode-based machine consisting of 48 worker Transputers (T800) each with 4MB RAM. The stereo matching software was first targeted to workstations using C with UNIX system calls. To ease the porting of this software to Transputers a POSIX conformant UNIX-like operating system IDRIS is used. The parallel stereo matcher is able to process a pair of SPOT images of size 6000×6000 pixels with 48 processors in around 3 hours. Extensive performance measurements and a quality assessment were carried out on this matcher for SPOT images (see [Holden et al. 92]). In this paper a parallel multi-resolution stereo matching technique is applied to aerial imagery. Results of comparison with serial matching at the original resolution level are presented.

[1] UCL 3D Image Maker won the 1990 British Computer Society award for technical innovation.

2 STEREO MATCHING

The stereo matching of conjugate points is an important step in the automatic extraction of 3D coordinates from planetary and industrial surfaces. Given a pair of stereo images of the same area of terrain, the aim is to find the disparity (change in pixel position between the two viewing points) at each point in a regularly spaced grid, known as a Digital Disparity Model (DDM). This can be transformed using a model of the camera geometry to ground height (elevation) [Muller, 89]. Disparities are found by stereo matching i.e. the correlation of conjugate points in a pair of images. The requirement is to obtain a dense disparity map (set of matched coordinates) for the whole scene accurately, quickly, reliably and without the need for the inclusion of information on the camera exterior orientation. To achieve these goals we must select a single or hybrid form of stereo matching algorithm [Muller et al. 88a] from the two main categories: feature-based and area-based. Area-based matchers use pixel data directly whereas feature-based matchers use features such as points and edges extracted from the scene and discard the remaining data. An area-based matcher (Otto-Chau) [Gruen, 85, Otto and Chau, 88] rather than feature-based (e.g. [Foerstner, 86], [Grimson, 85], [Nishihara, 84], [Ohta and Kanade, 85], [Pollard et al. 85]) is used here because it produces much more accurate and dense disparity field from stereo imagery [Day and Muller, 89]. However, the Otto-Chau matcher requires some initial estimates (seed points). These are provided automatically by a feature-based matcher which uses the Foerstner interest operator [Foerstner and Gulch, 87] and ALSC [Gruen, 85].

2.1 Otto-Chau Region Growing Stereo Matcher

The stereo matcher was initially designed as a sequential algorithm (coded in C) to run on Sun workstations under SunOS [Otto and Chau, 88]. Recently, it has been ported to Silicon Graphics and IBM workstations. It correlates a pair of square patches (typically 15×15 pixels in the left image), one from each image. For this it uses a numerical technique known as Adaptive Least Squares Correlation, ALSC [Gruen, 85]. An unique feature of this technique is that it returns an estimate of the "goodness" or "precision" of the match and a set of parameters describing the spatial distortion which is approximated to first order by an affine transformation. Given a left hand patch the transformation maps to a deformed patch in the right hand image. An estimate of how well the two patches match is derived from radiometric least squares correlation. This estimate is used to calculate an incremental change in the parameters of the affine transformation. A small adjustment is introduced causing re-sampling of the right hand patch during each iteration. The algorithm proceeds until either a good quality match is found or a pre-specified number of iterations are executed. The convergence radius of disparity of ALSC is about 2.5 pixels. Because the disparity surface is almost always continuous a technique known as "region" or "sheet" growing [Otto and Chau, 88] can be applied. To initiate this process a small number of known matches "seed points" are needed (see [Allison et al. 91]). These are matched and placed into a queue in a "best first" order, then their neighbours are selected, matched and added to the queue. This procedure continues until the entire image is matched.

2.2 Geometric Parallelism

The Otto-Chau stereo matcher involves intensive floating point matrix operations, typically 3×10^4 arithmetic operations per iteration per pixel [Otto and Chau, 88]. Because a single pixel is used in many patch matches there are as many as 10^5 arithmetic operations per pixel [Otto and Chau, 88]. Stereo matching images of size 6000×6000 pixels, with grid spacing of 5 pixels, has been reported [Holden et al. 92] to take up to 6 days on a Unix workstation hence the need for a faster parallel solution. A geometric form of data parallelism was chosen as the best approach for stereo matching large data sets [Holden et al. 92] with existing limitations of the hardware and operating system. The strategy of geometric parallelism is to distribute image data and process it locally ("divide and conquer"). Imagery is sub-divided into rectangular tiles by superimposing a regular grid on the left hand image (TLx, figure 1) and calculating a corresponding irregular one in the right (TRx, figure 1) using the seed points. The parent process ("gpsm", figure 1) runs on the root processor under a full IDRIS v4.2 kernel and each child process ("smx" in figure 1) is loaded onto its own worker processor. No inter-worker communication is currently implemented. The parent process has been designed to invoke a pre-specified number of child processes. This number, the number of tiles in the x direction and the number of pixels overlap in x and y are specified as command line arguments by the user. These parameters are used to calculate the coordinates of the left hand tiles that are passed to the child process. The child process uses them to calculates the coordinates of the corresponding right hand tile. Once both sets of coordinates are determined the matching phase begins: image lines are read on demand so that reading image data is interleaved with matching, thus spreading the disk load. The left image was subdivided into equal sized

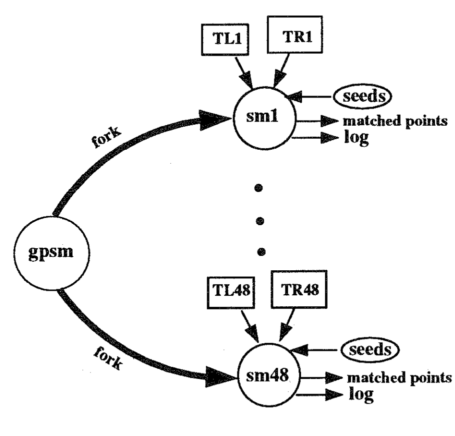

Figure 1: Process Organisation

tiles (6 in x direction and 8 in y direction, see figure 2). Figure 2 shows the left hand tiles (without overlap), their corresponding right hand tile boundaries and the automatically generated seed points [Allison et al. 91] superimposed on the original Gedney Hill images (see [Muller et al. 91] for detailed description of the data sets). Tiles must overlap by at least a patch diameter otherwise there will be an unmatched band around the tile perimeter. The selection of the corresponding right hand tiles was done during the parallel stereo matching as follows:

1. The left image is subdivided into equal sized tiles. These tiles are allowed to overlap by a user-specified amount (30-100 pixels have been used, see later).

2. Then, for each tile, the minimum and maximum disparities between the left and right image in x and y directions are computed from the seed points which lie inside the left tile.

3. The disparities are used to calculate the coordinates of the right hand tiles. These are then adjusted to allow for overlap.

The "best first" region growing strategy has a direct impact on the parallelisation of the algorithm. In the serial case, the next pair of patches are selected on the basis of the best quality match so far over the entire image space. In geometric parallelism only a small subset of the image space (pair of tiles) is available, so this is compromised to local best first. Comparison of the coverage and quality of the DEMs created using results from both the serial and the parallel version of the stereo matchers for SPOT satellite images showed insignificant differences (1%) with speedup of about 40 times with 48 T800 Transputers (see [Holden et al. 91a, Holden et al. 91b, Holden et al. 92]).

2.3 Applicability of Multi-Resolution

Also known as multi-scale, coarse-to-fine or pyramidal control strategies, multi-resolution techniques are applicable to a wide range of image processing applications. Because features within an image may only be significant at a particular resolution, there is no *a priori* basis for selecting an appropriate resolution. For this reason the analysis of multi-resolution image representations is a powerful tool for the examination of image features at multiple resolutions. In particular the use of coarse-to-fine search strategies for stereo matching first at a coarse resolution and then using these match estimates to converge to better estimates at a finer resolution. Figure 3 illustrates the data representation model used for coarse-to-fine control strategies.

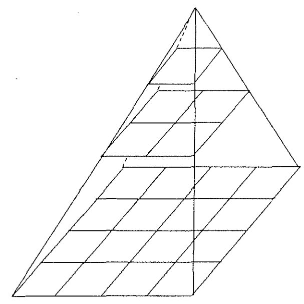

Figure 3: Three Level Pyramidal Representation

2.3.1 Pyramidal Representations

An image is decomposed into a set of representations, each one containing features at a different resolution. These are often referred to as image pyramids. Significant effort has been put into creating and analysing these representations, two widely used representations are:

- Gaussian Pyramids: Intensity values are smoothed or blurred over successively larger areas, so reducing the effects of noise. Linear smoothing by convolution is like applying a low pass filter. The size of the filter is determined by a Gaussian point spread function which takes the form:

$$G(x,y) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}} \qquad (1)$$

  where σ is a scale parameter

- Laplacian Pyramids: These can be formed by applying the Laplacian operator to a Guassian filtered image (i.e. LoG). This is equivalent to differentiating the smoothed intensity values to produce a set of band-pass filtered copies of the original. The LoG operator can be approximated by the difference between two Gaussians (DoG). Zero crossing pyramids can be constructed by extracting the points where the Laplacian changes sign [Dyer, 87].

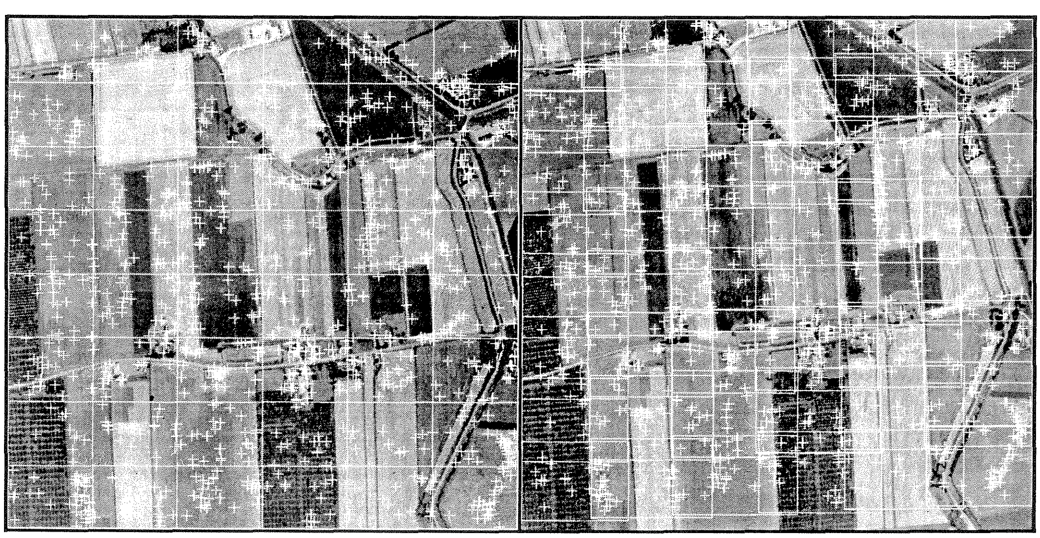

Figure 2: Gedney Hill Images with Tiles and Seedpoints

2.3.2 Multi-resolution Stereo Matching

A multi-resolution template matching algorithm is described in [Rosenfeld and VanderBrug, 77]. It starts by matching a coarse template against a coarse search image. This reduces the search space, and proceeds to a finer resolution only when there is a small degree of mismatch. Displacement estimates of pixels at the highest (coarsest) level can be used as initial estimates for calculating disparities at subsequent lower levels of the pyramid. [Chang and Chatterjee, 90] use a multi-resolution approach to improve stereo matching based on simulated annealing. It speeds up the convergence. Starting at the topmost level optimal solutions can easily be found due to the small search window. A preliminary study into applying multi-resolution for the matching of SEASAT SAR images was reported by [Denos, 91]. This is based on random initial disparity estimates at the top level.

In general disparity estimates are propagated to the next (lower) level as an initial state, and used to constrain the search space and substantially accelerate convergence. This is called constraint propagation between layers of the pyramid and can lead to a substantial reduction in the number of operations when compared with a single layer approach. The reduction is particularly relevant when large images need to be processed and therefore has good potential for real-time matching. Figure 4 illustrates the principle of our multi-resolution stereo matching; The automatically generated seed points from the finest resolution are scaled appropriately and fed to the top level of the pyramid. The disparity estimates (DDMs in figure 4) generated by stereo matching are scaled by a factor of 2 and propagated down each level of the pyramid. The seed points are found at the original resolution and scaled appropriately for each level.

Multi-resolution techniques are easy to implement in a parallel architecture. A simple technique is to dedicate a layer of processors for each resolution level. See [Williams and Anandan, 86] who has applied it to the Marr-Poggio-Grimson matcher.

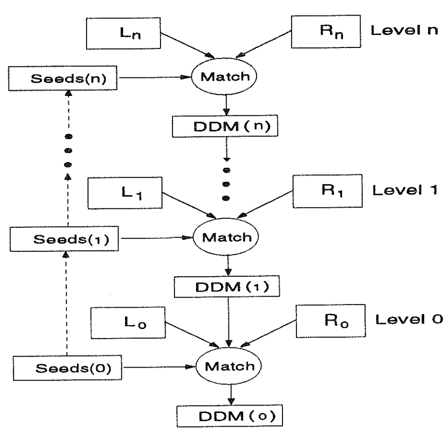

Figure 4: Multi-Resolution Stereo Matching

2.3.3 Evaluation of Multi-Resolution Stereo Matching

Although the stereo matcher was originally intended for use with satellite imagery it was designed as a general purpose algorithm capable of matching image data containing large amounts of high frequency texture. However, tests carried out on aerial imagery showed reduced match coverage coupled with blundering [Muller et al. 91]. A 2048 x 2048 scene of Gedney Hill (digitised at a resolution of 1m/pixel) has been used as a test case. This imagery has the following attributes:

- A large number of low or periodic texture areas such as fields of crops and tilled/ploughed soil.
- Small variation in height.
- A number of man-made structures: roads, buildings, etc. containing depth and radiometric discontinuities.

These pose problems to the stereo matcher and hence this scene was selected as the first test data. The first test carried out on this imagery with the serial version (with grid spacing of 3 pixels) resulted in a match coverage of only 70% (see figure 5). Tests with the parallel stereo matcher with a tile overlap of 30 pixels showed an even greater reduction of coverage down to 59%. This is mainly due to the partitioning of data. Increasing the overlap to 100 (at the expense of extra storage space and longer compute time) improved the match coverage to 67% so that any major differences between the serial and parallel versions were observed only around the image perimeter (see figure 6). Further inspection of figures 5 and 6 showed that some quite large areas (crop fields) were completely unmatched. This was considered unsatisfactory for the production of orthoimages.

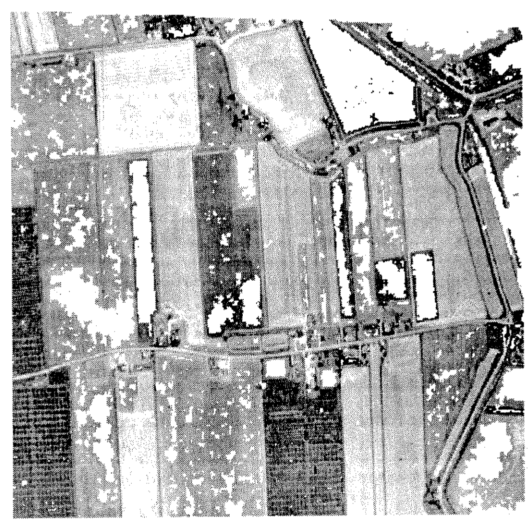

Figure 5: Coverage of the Serial Stereo Matcher for Gedney Hill left image (70%, white=no match)

Further tests were carried out to try to improve the match coverage. It was observed that the areas deficient in match coverage contained only small amounts of texture and it was supposed that this and the large number of discontinuities present in the imagery were the main factors contributing to the reduction in match coverage.

In order to further investigate this issue it was decided to evaluate the multi-resolution approach to stereo matching. The reasoning behind this is that at lower resolutions the discontinuities will be smoothed out and areas of low texture will be mapped to areas of higher texture. The matches at the top level (lowest resolution) can then be propagated to the next level and used as seed points to initiate region growing matching. A series of experiments were carried out to assess the merits of this approach. The image pyramid was formed by reducing the resolution by a factor of 2 at each stage. This was achieved by taking the mean of 4 neighbouring pixels at one level to represent 1 pixel at the next. The pyramidal matching is currently implemented using a Unix Bourne shell script so that it can be ported to IDRIS on the PARSYS Supernode in order to be parallelised. It was thought that the number of levels of reduction should be limited by the size of the images that could be practicably matched, i.e. 128 x 128 pixels.

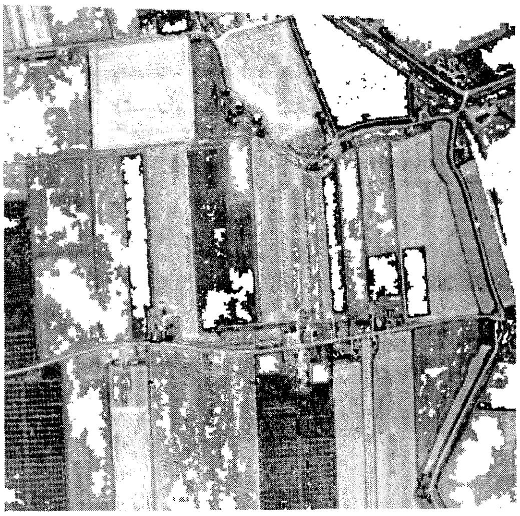

Figure 6: Coverage of the GPSM for Gedney Hill left image with 100 $\times$ 100 overlap (67%)

3 PARALLEL MULTI-RESOLUTION STEREO

In order to quickly evaluate the feasibility of this technique it was reasoned that the most difficulty would occur between the two highest resolution levels. Therefore the pilot study focussed on matching all low resolution stages on a Unix workstation and the highest resolution (no reduction) on the PARSYS machine. The matching of Gedney Hill images used an overlap of 75 $\times$ 75 pixels with 2 pixel grid spacing. The coverage (86%) is presented in figure 7.

Further improvement of memory allocation of GPSM is being considered. The present version allocates memory for an array to store the seed points first to calculate the minimum and maximum disparities needed for the redimensioning of the right hand tiles. The memory allocated for this array is freed only after all the seed points are matched and stored in the priority queue. For normal stereo matching (no coarse-to-fine) this array is small since the number of seed points is small. In the coarse-to-fine method thousands of matched points are generated at the top level and propagated down as seed points so the size of this array becomes significant. The stereo matcher needs some modifications to more efficiently deal with this. One solution is to use the space allocated to the priority queue initially to store the seed points, do the required calculation, match them and then re-order them best first. This will save the space allocated to the seeds array.

Figure 7: Coverage of the Multi-Resolution GPSM for Gedney Hill left image (86%)

3.1 Seed Points Generation

The aerial images used here were obtained by digitising aerial photographs on a Leica Kern DSR-11 analytical plotter. The Foerstner interest operator [Foerstner and Gulch, 87] was used to get a user-specified number of interest points for each image. Because the parallel stereo matcher requires at least one seed point in each tile, the Foerstner operator was forced to divide each image into sub-images and find sufficient interest points in each of them. The information obtained from the inner and relative orientation models set up on the DSR-11 and the maximum and the minimum expected heights within the scene were used to calculate the epipolar disparity limit in the right image where the corresponding point should lie. These heights can also be estimated automatically from orientation information obtained from the analytical plotter. A small error bound of ± 3 pixels on the original estimate is also allowed to cater for any imperfection. Any right interest points found inside this area are then listed against the left interest point for later stereo matching in order to select the correct conjugate pair by taking the best match. For further details see [Allison et al. 91].

4 RESULTS

Preliminary experiments were carried out on both a Unix workstation and a Transputer array for the Gedney Hill imagery. Image pyramids were formed using a local mean operator. In the first experiment 4 levels are used and the results are shown in Table 1. Although the coverage was

Reduction	Image size	Grid	Coverage	matches
8	256 $\times$ 256	1	69%	45359
4	512 $\times$ 512	1	83%	218244
2	1024 $\times$ 1024	2	87%	227670
1	2048 $\times$ 2048	3	86%	399017

Table 1: Results obtained from Multi-Resolution matching with 4 levels (patch-radius=7)

better than only one level (no coarse-to-fine) by 16%, it was noticed that in one of the uniform crop fields, it was still not sufficient. This was traced back to the lack of coverage in the lowest resolution (256×256), so another reduction level, i.e. 16, (image size of 128×128) was introduced with grid 1 in the anticipation that this would improve the coverage. This extra level gave little improvement. However, it indicated serious blunders were occurring at this level and were propagated to finer levels. In light of this, effort was then directed towards reducing this blunder rate. Tests were carried out using different patch radii at the top levels to study its effect. The blunder rate was studied visually at first by a rough estimation of the area of visual and significant blunders relative to the DEM area. Later a few Digital Disparity Models (DDMs) were created from some of the stereo matching results and these led to the same conclusions, in the areas where obvious blunders occur there are abrupt changes in disparity in the DDMs (shown as abrupt intensity changes in figure 8). Note that the DDM shown in the figure is histogram-normalised to highlight the blunder areas. The results of these tests for the 128×128 images (grid spacing=1), where the DDM grid contained 16384 points, are shown in Table 2 and suggested a patch radius of 5. Unfortunately there is no reliable automatic detector for quantitative estimation of the blunder rate so this was done by qualitative estimation of the percentage of blunder areas. This was done by noting either the holes matched only in one image and not matched in the other, or the abrupt changes in disparity found in the DDM. For 256×256 pixel images (grid spacing=1), the DEM contained 65536 points, the results are shown in Table 3 (best patch radius = 6). Similar tests were also carried out for 512×512 and 1024×1024 images and these showed that a patch radius of 7 for those image sizes performed the best in terms of apparent blunder rate and was used for these and the original (highest resolution) images.

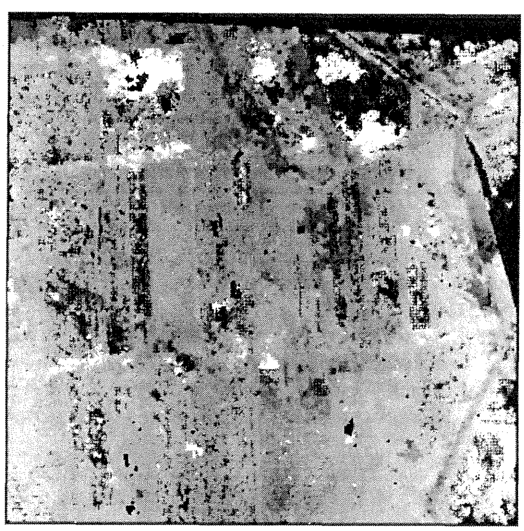

Figure 8: Histogram Normalised DDM Created from Multi-Resolution GPSM for the 2048×2048 Gedney Hill Images (large disparity=bright, low disparity=dark)

These results showed that although the coverage increased with patch size so did the blunder rate, and this was propagated down to the higher resolution levels. It was decided

Patch Radius	Obvious blunders	Coverage	matches
4	MODERATE $\approx 2\%$	46%	7560
5	OK $\approx 1\%$	55%	8947
6	MODERATE $\approx 3\%$	59%	9637
7	BAD $\approx 4\%$	62%	10151
8	V. BAD $\approx 5\%$	63%	10310
9	V. BAD $\approx 5\%$	63%	10292

Table 2: Effect of patch radius on blunder rate and coverage for 128×128 pixel images

Patch Radius	Obvious blunders	Coverage	matches
4	MODERATE	47%	30576
5	OK	57%	37186
6	OK	65%	42524
7	BAD	69%	45359
8	V. BAD	70%	46153

Table 3: Effect of patch radius on blunder rate and coverage for 256 images

that optimum coverage at the lowest possible blunder rate was required. Experiments have therefore focussed on identifying the optimum match parameters for optimum match quality. Another experiment with 5 levels 1, 2, 4, 8 and 16 and a suitable patch radius for each level (from the previous tests) gave slightly better results but had small blunder rate estimated at 1% to 2% of the total area of the highest resolution level (see figure 7). Table 4 shows the parameters used with each of the 5 levels as well as the coverage achieved.

Image size	Grid	Patch Radius	Coverage	matches
128×128	1	5	55%	8947
256×256	1	6	70%	45880
512×512	1	7	84%	219457
1024×1024	2	7	87%	227898
2048×2048	3	7	86%	400355

Table 4: Results obtained from coarse-to-fine matching with 5 levels

Other tests were carried out on 2229×2520 pixel Gwydyr Forest, Snowdonia (digitised at 0.8m/pixel resolution) imagery which is characterised both by large changes in terrain height and surface vegetation cover. The coverage of the parallel multi-resolution stereo matching is shown in figure 9.

The time spent matching the original resolution level of Gedney Hill and Gwydyr Forest images on the Supernode with Grid 2 was about 5.5 hours as compared to about 5 days on a Unix workstation. The current parallel version of the stereo matcher still has in general 1-5% less coverage when compared with the serial version if the texture is smooth near tile boundaries. This will be addressed using inter-process communications in IDRIS v5. The speed up achieved by parallelising the highest resolution level is only 24 times. This lower peformance is due to to the time spent reading the very large seed file which accounts for about half of the total time. Further substantial improvements in performance will be forthcoming with the T9000 Transputers which are 10 times faster than the current T800s.

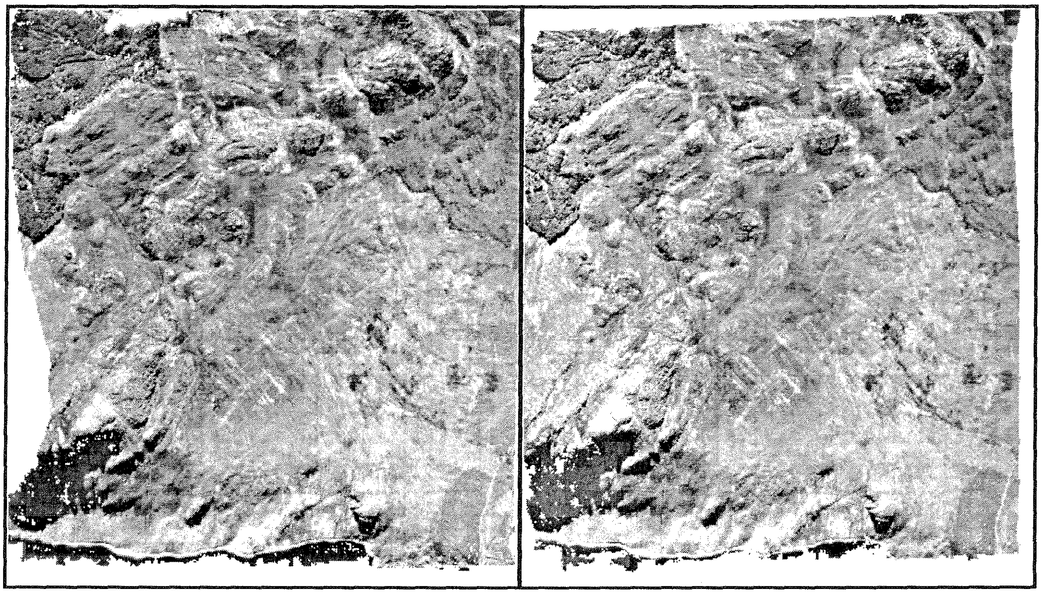

Figure 9: Coverage of the Multi-Resolution GPSM for the 2229 × 2520 pixel Gwydyr Forest left and right Images

5 CONCLUSION

This work has demonstrated the potential of multi-resolution approach when applied to stereo matching of aerial imagery with coarse grain MIMD systems based on Transputer elements to provide speed for a significant application area. During the course of this study a number of observations have been made:

1. Blunder detection - relation to lack of texture. The possibility of detecting blunders automatically needs to be further investigated.

2. Multi-resolution can lead to greater memory utilisation.

3. No Inter-Process Communications (IPC) between Transputers is used in the current parallel version. The availability of sockets in the next release of the operating system IDRIS, v5, will provide this capacity. This should improve the match coverage especially at the boundaries of the tiles, reduce the amount of overlap and save time in matching overlap. It also means that seeding every pair of tiles, also desirable, is not essential.

4. The best function to use when creating image pyramids (e.g. Gaussian, mean, etc.) should be investigated.

5. Matching all the pyramid levels in parallel ideally requires 2^{2N} Transputers for each level. It is probably worth considering increasing the number of Transputers to 64. Then, each process at the lowest resolution produces its own seed file. Only 4 processors need to read each seed file in the next level. This will save the need to search through one large file as at present. Thus the time needed to read the seeds could be cut to a fraction of this.

ACKNOWLEDGEMENTS

This research has been carried out under the Leverhulme Trust Project ("Monitoring Environmental Change Using Automated Processing of Aircraft Data"). Thanks go to David Allison for his help with photographic digitisation and seed point generation and to Mike Barnsley (UCL Geography) and NERC for providing the aerial stereo photographs.

REFERENCES

D. Allison, M. J. Zemerly, and J-P. Muller (1991). Automatic Seed Point Generation for Stereo Matching and Multi-Image Registration. In *Proc. Int. Geoscience and Remote Sensing Symposium (IGARSS 91)*, Volume 4, pages 2417–2421, Espoo, Helsinki, Finland, 3-6 June 1991.

D. Allison and J-P. Muller. (1992). An Automated System for Sub-Pixel Correction and Geocoding of Multi-Spectral and Multi-Look Aerial Imagery. In *Proc. XVII Congress, Int. Soc. for Photogrammetry and Remote Sensing*, Washington D.C., 2-14 August 1992.

C. Chang and S. Chatterjee (1990). Multiresolution Stereo- A Bayesian Approach. In *10th Int. Conf. on Pattern Recognition*, Volume B, pages 908–912, Atlantic City, New Jersey, 16-22 June 1990.

T. Day and J-P. Muller (1989). Quality Assessment of Digital Elevation Models Produced by Automatic Stereomatchers From SPOT Image Pairs. *Image and Vision Computing, AVC*, 7(2):95–101, 1989.

M. Denos (1991). An Automated Approach to Stereo Matching SEASAT Imagery. In *Proc. of the British Machine Vision Conf.*, pages 335–338, University of Glasgow, 24-26 September 1991.

C. L. Dyer (1987). Multiscale Image Understanding. In L. Uhr, editor, *Parallel Computer Vision*, pages 171–213, Academic Press, 1987.

W. Foerstner (1986). A Feature Based Correspondence Algorithm for Image Matching. *Int. Arch. of Photogrammetry and Remote Sensing*, 26(3):150–166, 1986.

W. Foerstner and E. Gulch (1987). A Fast Interest Operator for Detection and Precise Location of Distinct points, Corners and Centres of Circular Features. In *Proc. ISPRS Intercom. Conf. on Fast Processing of Photogrammetric Data*, pages 281–305, Interlaken, Switzerland, 2-4 June 1987.

W. E. Grimson (1985). Computational Experiments with a Feature Based Stereo Algorithm. *IEEE Trans. on pattern analysis and machine intelligence*, PAMI-7(1):17–34, 1985.

A. W. Gruen (1985). Adaptive Least Squares Correlation: A Powerful Image Matching Technique. *S. Afr. Journal of Photogrammetry Remote Sensing and Cartography*, 14(3):175–187, 1985.

A. Hamid and J-P. Muller (1991). 3-D CAD Modelling of Urban Scenes from Aerial Images. In *Proc ISPRS Int. Workshop on 3D in Remote Sensing and GIS: Systems and Applications*, Munich, 16-22 September 1991. Abstract only.

M. Holden, M. J. Zemerly, and J-P. Muller (1991a). Using a Transputer Array for Parallel Stereo Matching of SPOT Satellite Images. In T. S. Durrani et al, editor, *Proc of Transputer Applications*, Volume 1, pages 154–159, Glasgow, 28-30 August 1991.

M. Holden, M. J. Zemerly, and J-P. Muller (1991b). Applying Geometric Parallelism to Stereo Matching SPOT Satellite Images on A PARSYS Supernode. In *Proc. 6th Int. Conf. on Image Analysis and Processing*, Como, Italy, 8 pages, 4-6 September 1991.

M. Holden, M. J. Zemerly, and J-P. Muller (1992). Application of Transputers to Stereo Matching. In *Proc. European Workshops on Parallel Computing*, pages 40–51, Barcelona, Spain, 24-25 March 1992.

Y. C. Hsieh, F. Perlant, and D. M. McKeown (1990). Recovering 3D Information from Complex Aerial Imagery. In *10th Int. Conf. on Pattern Recognition*, Volume A, pages 136–146, Atlantic City, New Jersey, 16-22 June 1990.

J-P. Muller (1989). Real-Time Stereo Matching and Its Role in Future Mapping Systems. In *Surveying and Mapping 89*, University of Warwick, UK, 15 pages, 17-21 April 1989.

J-P. Muller, K. A. Collins, G. P. Otto, and J. B. Roberts (1988a). Stereo Matching Using Transputer Arrays. *Int. Archives of Photogrammetry and Remote Sensing*, 27(3):559–586, 1988.

J-P. Muller, J. Kolbus, M. Dalton, S. Richards, and J. C. Pearson (1988b). Visualisation of Topographic Data Using Video Animation. *Int. Archives of Photogrammetry and Remote Sensing*, 27(4):602–616, 1988.

J-P. Muller, M. J. Zemerly, and D. Allison (1991). Monitoring Environmental Change Using Automated Processing of Aircraft Data. Progress report I (1.10.89-31.3.91), Leverhulme Trust Project, August 1991.

H. K. Nishihara (1984). Practical Real-Time Imaging Stereo Matcher. *Optical Engineering*, 23(5):536–545, 1984.

Y. Ohta and T. Kanade (1985). Stereo by Intra- and Inter-Scanline Search Using Dynamic Programming. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, PAMI-7(2):139–154, 1985.

G.P. Otto and T.K. Chau (1988). A Region-Growing Algorithm for Matching of Terrain Images. *Image and Vision Computing*, 7(2):83–94, 1988.

S. B. Pollard, J. E. Mayhew, and J. P. Frisby (1985). PMF: A Stereo Correspondance Algorithm Using a Disparity Gradient Limit. *Perception*, 14:449–470, 1985.

A. Rosenfeld and G. J. VanderBrug (1977). Coarse-to-Fine Template Matching. *IEEE Trans. on Systems, Man and Cybernetics*, 7:104–107, 1977.

L. R. Williams and P. Anandan (1986). A Coarse-to-Fine Control Strategy for Stereo and Motion on a Mesh-Connected Computer. In *Proc. IEEE Computer Society Conf. on Computer Vision and Pattern Recognition*, pages 219–226, Miami Beach, Florida, 22-26 June 1986.